

TOWARD AN AUGMENTED FUTURE: MAPPING THE PROMISES,
CONTESTATIONS, AND SOCIO-TECHNICAL FUTURES PUSHING AN
EMERGING TECHNOLOGY

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

In Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Chung Li Liao

August 2014

© 2014 Chung Li Liao

TOWARD AN AUGMENTED FUTURE: MAPPING THE PROMISES,
CONTESTATIONS, AND SOCIO-TECHNICAL FUTURES PUSHING AN
EMERGING TECHNOLOGY

Chung Li Liao, Ph. D.

Cornell University 2014

This dissertation is an empirical examination of a group of people working on a set of technologies referred to as ‘augmented reality’ (AR). The dissertation examines who the stakeholder groups are, how they are conceptualizing and pushing the technology, and the future uses and applications they are advancing in order to build coalitions and make the technology a reality. Through participant observation, in-depth interviews, and document analysis, this work has followed the AR community for several years. The dissertation reports on five major areas of contestation: the first is a debate over the form that a commercial AR device should take, whether it is mobile or headworn. The second is the debate over what AR is, and the ways that the definitional debate helps understand the community itself and the tensions within it. The third is the relationship between a particularly strong stakeholder group, marketing and advertising, and the ways that other actors are adjusting and reacting to their influence. The fourth is a contestation over standards in this industry, and how certain groups are negotiating the standards that will enable certain uses of AR. The last debate is over the industrial AR space, as certain groups have been advocating a turn away from consumer AR to more industrial applications, and some of the implications for that move. This research is aimed at expanding our theoretical understanding of AR and the social, political, economic, and discursive dimensions surrounding its development, and what they tell us about how emerging technologies are formed.

The ways actors are envisioning technological innovation and development is contested through the structure of the community, who gets to be a part of the community, and the types of work that is valued within the community, all of which ultimately shapes the development and trajectory of the technology.

BIOGRAPHICAL SKETCH

Chung Li Liao earned two Bachelor of Arts degrees in International Relations and Political Science from the University of Southern California in 2007. He received his Master of Science degree in Communication in 2012 from Cornell University, after which he joined the doctoral program in Communication at Cornell University.

He has been the recipient of numerous honors and awards including the Glass Fellowship Award for Leadership and the College of Agriculture and Life Sciences Teaching Assistant of the Year Award.

He has also presented his work at major conferences such as the International Communication Association, National Communication Association, the Association of Internet Researchers, the ACM Conference on Human Factors in Computing Systems (SIGCHI), and the IEEE International Symposium on Mixed and Augmented Reality. His work has also appeared in major journals in the field such as *New Media and Society*, *Journal of Computer Mediated Communication*, and *First Monday*. Next fall he will be an Assistant Professor in the Department of Media Studies and Production at Temple University.

His dissertation, titled “Toward an Augmented Future: Mapping the Promises, Contestations, and Socio-Technical Futures Pushing an Emerging Technology,” was supervised by Dr. Tarleton Gillespie.

ACKNOWLEDGMENTS

When I was very young, my father earned his Ph.D. in physical therapy from New York University. At the time I was still too young to know what this meant except that these letters followed his name – Wen Shen Liao, Ph.D. To my father, the original Dr. Liao, it is only now that I know what it means to earn the title.

I want to dedicate this work to my mother, Li Yu Liao, who brought the family to the United States for our education and our future. For her unwavering dedication, sacrifice, and love, I only hope that I have made her proud.

To Grace and Vivian, who made sure I was never the smartest person in the room, like all older sisters should. I was always following their example, and they led the way even if they didn't realize it.

Special thanks goes to my committee, for all their support in making this possible. To my chair Professor Tarleton Gillespie, for first giving me the opportunity to study at Cornell, and for pushing me to ask the big questions. To Professor Lee Humphreys, for setting the example and showing me how to conduct research. To Professor Steven Jackson, for pushing me to make connections to broader literatures. To Professor Steve Hilgartner, for introducing me to many of the foundation literatures for understanding emerging technologies. From each of these people I have learned what it means to be a scholar, and I hope to emulate their mentorship style.

Lastly I also want to thank the New Media Society reading group, Josh Braun, Dmitry Epstein, Megan Halpern, Caroline Jack, Chris Langone, Elizabeth Newbury, Sandy Payette, Oya Rieger, and Leah Scolere, who read my work more than anyone and helped me think through the process of research.

This work was supported financially through grants from the National Science Foundation #1330297 and a Small Grant from the Institute for the Social Sciences.

TABLE OF CONTENTS

Biographical Sketch	iv
Acknowledgements	vi
Table of Contents	vi
List of Figures	vii
List of Abbreviations	ix
Chapter 1 – Introduction	1
Chapter 2 – Methods	29
Chapter 3 – Mobile versus Headworn AR	57
Chapter 4 – But is it AR?	124
Chapter 5 – Marketing, Retail, and AdvARTising	171
Chapter 6 – Positioning Before the Standards War	222
Chapter 7 – Industrial AR	271
Chapter 8 – Conclusion	297
Appendix	306
References	316

LIST OF FIGURES

Figure 1 – Four-Quadrant Map	7
Figure 2 – Sociotechnical Systems	15
Figure 3 – Communities of Practice	23
Figure 4 – AR Conference Map	36
Figure 5 – Interview Participants	40
Figure 6 – AR Adoption Curve	58
Figure 7 – Mobile AR System	62
Figure 8 – Metaio Augmented City	63
Figure 9 – Image from AWE2013	82
Figure 10 – GlassUp	86
Figure 11 – Using Mobile AR	87
Figure 12 – Facial Recognition and Social Networking	90
Figure 13 – Spectrum of Mobile Futures	92
Figure 14 – Spectrum of HWD Futures	95
Figure 15 – Google’s Project Tango	100
Figure 16 – Seebright	103
Figure 17 – Google Glass Advertisements	111
Figure 18 – MetaPro	115
Figure 19 – Mixed Reality Spectrum	129
Figure 20 - 1 st and Ten®	133
Figure 21 – User Centric Interactivity	136
Figure 22 – System Centric Interactivity	137

Figure 23 – Classification of AR Systems	165
Figure 24 - #RealityCube	166
Figure 25 – InsideAR Map	176
Figure 26 – Developer Investment in AR	179
Figure 27 – AR Software Developer Kits	181
Figure 28 – Cool Cola Orange	189
Figure 29 – Metaio Visual Search Demo	191
Figure 30 – Visual Representation of Marketing	196
Figure 31 – Augmented Hyperreality	204
Figure 32 – AR Ad Takeover	206
Figure 33 – Visual Representation of Marketing Opponents	211
Figure 34 – Band Aid for Kids	213
Figure 35 – IKEA Catalog	215
Figure 36 – Sterling Keynote	226
Figure 37 – Markup Language Comparison	234
Figure 38 – AR Interoperability	255
Figure 39 – Proposed Interchange Format	256
Figure 40 – ARML Specification	260
Figure 41 – Henderson and Feiner	276
Figure 42 – Vuzix & S&P	277
Figure 43 – AR Power Lines	280
Figure 44 – Vuzix M100	287
Figure 45 – Microvision Nomad	291

LIST OF ABBREVIATIONS

AMH – Arts, Media, and Humanities

AR – Augmented Reality

ARE – Augmented Reality Event

ARML – Augmented Reality Markup Language

ARNY – Augmented Reality New York

AWE – Augmented World Expo

AWENY – Augmented World Expo New York

CARIF – Common AR Interchange Format

CEO – Chief Executive Officer

CoP – Communities of Practice

CTO – Chief Technical Officer

HTML – Hypertext Markup Language

IARSC – International AR Standards Community

IEEE - Institute of Electrical and Electronics Engineers

ISMAR – International Symposium on Mixed and Augmented Reality

ISO – International Standards Organization

KML – Keyhole Markup Language

KARML – KML Augmented Reality Markup Language

Legitimate Peripheral Participation - LPP

OGC – Open Geospatial Consortium

S+T – Science and Technology

SCOT – Social Construction of Technology

SDO – Standard Development Organization

SSO – Standards Setting Organization

SWG – Standard Working Group

VC – Venture Capital

W3D – Web 3-D

W3C – World Wide Web Consortium

CHAPTER 1

INTRODUCTION

It was in 2009 that the application Layar application first launched, billed as one of the first augmented reality (AR) applications to be commercially available. In the first video advertising Layar, the scene opens with a panoramic view of city buildings with text proclaiming “this is reality.” It then shows a mobile device, running Layar, panning over the scene with the text now proclaiming “this is augmented reality.” As the mobile phone moves across the physical setting behind it, information begins to pop up about various locations it identifies in the scene. The advertisement showed that by moving a mobile device, one could get information about your surroundings such as which houses are for sale, the distance of the house, the price of the house, and photos of the house. Using the touchscreen, one can adjust the settings of the search, show an overhead map, and switch between themed layers.

The video advertisement was impressive, and raised many questions about AR’s potential as well as possible implications. First, the potential psychological implications of AR were unknown, as well as how people might perceive augmentations in relation to their physical surroundings. Second, these applications might alter people’s perceptions and how people communicate with one another. Third, the content shown was relatively uncontroversial, but the potential for many different kinds of AR content was there, depending on how existing media industries might adopt/integrate AR technologies. Some of these might subsequently raise complex legal issues, from privacy and property rights, where AR could be the next area where some of the content production and distribution issues raised by the internet were contested.

It quickly became clear, however, that the rush to study the content and consumption of emerging technologies was premature and would only offer an incomplete understanding of the technology. AR was (and arguably still is) in a stage similar to the internet in the 1990’s, a period

when ‘punditry rides rampant,’ or the optimistic celebration of its transformational potential peppered with dystopian prognostications from the skeptics (Wellman, 2004). It is in this moment that many decisions are made that push the technology forward, which ultimately inform an understanding of why the technology developed the way it did, what is possible with the technology, and the implications of those decisions. At the local AR meetings there is a small but enthusiastic group of supporters doing just that, trying to celebrate its remarkable innovation and technological breakthrough while warning about potential dire consequences. But their discussion was framed in terms of projections about AR, with an inflated certainty about the consequences that AR will have, the ways that it would change existing industries, and the content issues that the technology will raise. These promises flowed quickly, moving from the possibility of AR to certain timelines of when technological breakthroughs would happen and finally to its ‘impact’ on users. The tone was optimistic and declarative, with a strong sense of inevitability, social change, and technological determinism being celebrated and championed. This despite the fact that AR, like all media, is the combination and confluence of many technologies, each of which is subject to its own economic, social, and cultural conditions of use, which make it impossible to predict how it will develop let alone isolate any specific impacts or effects.

Instead of AR and its use being the object of study, it became clear that the research focus should be on the people who were trying to bring AR into being, the work they did at these meetings/conferences, the stakeholder groups they represented, their priorities and their interactions, the structural forces that brought these events and people together, the assumptions and expectations they had for the technology, the promises they were making, and the visions that they advanced that motivated them and others to action. Unpacking these entanglements and

understanding the community was a prerequisite to all of the other questions that might follow, because these stakeholders were making determinations about, designing for, and debating those very issues. More broadly, what kind of industrial relationships do emerging new media technologies bring into existence, and how – what coalitions form, what are the communicative practices that make the technology persuasive and push for certain developments, what strategic actions and decisions do they take, what factors and commitments guide these coalitions, and how does that ultimately shape the technological artifact?

The AR industry does not exist in isolation, rather it is built on a confluence of existing information and communication technology industries. Many of those stakeholders are moving into AR, with their own sets of market and economic logics, which makes the AR industry both a place to see how existing industries are interested in the technology as well as how AR actors respond and shape the technology to fit those demands. Two areas that I examine in more detail are marketing and advertising relationships in AR, as well as the standards creation process.

Beyond the shape of the industries themselves, it is important to examine the constructed and contingent nature of technologies, the stakeholder groups that work on them, the larger social, political, and contextual factors that drive technologies, and the interrelationships between these. Understanding the phenomena as it was happening proved to be a difficult task, particularly because AR existed in between and amidst projection and speculation, with many promises of possibilities and threats of peril. Often it is through these promises and expectations that people express their intentions and commitments. In the case of AR, mapping these future visions in real-time proved to be an important theoretical and methodological tool, as well as being one of the first inquiries into the futures of the burgeoning AR field. Another way of understanding the AR community, in particular their practices, rituals, and disputes, is to view

these actors as a Community of Practice (CoP), which offers a set of theoretical concepts for making sense of members' interactions and the underlying motivations for their actions.

While one of the exciting things about studying emerging technologies is that it is constantly changing, these types of studies are not predefined by any literature (unless there are naturally occurring experiments within the case), and what one might find in those settings is uncertain. This chapter highlights the ongoing discussions from a variety of fields, and draws from them analytical and methodological tools to help illuminate this specific case. Some of what happened in this case is unique to AR technologies or to certain stakeholders and the decisions they made, while some of it is the latest manifestation of longstanding pressures and trends in new media development. The empirical findings from these cases first make a contribution to our overall understanding of emerging technology, but also builds on these literatures and frameworks as an emerging case where these theories, concepts, trends, and relationships are observed, documented, and analyzed.

New Media Industries

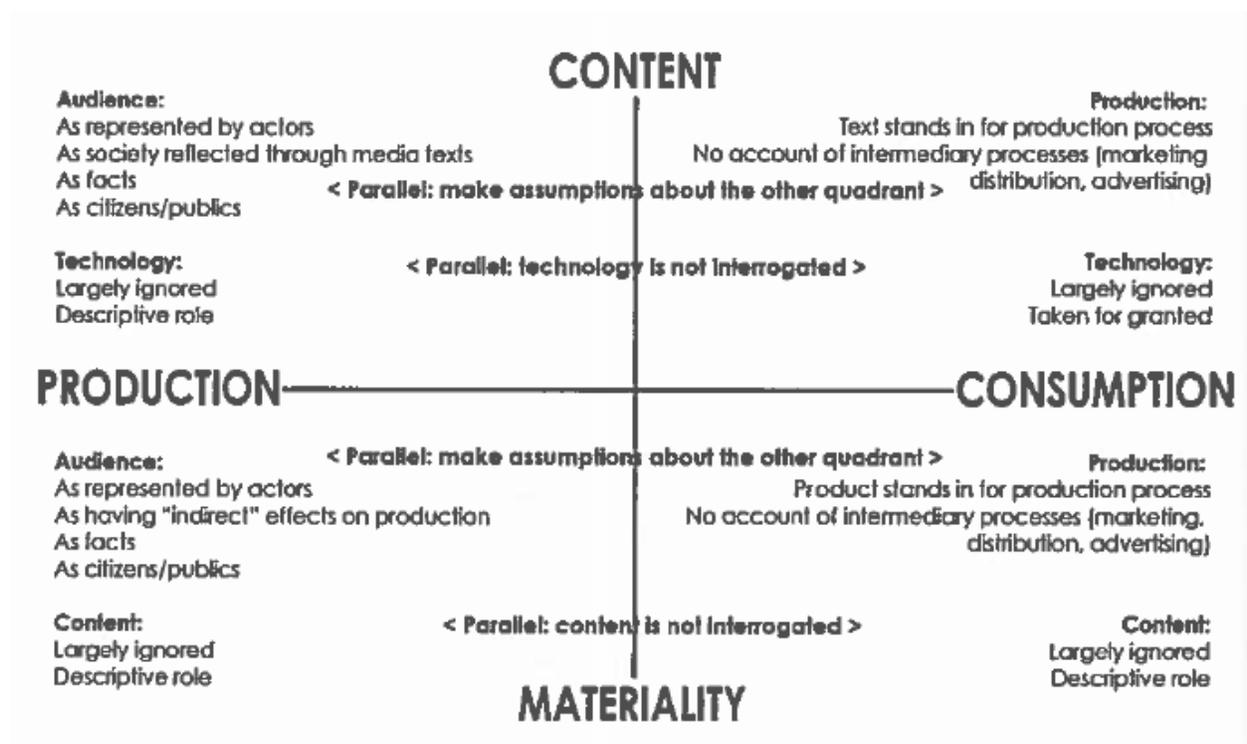
Lievrouw and Livingstone (2006) provide a helpful conception of what “new media” even is, one that goes beyond a classification of systems and features, to include “the artifacts or devices that enable and extend our abilities to communicate, the communication activities or practices we engage in to develop and use these devices, and the social arrangements or organizations that form around these devices and practices (p. 23).” This dissertation begins from Lievrouw and Livingstone’s perspective: that that these three elements are interrelated, but what is unknown is the precise relationship amongst the artifacts, practices, and social arrangements: “For communication research and related fields, the central question concerning technology

today is whether the particular configurations of artefacts, practices, and social arrangements differ, and in what ways, from those that characterized older information and communication technologies (Lievrouw & Livingstone, 2006; p 2).” The remaining chapters explore each of these three components, the design, practices, and social arrangements of the AR industry, and discuss those developments in relation to other communication technologies.

In order to do so, this study first has to understand the materiality and histories of the systems that gave rise to AR and made it possible. Focusing on new media as an area of study requires looking at these elements as inextricable and mutually determining, first in how they are the instrument and product of social shaping, and second their particular social consequences. In calling scholars to study new media, Silverstone (1999) reminds us that: “it needs to be remembered that media technologies emerge as material and symbolic objects and as catalysts for action, and are effective as such only through the deeds of individuals and institutions (p. 145).” Cultural historians of media technologies have worked to understand media technologies, media industries, and media uses as a holistic enterprise, one that is the product of economic, social, and cultural forces. In particular there are many similarities to Marvin’s (1988) work on electricity; in her story, as with AR, actors struggled to define the field and demonstrate their expertise, all while putting forth various social meanings and uses for the technology into society. At the moment of emergence, “the introduction of new media is a special historical occasion when patterns anchored in older media that have provided the stable currency of social exchange are re-examined, challenged, and defended” (Marvin 1988: p. 4). It is toward these early moments that scholars need to turn their attention, as “very little work either in the field of science and technology studies or in media studies has addressed the innovation process as a whole (Silverstone & Haddon, 1996; p. 70).” William’s (1990) work on television also turned its

focus onto the industry itself, examining the multiple meanings that television could have and the various pressures that altered the technology, the content, the genres, the form, and ultimately our experience with the technology. Other works that cast a historical eye on the development information technological systems such as the internet (Abbate, 1999), computing (Edwards, 1996), radio, film, and television (Wu, 2010), have examined the multiple ways that different social groups attempted to influence the technological, commercial, and ideological elements that would ultimately characterize those systems and the way they developed. New media development is comprised of a complex web of economics, human actors, technological actors, institutions, and social structures.

Because so much of the discussion surrounding technology still focuses on the innovation timelines, the technological advances, and the trajectories of technology that can seem natural/inevitable, it is these kinds of mappings that shed light on social, political, local, and temporal dynamics that might otherwise have been lost. These serve as not only historical and methodological analogues to this work, but also as a point of comparison to see how AR as an emerging field struggles with the same issues, faces similar pressures, and takes tangible actions to resolve and push past those issues. This is still a problem today, as Boczkowski and Siles (2014) further articulated the need for these kinds of studies in media technologies through their four-quadrant map of these relationships, as each quadrant makes assumptions and overlooks important factors (see Figure 1).



(Figure 1 – Four Quadrant Map, Boczkowski & Siles, 2014)

This dissertation is an attempt to bridge these divides, by looking at how actors produce the technology through a series of assumptions and discussions about the users, how that influences the content production possibilities for the technology, how the material technologies are adopted by early users, and finally the types of content that people engage with and how those are produced by certain intermediaries. Because of when this inquiry began (and the relative lack of users), it focuses primarily on the production side, but it does begin to consider how some of those decisions could affect consumption, while also opening up space for additional work in AR that relies on understanding how the production shapes users' ability to create and consume content (see, for example, Liao & Humphreys, 2014). Through this case, we begin to see the strategies that actors choose to engage in, as well as the implications of some of these moves. In doing so, we can begin to understand how AR might become an integral part of the media

environment, how it gets modified by economic, social, and cultural practice, and the range of potential social outcomes that result.

In looking at the production of existing new media industries, much of the recent focus has been primarily about how certain actors are adjusting to, capitalizing on, and pushing internet and digital technologies. One line of scholarship looks at how these networks of production have altered the structure of existing media industries and their business models (Boczkowski, 2004; Hesmondhalgh, 2002). In particular, even as the internet has allowed for more personalized access to information and content, the trend has been toward a consolidation of media industries and media companies, a convergence of many different forms of technology, and an emphasis on transmedia production (Hesmondhalgh, 2002; Jenkins, 2006). Advertisers have also been adjusting to these new possibilities, as traditional media may not reach the same viewership or market segments they used to (Turow, 2005). The option for advertising online has fundamentally changed the business models of many media businesses, while also providing businesses with a more efficient means of reaching consumers: “online advertising is disrupting all aspects of the global advertising industry [...] from how creative work is done, to how advertising campaigns are run, and to how advertising is bought and sold (Evans, 2009; p. 38).” These changes in the industry have been well documented, in particular the ways that the internet has altered the production possibilities and organization of global economic structures (Castells, 2000; Flew & McElhinney, 2006). These examinations of the political economy of new media and the dynamics within offer an important starting point for understanding the way that new media innovation is infused with new and shifting relationships of power (Mansell, 2004). The AR industry stands alongside these massive industries undergoing rapid transformation as one part curiosity and one part possibility, offering a potentially new way of accessing media. The

first movers in the space have been the marketing and advertising industry, which has important implications for the technology. Given that much media is funded through an advertising model, the presence of these stakeholders have the potential to push AR technology in particular directions, as well as provide the necessary resources for the industry to grow. Because of the rapid and dramatic changes in media structure around some of these new media possibilities (Spurgeon, 2008), stakeholders have begun to realize that embracing the ‘next next’ thing, the newest form of media, might ultimately keep them ahead of the competition and serve their best interests.

The current new media industries are characterized by particular organizational forms, economic resources and logics, as well as a certain culture. First, the consumer technology space began as a collective of small organizations: “Much of the nexus of businesses around new media is entrepreneurial, small or very small scale in terms of human capital, but creative and knowledge-driven (Cooke, 2006; p. 267).” Many of these are singularly focused on a particular product, application, or service, smaller, newer, with younger employees, more risky, and less structured (Marwick, 2012). Many of these startups and technology companies that are prominent in the space are clustered around creative problems or emerging technological innovation. Some of these startups have now become the large technology companies that play a major role in the technology landscape, like Google and Yahoo, through a series of organizational, managerial, personnel, and investment decisions (Levy, 2011). Their emphasis on acquiring smaller startups and trying to expand into new technology areas becomes an important part of this story, as well as the primary actors in the AR space, the economic logics that drive many of these companies, and the resources that are available to them. Often these companies are too risky for traditional investment sources to support, given the lack of organizational

structure and resources, relatively unknown executive teams, comparable lack of experience, experimental business plans, and high levels of competition. Venture Capital (VC) firms often fill this void for funding innovation, but requires the promise of much higher returns than other investments to attract private equity (Zider, 1998). In exchange, these funds help them grow and build the infrastructure, but in doing so they need to outgrow their competitors, obtain first mover advantage in the AR space, and attract developers to their platform. Sometimes the VC companies even take over executive positions or send over supervisors to help steer these projects in a particular direction. Through VC or not, these startups are attempting to strike it big either through an independent public share offering (IPO), or through acquisition by one of the larger technology companies (Crain, 2014). Lastly, these are the organizational and economic pressures in this space, but these are people working in these industries who have not only specific skills and expertise, but also certain cultural values. These startups are characterized by high levels of risk to both the company and the people working at these places, but that risk becomes normalized, internalized, and valorized as part of the entrepreneurial culture of venture labor (Neff, 2012). Because they bear much of the risk, this is not only easier for the economically privileged but also for those who believe they are responsible for their own skills, career development, and successes (Neff, Wissinger, & Zukin, 2005; Pratt, 2000). This culture of cool, the desire to chase the next big thing, and the cult of personality and social status/networking is a critical element of the emerging technology space (Marwick, 2013). These organizational structures, economic logics, and individual and collective values are not only present in the AR space, but are important for understanding both the people who are there and why they undertake certain actions. At the same time, this study of AR builds on this line of

research in understanding not just the practices and cultural values of this industry, but also how those get picked up by and enacted around new technologies.

Social Construction of Technology

There has been an ongoing discussion in the field of Science & Technology Studies about how people make technologies, the relationships that result, the factors that are most important in the process of innovation, how to understand the technological objects produced, and the ways they are shaped by social pressures. Scholarship on the production of technology has often been the domain of STS-inspired analysis of the political economy of these artifacts (Bijker, Hughes, & Pinch, 1987; Suchman, 2006; Williams & Edge, 1996).

The social constructivist position argues that the making of these technologies, the resolution of conflicts about their development, and the final understanding of them rests not on technical elements, but rather on the social meanings of these things by many stakeholders that force its development (Pinch & Bijker, 1987; Williams & Edge, 1996) To critique the one-dimensional, linear stories commonly told about technological development, “social shaping of technology” (SST) and “social construction of technology” (SCOT) posit multiple dimensions along which a technology evolves simultaneously and in different directions. For each artifact one could determine the site within which relevant social groups of actors “constituted” the object and created/defined its uses, a process defined as ‘interpretive flexibility’ (Pinch & Bijker, 1987). Each relevant social group with a stake in the artifact defines the problems it has encountered and proposes solutions for them, each operating within a set of ‘technological frames,’ defined as “the goals, key problems, current theories, rules of thumb, testing procedures, and exemplary artifacts that structure group members’ thinking, problem solving, strategy

formation, and design activities (Bijker, 1995; p. 125).” These frames promote certain facts and prescriptions over others, and invite actors to draw upon cultural and historical elements that may resonate with similar social actors.

SCOT is particularly instructive in cases where artifacts are forged in controversy, “where a social group or a set of groups imposes its solutions on other interested parties (Law, 1987; p. 112).” Scholars of SCOT also wrestled with questions of how all of these contestations gets settled into artifacts, a process described as ‘stabilization’ and ‘closure’ (Bijker, 1995; Pinch & Bijker, 1987). Closure is the last stage of stabilization, the point in which the relevant social groups see the problem as solved. There are multiple ways that these problems can be solved, either theoretically (through some combination of technology or discourse) or by shifting the problem. The manufacturers of the ‘easy’ bicycle, for example, were able to claim that their design technologically solved the problem of safety. Controversies are also not only technical but moral, as Bijker (1995) showed that manufacturers designed certain bicycles specifically for women wearing dresses. Their argument is that these stakeholders make every element of technology social: “it is claimed that there is nothing but the social: socially constructed natural phenomena, socially constructed artifacts, and so on (Pinch & Bijker, 1987; p. 109).” The concept of closure features less prominently in this account, because what the actors are still negotiating is the process for stabilizing meanings around the artifact and the problems have yet to be universally recognized as ‘solved.’ What this case does offer, however, are ways of understanding the strategies that actors undertake to stabilize certain meanings and designs – either by denying or redefining the problem, implementing policies to minimize social critiques, and attempting to modify users’ behaviors.

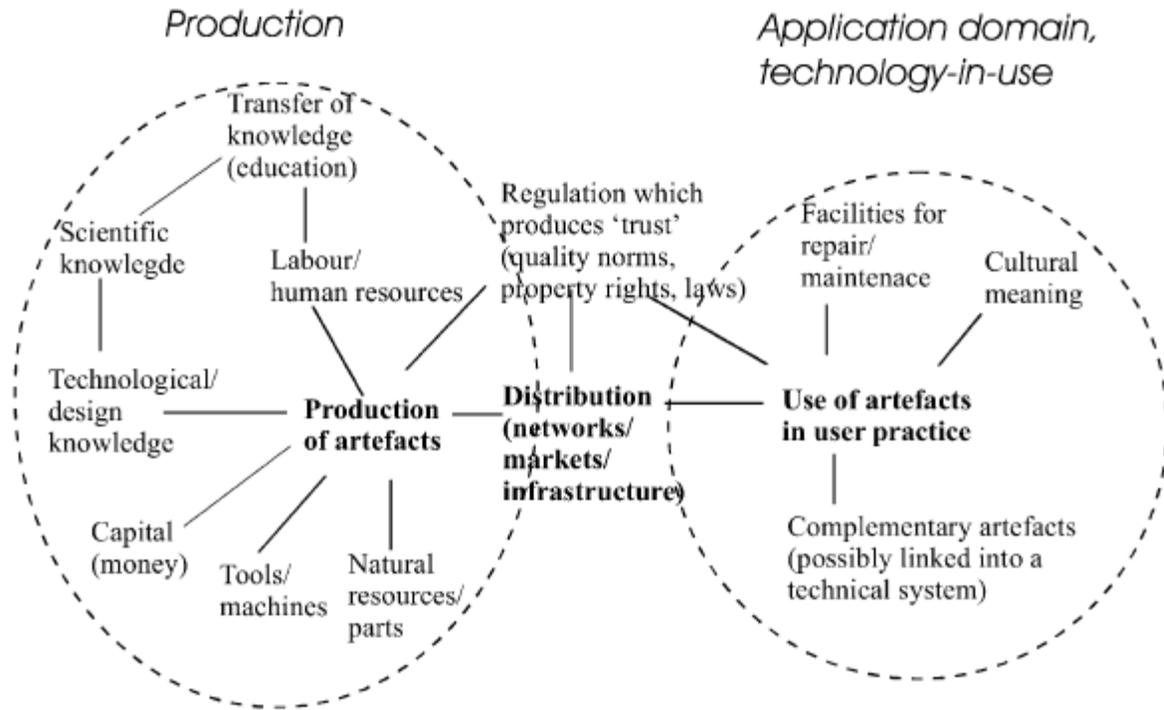
Some have challenged the implementation of SCOT, arguing that the notion of a relevant social group is vague and often amorphous (Flichy, 2007). While in some instances the producers and users are related or part of the same community, in other instances the break is much deeper, causing some to ask: “who says what are relevant social groups and social interests?” (Winner, 1993; p. 369) Others charge that SCOT does not do enough to understand the role of power between these actors, and the relative power disparities of particular social groups and their influence on the artifact (Russell, 1986). While it is important to consider both designers and users in the analysis, that should not conflate their degree of participation in the artifact and the ways that they engage with it – designers negotiate with the representations of users in the construction of the machine, while users negotiate with a machine that has certain prescriptions that they can deploy, appropriate, modify, or reject, each with its own resistance (Akrich, 1992; Flichy, 2007). Questions have also arisen within SCOT about the role of agency and structure in analytical interpretations of closure and stabilization, when we know those things are happening, and the extent to which those concepts exist (Klein & Kleinman, 2002). Recent works have refined these categories to address some of these critiques, showing how there are multiple ways these technologies can be renegotiated and reopened as they are taken up by different social groups, as well as how the identities of those groups are reconstituted in the process (Kline & Pinch, 1996; Oudshoorn & Pinch, 2003). Other work has linked SCOT to the broader political discourse and security policies of the day, examining how these artifacts contribute to the “technological construction of social worlds (Edwards, 1996; p. 34).”

The AR industry presents a case to explore the emerging technology innovation space using SCOT, in terms of the material elements and broader discourses that actors are trying to link to the technology. AR offers a unique case to explore these tensions, as it currently sits in a

transition phase between imagination and materiality, fiction and reality, and design and use. Using SCOT to understand the innovation process calls for an analysis of the relevant social groups, the problems and issues that are raised by an artifact, the ways that designers attempt to address those and imagine the users, the ways that users engage the artifact, and the relationships between these.

In some ways, historians have the benefit of knowing the outcome of a particular socio-technical arrangement, and the work comes from explaining and linking together certain events and occurrences which they deem to be explanatory (Flichy, 2007). Contestation over the models and theories of technological innovation that arise from these historical accounts may argue that certain perspectives and factors were overlooked, overemphasized, or taken as overly predictive. Real-time mappings do not have this knowledge, but are able to avoid the pitfalls of telling the story backwards, and obscuring the various contingencies and disagreements that gave rise to a particular technology. The various actors, stages, frames, systems, and boundaries get debated as analytical categories and for their explanatory and predictive potential. Unlike other pieces that have taken a historical angle to unpack these categories (which elevates the interpretive questions of whether designers were actually considering other relevant social groups as part of the negotiation, whether the social supersedes the material, whether certain factors were indeed the most important), this piece aims at an active, real-time mapping of the stakeholder groups, their social concerns, and the strategies and approaches they undertook to solve these.

A model of factors that contribute to contemporary socio-technical innovation from Geels (2004) shows the interrelationships between several production factors, distribution, and the use of artifacts (Figure 2)



(Figure 2 – The Basic Elements and Resources of Socio-Technical Systems, Geels, 2004)

These relationships are not static, and are increasingly distributed over a wider set of actors and greater geographic distances, such that these regimes must coordinate amongst actors, activities, and technologies. Intermediary places like industry conferences are where much of this coordination work gets done, from tangible business partnerships and standards setting, to more abstract work like coalescing around certain visions of the technology and teaching new members how to support and promote the technology. Geels (2004) explains that inter-organizational places are important to understand socio-technical innovation, as they are organized by members so that resources can be circulated (artifacts, knowledge, capital, labor, etc.) but also as a place where social groups come together around certain goals, roles, responsibilities, norms, and perceptions. It is in these places where we see actors engaging in much of the work that the SCOT approach believes is critical to technology creation; the appearance and coalescence of relevant stakeholders, the exchange of problem spaces and

solutions, and the ways these affect the material outcomes. These intermediary spaces are where networks form to facilitate the distribution, infrastructure, and human resources that ultimately shape the application domains and uses of the technology. This dissertation attempts to fill a gap in existing research on emerging technology, first to understand the critical transition between the production and distribution, and second at these intermediary places because “the background conditions of group interactions, such as their relation to each other, the rules ordering their interactions, and factors contributing to differences in their power, remain largely invisible (Klein & Kleinman, 2002; p. 30). Mapping these intermediary places, the conferences, meetings, and standards work of an industry, becomes a way of answering the question of who the relevant social groups are, who they believe is important, what they believe the problems are, who they attempt to account for in their design, and who they point to as indicators of progress. It is also in these moments that we can understand the interpretive flexibility of the technology, as people actively construct these artifacts, negotiate the form they will take, make assumptions about users, and build coalitions amongst groups in an effort to stabilize these technologies for distribution and use.

This dissertation builds on works that have recognized the need for early real-time mappings that follow the journey of an artifact (Boczkowski, 2010; Rip & Schot, 2002). Some of the retrospective work surrounding technology may overlook or overprivilege certain perspectives, not be in a position to capture certain strategic decisions made to avoid certain possibilities, and rely on outcomes as the guide for understanding the space (Rip & Schot, 2002). Instead, following the ‘journey’ of these technologies from the beginning “can contribute to identify critical moments, sites, and dynamics that mark the life cycle of media technologies (Boczkowski & Siles, 2014; p. 72). In these moments you find actors attempting to mobilize

resources to be able to work on the possibility, building coalitions with other actors in the chain, even tentatively introducing new products and processes that broaden the social actors that may exert pressure on the technology (Deuten, Rip, & Jelsma, 1997; Rip & Schot, 2002). In this nascent state, the technology exists in what Flichy (2007) defines as a ‘catchall phase,’ where a variety of actors were simply imagining new uses and technical utopias for the technology, given certain experimental possibilities. In this time there is a significant amount of indeterminacy, and a wide range of possibilities remain open.

This real-time mapping can also address the question of power, which some have argued is underdeveloped in applications of SCOT (Klein & Kleinman, 2002). Rather than have to speculate on the role of power, this mapping shows tangible instances where powerful groups push the technology in a particular direction, as well as how actors perceive and react to such efforts. The recognition that different relevant social groups have disproportionate authority over the artifact comes through in the movements that people make, how the technology shifts to accommodate those spaces, and how certain actors attempt to attract and align themselves with powerful entities. Some of the clearest examples reveal how the less powerful compensate for their lack of access to certain resources through design, or confront power by trying to build a coalition and appealing to other powerful institutions.

As the technological artifacts begin to materialize, the contestation manifests not only at the level of discourse but also at the level of technological design (Flichy, 2007). It is in these moments where “technological artifacts are sufficiently undetermined to allow for multiple possible designs, so whatever the design that finally results from the process, it could have been different (Klein & Kleinman, 2002; p. 29). Exploring this intersection helps reveal how these mediators come together to negotiate, articulate, and align specific technical choices and user

needs (Schot & de la Bruheze, 2003). It is also in this phase when some early designs are proffered as tangible technologies, making visible how users and uses are being presumed (Akrich, 1992). This dissertation is an attempt to bridge the constructivism versus materiality issue, over which perspective scholars should focus on to explain technology (Sterne, 2014), by focusing on this particular moment in between these phases. To fully understand what is happening, it is not enough to look only at the material design or only the social construction, rather we have to examine both how the discourse surrounding the technology gets embedded into the design of the technologies, and the ways the technological designs attempt to enter, affect, and mediate ongoing discussions. It is in these spaces, in its articulation and deployment, that AR as a sociocultural artifact is being actively imagined, negotiated, and constructed (Pfaffenberger, 1992). At various points, we can begin to see how certain areas of the social contestation corresponds to a partial understanding of the design of the technology.

Futures, Promises, and Expectations

Understanding the conditions, procedures, and principles of emerging technological innovation requires dealing with a shifting field of contending stakeholders and institutions, marked by an entrepreneurial drive for commercial implementation and success amidst changing economic conditions (Callon, 1998; Freeman & Soete, 1997), ongoing institutional innovation (Barry, 2001; Jasanoff, 2005), and problems of political legitimacy (Brown, Rappert, & Webster, 2000; Van Lente & Rip, 1998). It is in these transition phases when technology moves from the experimental phase in a laboratory or company into a stable and marketable object where new questions become negotiated. This move from utopia to reality, from abstraction to materiality,

from promise to fulfillment, from singular to reproducible, is being constructed and negotiated by the actors themselves.

One way of making sense of these is to focus on how actors attempt to discuss the future. There has been growing recognition that future visions, or specific normative accounts of what the future should look like and what might be possible, are a crucial and constitutive element of the social and political development of emerging technologies. Some have argued that the articulated future visions of the relevant stakeholders are central to understanding the political economy of emerging science and technology fields (Brown, Rappert, & Webster, 2000; Fortun, 2008; Hedgecoe & Martin, 2003; Pollock & Williams, 2010; Shapin, 2008; Van Lente & Rip, 1998). So far, this futures research has been largely focused on specific areas of life sciences such as genomics (Fortun, 2008), membrane technology (Tutton, 2011; Van Lente & Rip, 1998), pharmacogenetics (Hedgecoe & Martin, 2003); nanotechnology (Lösch, 2006), stem cells (Brown & Kraft, 2006); and xenotransplantation (Michael & Brown, 2004).

This research suggests that future visions are performative; that speculative visions advanced about the future have economic and epistemic value in the present, and work to create social networks and alliances (Borup et al., 2006; Brown, Rappert, & Webster, 2000). Future visions not only create the institutional and organization momentum supporting the technology, but they matter on a design level, with consequences ultimately for the impact the technology can have. Callon (1987) argues that specific future visions and conceptions held by engineers and designers of technology determine not only the precise characteristics of the technology but also the social universe in which the technology would function. These characteristics become inscribed with power relations both in its design and its relationship with actors (Latour, 1991), which underscores the importance of the process and the specific persuasive futures that move

the technology. It is often difficult for advocates of AR to argue that people ‘need’ augmented reality, considering that people have lived for thousands of years without it. Therefore the affirmative burden falls onto the people promising the technology to demonstrate that it is necessary in an important context or that there is an intractable problem that only AR can solve.

If ‘future visions’ are the ultimate goal, ‘promises’ are the instruments by which actors stress normative actions and wishful enactments toward that desired future (Van Lente, 2000). Essentially, a future vision is the endpoint, whereas ‘promises’ about the development of supporting technologies or organizations are the steps to realize that future. Therefore in addition to mapping futures, it is important to also analyze and understand the promissory work that actors engage in regarding AR. Fortun (2008) notes that there are several different types of promises, such as performative utterances (e.g. “I will do x...”), statements that carry promise (e.g. this technology is becoming x...), objects that suggest promise (e.g. promising markets/outlooks), promises that seek to manage/bound uncertainty, and promises about people/organizations. Categorizing both the subject and the consequence of the promise allows us to understand what is being made responsible or heralded. For example, the utterance “I promise x” not only provides a description of the promise, but also places upon the subject a sense of accountability and an expectation to justify their future oriented claim (Borup et al., 2006). Fortun (2008) argues that promises are not limited to an individual; rather, individuals often make promises external to themselves that hold value. These different types of promises work to raise expectations of other persons, technologies, and organizations which are important to analyze.

Alongside positive promises that work to advance a particular technology, there are usually pessimistic visions associated with emerging technologies as well. Pessimism is an

understudied area of futures research, but can play a major role in suppressing competing futures. Geels and Smit (2000) have looked retrospectively at reasons why a variety of positive predictions fail to materialize, and lessons from these ‘failed technological futures’ can be seen in certain pessimistic accounts of AR. Tutton (2011) recently explored the concept of pessimism in to the field of biotechnology, arguing that these darker imaginaries play an important role in constructing futures both in terms of the discursive character of the vision itself as well as how actors account for and react to these negative visions. Beyond understanding the different positive promises being made about AR, this dissertation builds on Tutton’s (2011) call for exploring pessimism and how those futures are leveraged against technologies from opponents as well as supporters with differing agendas.

An important distinction to make here is between promises and expectations. Promises are claims that people make about what they will do and what they wish to happen, whereas expectations are what actors anticipate actually will happen (Borup et al., 2006). Expectations often overlap with future visions, but can also diverge – e.g. an actor might wish for something to occur, but fully expect something else to happen. This divide between promises and expectations is particularly important for understanding emerging technology, where it is not uncommon for groups to overpromise and ‘overhype’ certain developments (Pollock & Williams, 2010) while simultaneously pursuing other strategies. Thus, both the promises and the expectations of a community are valuable to map because stakeholders hold certain perspectives about how those visions come to pass. They are not isolated, but rather are closely related; promises and futures are intended to shape expectations, which in turn aim to change people’s actions, decision-making processes, and interactions with other actors in the community (Borup

et al., 2006). As Borup and colleagues (2006) conclude, we need studies that look *into the future* while also contributing to knowledge about the future already underway in the present.

Communities of Practice

Despite AR being an emerging technology, the people at these conferences were not just an assembly of new stakeholders, but rather a mix of new and old participants working across academia, industry, and policy. It was also an ongoing process, where conferences would be held throughout the year, along with smaller regional meetings where people would gather. In these settings I began to observe practices consistent with what has been described as a Community of Practice (CoP), in terms of how people relate to and teach one another, and how that explains the kinds of interactions and intelligence they produce (Lave & Wenger, 1991; Wenger, 1998). Lave & Wenger (1991) first defined CoP as: “participation in an activity system about which participants share understandings concerning what they are doing and what that means for their lives and for their communities (p. 98).” These CoP describe a complex set of social relationships, organized around a particular domain, where knowledge and skills are circulated, often through a practice of legitimate peripheral participation (LPP). LPP refers to actions and social practice that people in the community engage in as a way of gaining access to sources of understanding through growing involvement (Lave & Wenger, 1991).

The CoP concept has undergone a few modifications since its first articulation, as different authors have focused on different elements of the CoP. Lave and Wenger (1991) are primarily focused on legitimate peripheral participation, how these acts can be an analytical tool for understanding how people learn, and how these strategies are what contribute and alter a broader CoP. Wenger (1998) further elaborated on the idea of a CoP itself, what those group

structures entailed, and a lengthier list of characteristics that describe CoP (see figure 3). Wenger (1998) focused more on the identity that is mutually defined in the community, and the various trajectories that members can take through their participation in the community.

The concept of CoP was clarified further to reflect the changing nature of innovation:

“Communities of practice are groups of people who share a concern, a set of problems, or a passion about a topic, and who deepen their knowledge and expertise in this area by interacting on an ongoing basis (Wenger, McDermott, & Snyder, 2002; p. 4).”

Table 3 Wenger’s indicators of community of practice

1) sustained mutual relationships – harmonious or conflictual
2) shared ways of engaging in doing things together
3) the rapid flow of information and propagation of innovation
4) absence of introductory preambles, as if conversations and interactions were merely the continuation of an ongoing process
5) very quick setup of a problem to be discussed
6) substantial overlap in participants’ descriptions of who belongs
7) knowing what others know, what they can do, and how they can contribute to an enterprise
8) mutually defining identities
9) the ability to assess the appropriateness of actions and products
10) specific tools, representations, and other artefacts
11) local lore, shared stories, inside jokes, knowing laughter
12) jargon and shortcuts to communication as well as the ease of producing new ones
13) certain styles recognised as displaying membership
14) a shared discourse reflecting a certain perspective on the world.
Extracted from Wenger 1998: 125-6

(Figure 3 – Cox, 2005)

CoP offers an analytical lens to understand and explain how people learn and become socialized in certain communities (Lave & Wenger, 1991), how behaviors and attitudes are modeled (Lave & Wenger, 1991; Wenger, 1998), how identities are mutually constituted (Lave & Wenger, 1991; Orr, 1996), and how organizations might deploy them strategically to foster innovation (Hildreth & Kimble, 2004; Wenger, McDermott, & Snyder, 2002). I chose to use CoP as an analytical tool because I believe it more accurately describes the particular structural dynamic of the group, where it is both broader than an ‘occupational community’ (Van Maanen & Barley, 1984) because of the diverse set of actors (even though they nominally work on AR),

and ‘networks of practice’ (Brown & Duguid, 2001) because those are characterized by groups who may never know or know of one another, whereas members in the AR CoP have extensive knowledge of other members in the community and significant overlap in the description of who belongs. ‘Epistemic communities’ (Haas, 1992) describes a group that is far more unified in normative and descriptive beliefs than this one, as well as an authoritative policy perspective that is not yet present in this industry. Other descriptors like ‘occupational cultures’ (Schein, 1996) are useful in mapping particular organizations, whereas the AR CoP is more distributed and diverse with many different cultures coming into the community. These other frameworks are not mutually exclusive, however, rather they correspond to similar ideas as CoP (Wenger, McDermott, & Snyder, 2002). CoP is most appropriate in this instance, while also capturing the dynamic nature of this community over time as newcomers enter and attempt to move toward full participation.

CoP can describe a wide range of groups, from small to big, long-lived to short-lived, collocated or distributed, homogeneous or heterogeneous, inside or across organizational boundaries, spontaneous or intentional, and unrecognized to institutionalized (Wenger, McDermott, & Snyder, 2002). In this case, the AR community is moderately sized, not a few members but not necessarily in the multiple thousands. It is relatively short-lived, particularly on the span of generations that some communities have endured (e.g. midwives). It is a distributed community that comes together at these meetings, and is heterogeneous in terms of discipline, function, and background. It lives across organizational boundaries, is intentionally planned, and is recognized and supported by many member organizations as a valuable entity.

These conferences are a gathering of people working in the same domain of technology, but their efforts clearly went beyond just meeting tangible goals of networking and potential

business partnerships. Different stakeholders are trying to accomplish different goals, and the various events, rituals, and actions all serve multiple professional and social purposes. On one level they can be understood as serving a functional purpose, where people in the industry can meet one another and coordinate activities like standards. Here they function as formal meetings with turn taking and deliberation toward a common goal, as well as learning how to participate in these meetings under the tutelage of a longstanding member. These member organizations are increasing providing resources to the AR community, by sending people to these places to fill certain community roles and manage the community activities, hosting activities to bring the community together, and tools to manage and share the explicit knowledge that the community creates (Fontaine & Millen, 2002). Another way of making sense of these spaces is for their strategic and networking functions, where time is blocked off for people to build partnerships and professional relationships. The hope is that they can access some of the benefits of forming a community, in particular the ability to see an emerging market, product, technology, capabilities, and opportunities, as well as identify and access to experts and knowledge (Hildreth & Kimble, 2004). Still another is to look at these conferences as a ritual recognition of themselves as a group, a celebration of sorts that the industry is at a place that needs these kinds of big showcase events and to continually hype up the industry. Lastly, there are efforts to introduce newcomers to the technology, teach them how to promise the technology, and socialize people into the community by sharing practices and tools for contributing to the community.

In many ways, the more subtle purpose was to help newcomers learn how to be a member of the AR community: how to talk about the technology, how to promise the technology, and how to do work that might ultimately support the technology. The goal was to build a larger, more sustained group of active participants who would support the technology but also be able to

do some of the things necessary to help, whether that be through advocacy, development, resources, or creativity. These demonstrations of learning, from software kit tutorials to design competitions, were how newcomers got socialized through various forms of legitimate peripheral participation (Lave & Wenger, 1991). AR meetings are places where people are organized into certain tasks, learn how they can contribute and engage in those, share methods and practices, and in doing so gain the experience to become full participants. Language, discourse, and social practice form the ties that solidify the CoP (Lave & Wenger, 1991; Wenger, 1998). Many organizational scholars have argued that it is in this everyday talk, conversation, and interaction that meaning is made, and where cognitive processes and intelligences are linked between people engaged in practical work (Boden, 1994; Taylor & Van Every, 2000; Weick, 1995).

While these seminal works have offered many examples of communities and the different forms that they take, there are several important elements of CoP that require further examination, and which this case illuminates. The first is the issue of power, which the authors admit is an important part of CoP that requires further study: “unequal relations of power must be included more rigorously in our analysis (Lave & Wenger, 1991; p. 42).” Other scholars have advocated examining power inequality in CoP less as a result of presumed structural forces, but more as locally situated entanglements between protagonists and allies (Fox, 2000). In this particular case, we see empirically how much of the contestation is locally situated over important resources, epistemic authority, and other presumed actors. This way of understanding the community’s interactions also contributes to an understanding of power between relevant social groups, that they are not flattened or equal in their influence on the artifact, but that there is contestation amongst them over authority, attracting incoming supporters, and teaching them how to participate and support the CoP.

CoP also needs more empirical examination into the different circumstances under which the a community receives newcomers. Lave and Wenger (1991) talk about how different cases have differing levels of formal recognition from the masters and how there are many ways one can enter the community, but all of those cases were of controlled growth (e.g. tailors selected apprentices only when they wanted to) or need-based growth of the community (e.g. quartermasters selected people that they required for a trip). This case involves an instance where there is a mass influx of members, and how they bring with them their own visions and authority. The ways that changes the constitution and practices of the community, the way the community tries to sustain itself and hold on to their old visions, and the strategic push and pull between the old and the incoming helps understand a particular form of recruitment dynamics within a community. Lastly, the issue of conflict and displacement in these community is a critical one, which takes place both between individual groups as well as generationally. Because both newcomers and old-timers have a role in shaping the direction of the community as they establish their identity in the future, these conflicts are fundamental for understanding the displacement of certain norms and practices of social production.

Chapter Outline

Drawing from these core areas of literature, Chapter Two explains the methods for gathering and analyzing the data. Chapter Three first is an exploration of one of the key conflicts plaguing the AR community, specifically the form (mobile or headworn) that AR should take as a consumer device. That chapter draws on the futures literature to explain how different coalitions are forming, before looking at how that discussion helps us understand some early designs of the technology. Chapter Four delves more deeply in the AR community, specifically

looking at how they are dealing with newcomers in the space and how some actors are contesting the definitional inclusion of certain devices, applications, and people. Chapter Five examines one of the early clients for AR technology, specifically the marketing/advertising/retail industries, and how their influence has shaped both the structure of the AR industry and material technological capabilities. Chapter Six is a discussion of standards, namely how different groups attempt to rally support for a particular approach, how they strategically build on existing technologies, and the institutional alignment they undergo to strengthen their standard. It then analyzes the circumstances that led to a standards agreement, as well as certain decisions for what was included in the AR standard and the implications of those decisions. Chapter Seven looks at a recent turn in the community toward industrial applications, and the implications that might have for the technology.

The goal of each chapter is to illuminate not only the intermediary space and the work that happens there, but also how the contestations that occur demonstrate the social construction and shaping of these technologies. Lastly, these decisions and attempts to shape the community are not neutral, rather they have tangible implications for the material artifact, the content that is possible, and the uses people can and should engage in.

CHAPTER 2

METHODS

Fieldwork

This dissertation is based on over three years of fieldwork and participation observation in the AR community. Because my question attempts to make sense of situated activities and people's perspectives of the world and technology, qualitative methods are employed to “study things in their natural settings, attempting to make sense of, or interpret, phenomena in terms of the meanings people bring to them (Denzin & Lincoln, 2005; p. 3).” As a method for uncovering those relationships, “fieldwork’s stress on the taken-for-granted social routines, informal knowledge, and embodied social practice can yield understanding that cannot be obtained through [other methods] (Gupta & Ferguson, 1997; p. 36). This is at the core of the Communities of Practice literature, where much of the understanding comes not only from unspoken and unacknowledged practices, particularly in terms of how people learn, think, and act in relationships among people in, with, and arising from the socially and culturally structured world (Lave & Wenger, 1991). Examining how these social actions are structured, how systems of relationships are formed, and the implications of those is the focus of this dissertation, which calls for qualitative observation of the phenomena.

I first became interested in AR when I heard about the Layar browser, then heralded as the world's first publicly available AR application (although there has since been some discussion over who was first). When I saw a demonstration of their application online, I wanted to find out more about the people who were developing these technologies, so I began going to New York City to attend a meeting of Augmented Reality New York (ARNY), a monthly meetup group. This particular group was founded in November of 2009 and asked “developers, designers, entrepreneurs, scientists, educators, investors, artists, marketers, hackers, journalists,

and more - to bring the augmented reality experience to consumers in a fun, productive - and lucrative way (ARNY site).” I first attended an ARNY meeting in November of 2011, marking my entry into the AR community and the people involved. This was my way of ‘starting where I am,’ in finding the place that was locally and practically relevant (Lofland et al., 2006)

Those meetings were always structured in the same way, with an hour for social mingling and networking, followed by three or four presentations by AR companies, developers, or artists, followed by some time at the end for questions and more networking. It was there I first met many of the people who were already established leaders of the AR community, who would end up guiding me through the fieldsites and helping me understand the landscape of the AR space. In the end I made it to four ARNY meetings between November 2011 and April 2012 (some of them were cancelled by the organizer due to weather and other obligations).

It was at these meetings that I met the host of ARNY but also organized the largest AR conferences in the world and later founded AugmentedReality.org. I also met one of the world’s leading researchers in the field of AR. I met a police captain who works in cybercrime and is a leading expert on AR and legal issues. I met one of the leading artists working with AR and creating AR protest layers. I saw presentations from many AR companies as well as industries that were adopting AR or possibly funding AR. The people who would come to ARNY meetings were often based in New York, and while it is a major city and hub for AR development it also became clear that this conversation was also happening at a national and international level. This first entry point introduced an already diverse set of actors and interests in the AR space, but one that was clearly only part of a larger network of actors. This quickly began to push me toward other venues besides ARNY.

Simultaneously, the group was so diverse that I decided not to the specific production of any one device or application, but instead the collective spaces where the technology is made persuasive and how people come together in these spaces to do important practical and symbolic work surrounding the technology. With an emerging technology that is geographically distributed and diverse across fields, these are the places that bridge the work of the community and serves to connect them through social practice.

Van Lente (2000) explains that with emerging technologies, groups of technologists and early leaders often attempt to secure a social mandate to speak about these technologies. The technology and the people are at the heart of the analysis, but my goal is to focus not only on technological facts but on constitutive action, the intentions and deliberations which precede action and negotiate action between different actors. In particular there is a need to study how these groups coordinate their actions, by going to the places where they come together to form networks and mutual dependencies around the technologies they produce (Geels, 2004). With an emerging technology, considering these places and processes can “identify crucial moments, sites, and dynamics that mark the life cycle of media technologies (Boczkowski & Siles, 2014).” Over the course of three years, my fieldwork included extensive participant observation at a number of conferences, meetings, and workshops, as well as in-depth interviewing and secondary document analysis.

Participant Observation and the Importance of Industry Conferences

Participant observation is considered most appropriate when the research examines something that occurs in specific, physical locations (Lofland et al., 2006), and there were clearly places where the discursive articulation of AR was occurring. At the same time, important

decisions needed to be made about the choice of fieldsites. As a methodological approach, SCOT advocates finding the most relevant social groups and stakeholders that shape the technology (Bijker, 1995; Pinch & Bijker, 1987). With emerging technologies, however, identifying relevant social groups is difficult because these groups are still emerging and coalescing. I briefly considered approaching this problem by embedding myself in a particular company, but in the end I decided that this would have been too rigid a predetermination of the important actor groups. Instead I adopted the suggestion from SCOT to ‘follow the actors,’ and I found a place where those actors were congregating. Beginning with my first contacts at ARNY, it was clear that one major stakeholder group was the commercial companies who were forming around AR technologies. This business focus was made clear even in the description of the ARNY meetings that they hope to continue this development in a “lucrative way.” A number of the presenters at ARNY were demoing their companies AR product, either a game that they had developed or an enabling technology that would make AR experiences possible. It is from these early interactions with the community that I ended up selecting my data sites, where I could best observe the ‘promise champions,’ or the people whose futures are intricately linked with articulating the promise of the technology (Van Lente & Rip, 1998).

Around the time I was starting this project, some of the first industry centered AR conferences and trade shows were being held. The founders of ARNY also organized a conference called the Augmented Reality Event, which I decided to attend in the Spring of 2012 (ARE2012) because it was billed as the world’s largest AR conference. That was a full two day event, with keynote speeches and exhibits all day, an exhibition area that ran throughout the conference, and an ‘Auggies’ contest for best application/product, followed by more networking and parties in the evenings. Sessions were separated by track, depending on if it focused on the

‘business’ side of AR, the ‘technology’ side of AR, or ‘production.’ Here was where representatives from many stakeholder groups were converging and congregating, from mobile hardware companies, software development companies, app developers, academic researchers of AR, advertising agencies, artists, investors, hobbyists, and others. This umbrella event was in a way a legitimation for many of the participants, that the industry had grown so large that they needed such a conference.

ARE2012 became the first place where I began to see the multitude of stakeholders within AR interact, and it was a really important place for me to not only build connections but deepen my understanding of the event and the community. While there were many stakeholder groups present, the business/technology/production focus of ARE2012 was geared toward the commercial side of AR and the creation of an industry, which is why that became one group that I examined. I returned for the conference in 2013, which had been renamed Augmented World Expo (AWE2013) to reflect the international scope of AR and the spectacle of the conference. It had grown noticeably from the year before – the conference organizers boasted that there were over 1100 attendees that year, with over 100 demonstration booths that took up the entire exhibition floor of the Santa Clara Convention Center. Later that summer I attended the AR Summit 2013 in London, billed as the preeminent European conference on AR. It was smaller than ARE2012 and AWE2013, but had the same general structure, with an area for speakers and panels and a separate exhibit area with demonstration booths. In the Fall I attended InsideAR 2013 hosted by Metaio, a recently begun annual conference with about 750-800 attendees: where the ‘AR industry meets.’ The keynote address by the CEO of Metaio explained that it was a corporate event for the AR industry worldwide, so that developers, customers, tech partners, and even their competitors could come together and share their ideas and solutions. The most recent

industry conference I attended was Augmented World Expo New York 2014 (AWENY 2014), a New York City based conference held in March 2014 that served as a precursor to the larger AWE 2014 conference happening in Santa Clara that May. These conferences, along with the earlier ARNY meetings that I attended, were my primary fieldsites for understanding the industry actors of AR and how they were interacting, competing, envisioning the industry, and forming alliances, but represented places where all groups of the CoP were represented and came together.

From my first ARNY meeting, it became clear that there was another side of AR, specifically driven by academic researchers. The primary place where the academic discussion of AR was happening was the International Symposium on Mixed and Augmented Reality (ISMAR), which started almost a decade ago. In the late 1990s, there was some early academic organization around AR, with the International Workshop on Augmented Reality (IWAR), the International Symposium on Augmented Reality (ISAR), and the International Symposium on Mixed Reality (ISMR). ISMAR was formed in 2002 as the result of a merger between IWAR/ISAR and ISMR, bringing everyone together in a single location once a year to exchange key ideas and concepts, and becoming the premier conference on AR. The first ISMAR I attended was in November of 2012, in Atlanta, Georgia. Including the preconference I attended on classifying the AR space, it was a full four day conference with many paper sessions, an exhibit area and demo contest, a tour of the laboratory space at Georgia Tech, a banquet, and many coffee sessions in between.

This event had a decidedly different feel to the industry conferences, even though some of the same companies presented in the exhibit area, as it was primarily focused on the research papers that were being presented and discussed. The conference was also older than the industry

conferences, and rather than networking it seemed that many people had known each other for a long time, and students could find common ground based on the laboratories in which they worked. I returned to ISMAR 2013 the following year, this time in Adelaide, Australia, hosted at the University of South Australia. This conference was similar to the previous year, with paper sessions and an AR laboratory demonstration at the university. It also featured an ‘Industry Day’ before the start of the conference, which was intended to facilitate a dialogue between AR academics and industry needs and desires. This was one place where stakeholder groups were strategically reaching out to one another and attempting to coordinate their actions with one another. There was also an art exhibition featuring work from AR artists from around the world. These two conferences were where I observed the academic discussion surrounding AR, and how they were thinking about these technologies as a group and in relation to other stakeholders.

Lastly, it was at ARE2012 that I heard about the standardization efforts taking place with AR, where there was a panel scheduled about the standards development process for AR browsers. Some of the work was happening between companies and academics, and it seemed to be another kind of site where different stakeholder groups came together to negotiate and enact certain visions of AR (as well as protect their own interests). This seemed to be a place where certain goals were being contested and materialized, as standards are often politics by another means (Bowker & Star, 1999). Because of this, I decided to follow the standards discussion surrounding AR and attend those conferences. Having met one of the leaders of the standards movement, Tina, at ISMAR 2012, I began attending meetings where AR standards were being discussed and groups were updating one another on their progress with certain enabling standards. I attended the Open Geospatial Consortium technical and planning committee meeting in Redlands, California, in January 2013. That conference was my first exposure to the standards

setting process, in particular the organizational structure of the particular consortium that was developing AR standards. Later that year I attended the 9th International AR Standards Community (IARSC) meeting in Barcelona, Spain, a two day event with speakers from many different standards setting organizations (SSOs) as well as companies discussing what standards were necessary. In the summer of 2013 I attended the 10th IARSC meeting at Columbia University, and later that Fall I went to another OGC technical and planning committee meeting in Frascati, Italy. The last standards conference I attended was the 11th IARSC in Crystal City, Virginia. It was at these fieldsites I was able to observe the people who were participating in the discussion surrounding standards, as well as the ways they attempt to build consensus around these standards. All told I attended over 15 meetings and conferences in 10 different cities, spanning 6 different countries (see Figure 4, industry in blue, standards in green, academic in red)



(Figure 4 – AR Conference Map, Created by Author)

In the end, my fieldwork was conducted at meetings that were hosted around industry interests, standards activities, or academic research into AR, although there ended up being

significant overlap among the participants at these conferences. While many social studies of technology look to the design space or the creator space, these conferences are the coordinating space between those points, where products, ideas, strategies, and resources come out of the design space and are presented to others in the people in the industry for a variety of reasons. Drawing on SCOT as an analytical lens, my choice of conferences was based on where actors were gathering. The fact that they were organized around industry, academia, and standards activities is simply the broadest categorization of these meetings – indeed, within these groups there are many diverse interests represented. These conferences are not comprised of any one organization, rather they are a collective meeting ground for people from many relevant social groups to work around these technologies.

These sites are also incredibly complex in terms of what they are trying to accomplish, for different groups and members of the community. To highlight an example that illustrates this complexity, I'll just describe one session at AWE2013. This particular session was on the main stage (there were three staging areas for panels), and it was a 'Tech Start-Up Competition.' During that session, representatives from five different startup companies pitched their product to a panel of judges, all of whom worked for major venture capital (VC) firms. After their presentation the judges questioned them on stage about their product, their business model, their target audiences, their vision for the technology, and any other factors the judges deemed important in thinking about investing in the company. After deliberating, the judges announced the winner of the contest and gave them a \$5,000 prize. While this type of contest is common at these technology conferences, this one session can simultaneously be understood from a variety of perspectives. For the conference organizers, it can be understood as a way to encourage companies to come to the conference and reward them for innovation. For the companies who

choose to compete, the monetary prize might be paltry in VC terms, but the practice and experience of going through the VC pitch process is valuable and the awards could help build the reputation for future investments. The organizers also made the decision to hold this on the main stage, allowing the conference attendees to watch these pitches. This gave the startup companies more exposure than just to the VC judges, but, perhaps more importantly, was also a way for the organizers to educate the audience about the VC process by seeing it enacted. The session offered potential designers and entrepreneurs a way of seeing what factors they should consider when creating/evaluating an AR product, how to think about innovation through these VC frame, and the importance of business models as factors in garnering financial support for AR.

These multiple goals highlight the importance of physically attending these events, but also illustrate the types of interpretations of legitimate peripheral participation I was making in my writeup of these conferences (Lave & Wenger, 1991). Giving everyone access to this resource demonstrated that attracting funding is a desirable and necessary goal for the community, while also showing how it could be done. It is only by going to the places where they are can a researcher hope to understand who the stakeholders are, the stakes of the contestation, and why certain fields develop the way that it does (Gieryn, 1999). Within and across organized collectives of scientists, academics, corporate actors, and standards organizations, this debate is taking place at conferences and meetings where the work of making the technology persuasive takes place. Writing these interactions in fieldnotes is particularly important in pinpointing the indigenous meanings and interactional processes underlying those events (Emerson, Fretz, & Shaw, 1995). From these conferences I have over 300 pages of fieldnotes, from observations about formal sessions and events to social settings and mixers, all of which document elements of the social practices, values, and artifacts of these conferences.

In-Depth Interviews

While I was able to observe a great deal just by attending these conferences, I supplemented these observations with follow-up interviews with key leaders of the AR community and speakers/organizers of these events. According to Lofland et al. (2006), the specific advantage of in-depth interviewing is to “plumb the depths of meaning of action beyond group members’ talk in the field (p. 88).” Interviewing allowed me to learn more about their motivations, intentions, thought processes, reactions, and decision-making, as well as getting different stakeholders to weigh in on the same issue. Between November of 2011 and June of 2014, I conducted semi-structured in-depth interviews with 48 members of the AR technology scene and spoke to many more in informal face to face settings. I began these interviews shortly after I attended the first ARNY meeting, recruiting from among the presenters and attendees at those meetings. I continued to recruit at ARE2012, contacting speakers who presented at the conference. With each of these interviews, I also developed a snowball sample by asking these speakers who they felt to be authorities on AR. In particular I was looking to interview individuals who are recognized and identified as leaders by other members of the community. The consent document ensured confidentiality, so I created a table of all the participants with their pseudonym, gender, job description (without specific companies or schools mentioned), and the type of conference (industry, academic, standards) they were in attendance at (see Figure 5). Interviews lasted between 30 minutes and 150 minutes, with an average of 65 minutes.

Pseudonym	Gender	Conference	Professional Description
Bran	M	I	Law Partner Specializing in Social, Mobile, and Emerging Technologies
Karl	M	I	Director of Project Management for AR Hardware Company
Will	M	I	Chief Financial Officer of Major AR Company
Adam	M	I	Founder of an AR Consulting Firm and Blog
Ramsey	M	I	Police Chief Specializing in Cybercrime
Walter	M	S, I	CEO of a Major AR Software Company
Santos	M	I	Founder of an AR Media Website
Linus	M	I	CEO of a Major AR Software Company
Aaron	M	A, S, I	Founder of AugmentedReality.Org, Co-organizer of ARE, AWE, ARNY
Rodney	M	S, I	Programmer at Major AR Software Company
Fred	M	A, S, I	Professor of Computer Science at a University in Northeast U.S., ISMAR Steering Committee Member
Dennis	M	A, I	Ph.D Student of Computer Science at a University in Northeast U.S.
Luke	M	A	Ph.D Student of Computer Science at a University in Singapore
Mary	F	S	Professor and User Experience Researcher in Portugal
Arthur	M	S	Researcher and Software Engineer at a University in the U.K
Stan	M	A, S, I	Director of AR Laboratory at a University in the Southern U.S., ISMAR Steering Committee Member
Cary	F	I	Director of R+D for Location Based Mapping Company
Giovanni	M	S	Post-Doctoral Researcher in Computer Science at a University in the UK
Seth	M	A, I	ISMAR Steering Committee Member, Technologist in Residence for California Based Startup
Liu	M	A	ISMAR Steering Committee Member, Professor of Computer Science in New Zealand
Brian	M	A	Professor of Computer Science at a University in Midwest U.S.
Greta	F	A	Professor of English at University in Sweden
Colleen	F	A	Co-Director of AR Laboratory in the Netherlands
Victor	M	A	Co-Director of AR Laboratory in the Netherlands
Nathan	M	A, I	Digital Technology Artist
Nick	M	S	Technical Director of a Computer Software Company in the UK
Percy	M	S	Chief Engineer for a Standards Setting Organization
Mark	M	A, S, I	Chief Technical Officer for an AR Software Company
Boris	M	A, S	Founder of an AR company, Head of an AR Standards Development Group
Richard	M	S	President of a Large Standards Setting Organization
Simon	M	S	Participant in the ARML Standards Process

Evan	M	S	Participant in the ARML Standards Process
Tina	F	A, S, I	Head of an AR Standards Group
Hillis	M	I	CEO of an AR Application Company
Ray	M	I	CEO of a Construction Company adopting AR
Terry	M	I	Consultant for a Major Investment Company
Merton	M	I	CTO of an AR Sports Company
Gerald	M	I	Co-Founder of a Major AR Software Company
Morris	M	S, I	CEO of an AR Computer Vision Company
Dmitry	M	S, I	CEO of an AR Navigation Company
Josephine	F	I	Marketing Director for an AR Software Company
Murray	M	I	CEO of an AR Hardware Company
Rick	M	I	CEO of an AR Creative Agency
Ashton	M	I	Project Manager of a Major Computer Vision Company
Sven	M	I	Digital Technology Artist
Justin	M	I	Analyst for a Major Marketing Research Company
Riley	M	A, S, I	Professor of Computer Science at a University in Southeastern U.S.

(Figure 5 – Interview Participants and Descriptions, Created by Author)

Given that I was interviewing people from many different professional capacities, semi-structured interviews provided the best balance between guiding the interview in a direction that could yield rich data, while also allowing some flexibility for the interview to take form organically so as not to impose any a priori categorization that might limit the field of inquiry (Fontana & Frey, 2005). For the industry representatives, My questions included what their goals are, why they decided to invest in AR, how they foresee the industry playing out, what external forces they deal with, and their approach to competition. For academic AR researchers and hardware designers, I was interested in why they make certain technological decisions, what assumptions they make about their users, and toward what purpose they base their designs. For standards advocates, my questions were about the standards they were working on, the vision that motivated the to work on those standards, and the decisions they made regarding their development. For all interviewees, I asked them about their ideal visions for AR, the types of promises that are made surrounding AR, pessimism surrounding AR, and their expectations for

how AR will develop. I also asked them conference specific questions, e.g., to elaborate on a particular point they were trying to make during their presentations as well as their reactions to other speakers. The interview guide is attached in an appendix..

Documents, Materials, Listserv Discussions

The AR community congregates at these conferences, but there is quite a bit of work that goes on both before and after these meetings, as well as many secondary documents that are artifacts produced for and from these events. For the industry conferences, there are the programs and materials for the conference, as well as the press releases and promotional literature put out by the companies themselves (some are available at their booths, while others come included in the conference folder/tote bag). For the academic conferences, there are the published papers and posters that are the outputs of scholars' research. At the standards conferences, it is the draft documents written up about the standards, including specifications or glossaries or classification schemas, which represent their collective work. In particular standards setting organizations, much of the work happens remotely before and after the conferences via listservs. As a member of the Open Geospatial Consortium I had access to the online discussions about ARML 2.0, and was able to publicly access the standards conversation at the W3C on their website. These secondary documents are additional points of data that clarify the goals of these actors as well as their discursive efforts to align members of the community. Taking into account this documentary data provides another 'perspective of action,' revealing not only what certain actors are hoping to accomplish but also how certain disputes emerge, get debated, and ultimately get resolved or lead to compromise.

Data Analysis

I analyzed the data using the qualitative coding software Dedoose. What I began looking for were the future visions that people were advancing, what Berkhout (2006) defines as: “collectively-held and communicable schemata that represent future objectives and express the means by which these objectives will be realized (p. 6).” I coded for these as any time actors expressed or advocated particular futures for AR or pushed for developments in order to make those futures possible. Futures can take many forms, ranging from science fiction with no claim of feasibility to guiding visions and goals that are more concrete aims for technology development. (Grunwald, 2004). Because this inquiry is about the tangible promises and expectations driving the technology, I am focusing on guiding visions which have some claim to feasibility. On a general level, Michael (2000) explains that there are four characteristics of futures that may be present – whether they are positive or negative in valence, how distant or proximate they are, the degree those futures are tied to the present, and the speed at which those futures are approaching.

For each of these future visions, I coded for these characteristics using Dedoose. Valence was coded based on how actors were articulating those visions. Some visions are ambiguous as to whether they would ultimately be beneficial (e.g. police officers using AR), but if the actors were articulating that as a positive future it was coded accordingly, so that I was not making my own value judgments. Proximity was more straightforward, as actors would often caveat their visions as something that may be far in the future. I coded for the degree to which certain futures are tied to the present, a characteristic Michael (2000) classifies as instrumental or substantive, i.e., whether they assume substantive changes in the way society or humans are constituted (e.g.

human perception would adjust to the speed of augmentation). Lastly, speed was coded as fast or slow, depending on the rate in which actors were predicting these changes.

Delving into the specific content of the futures, I coded based on Berkhout's (2006) classification system of "*objectives*, the qualitative or quantitative expression of novel future outcomes; *orders*, a set of social and institutional relationships in which these objectives can be met; and *technologies*, the means for achieving objectives (p. 7)." Several quantitative objectives would be given for futures, like the AugmentedReality.org goal of one billion users of AR. Other futures would be more qualitative in nature, specifically that AR would be useful in people's everyday lives or fundamentally change certain social interactions or work situations (e.g. Metaio's theme of 'The Augmented City.'). The *orders* of these futures were coded as needing either social and institutional relationships to make those objectives a reality, such as the need for a change in social patterns, or organizations that need to either take certain actions (e.g. create an open process for standards development). AR is the technology that they believe will bring about those objectives, but the form of that technology varies a great deal depending on the vision. To understand these, I created a broad coding scheme that was generated iteratively from the data itself. I began by parsing out the distinct technological promises that actors were making about AR, and sorted those into seven broad categories: 1) obtaining information, 2) aggregating data, 3) recognizing real world content, 4) visualizing three-dimensional objects, 5) interacting with augmentations, 6) enhancing human capacity, and 7) restoring senses (Liao, 2012). After that, I coded for features that people were ascribing to the technology, the form of AR, as well as certain technological barriers or concerns that people had. From this perspective, the visions that actors hold for AR technology and the promises they choose to advance are an important unit of analysis, because they are "constitutive" or "performative" in attracting the interest of necessary

allies (various actors in innovation networks, investors, regulatory actors, users, etc.) and in defining roles and in building mutually binding obligations and agendas.” (Borup et al., 2006; p. 289).

Coding in this way began to reveal a number of distinctive visions for AR, which began to illuminate the ways that actors are deploying these futures as problem-defining frames, ways of monitoring progress, building actor networks, and narratives to focus resources. The next step in my analysis was to identify the places where people are communicating these visions and the patterns and interconnections between them. Charmaz (2002) explains that these constitute the “intermediate step between coding and the first draft of your completed analysis (p. 347).” I began writing qualitative memos about the people and organizations I observed, the interactions between them, and the physical layout and structure of these conferences. In particular I was looking at how certain groups were forming and championing particular visions and the relevant social groups that were enlisted in those visions. I was also examining the differences between visions, and the identification of fundamental assumptions between them. Where I noticed strategic ways that actors attempted to influence others to align behind those visions, I documented those as well. Lastly I documented the ways in which the community was engaging in social practices that attempted to pass on skills and knowledge to other members, Lave and Wenger (1991) explain that one should look for the structure of access that newcomers have to ongoing activities, resources, and skills, as well as the conflicts, interests, common meanings, intersecting meanings, and motivation of participants.

These memos form the basis of the upcoming chapters and the particular places where the contestation is occurring. As the areas of contestation were becoming clear, I went back to analyze all the data more closely to look at particular areas. With the debate over mobile or

headworn devices, I excerpted and analyzed each statement regarding the form of AR and how people can access it. With the definitional debates, I pulled all the data about those issues and reexamined those. With the contestation over marketing, I isolated and examined all of the perspectives and reactions to those uses. For the standards conversation, I reanalyzed the interviews with the standards participants as well as other actors' statements about standards. Lastly, with the push for industrial AR, I sorted specifically for visions specific to industrial contexts, to pull out similarities and patterns across those visions and justifications. These are examples where 'refocused questions' emerged from the data, to which multiple codings can further break down the data and yield more detailed insights (Lofland et al., 2006)

Methodological Choices

While the sections above outline the procedures I went through and how I gathered and coded data, it is also important to be reflexive about the implications of these decisions, the way I had to gain access, the unique way I situated myself in this community, and the community that I was studying. The first was in attending and being a participant at these meetings. The ARNY meetings that I attended were open to the public, which is part of the reason (along with proximity) that I started there. The large industry AR conferences such as ARE, AWE, InsideAR, and AR Summit conferences were open to anyone who could travel to those locations (which meant hundreds of dollars for domestic, multiple thousands of dollars for international) and pay the registration fees (which ranged from \$300-600). Access was more difficult with the standards organizations, as I had to contact the host of the AR standards community as well as negotiate membership in the Open Geospatial Consortium. These meetings were only open to members, and membership fees varied depending on the type of organization affiliated, as well

as the costs of attending their conferences. Access to ISMAR was also nominally open as well, but for Non-IEEE members the costs could exceed \$1000.

To attend all of these conferences one large barrier was cost; the way I was able to participate as a researcher was through grant funding. As part of the grant proposal process for the National Science Foundation I had to obtain permission letters from the organizations (OGC, ISMAR, and the International AR Standards Community), which was in part accomplished through personal connections but also required a thorough explanation of the project and a copy of the IRB form. The account of the work required careful examination. In doing so I adopted a learner role, hoping people would help me understand the lay of the land (Lofland et al., 2006). In all three instances I was granted access, but the communities I was joining are comprised of elites, which meant that I needed to justify my presence in order to gain access. My justification explained the project and the theoretical basis for it, which may have implied that they would benefit these organizations. The choice of research sites involved negotiating gateways, guards, and practical limitations, even in places where the purpose is to share ideas and work with others.

The second important decision was how I to make my presence as a researcher known, and how that shaped my relationship to the community. While I did my best to make my work known, at these conferences it was not possible (without drawing undue attention) to let all the speakers or the other hundreds of conference attendees know that I was a researcher. For the people I interviewed or spoke to for any extended period of time, I would identify myself as a researcher, studying how people were promising and pushing AR forward. At some conferences it was known after repeated visits that I was a researcher, because people had become familiar with me and my project. For those I was unable to inform, I do not use their names when quoting them, but I believe it is still ethically defensible to use that data because they were speaking in

their professional capacity, in a semi-public space, as representatives of themselves and their organizations.

One important consideration is how I became situated in the community. I initially worried about how I would be received by the community, as someone who neither codes, designs, builds, nor adopts AR technologies. The barrier for me was not necessarily the coding, but rather learning the language of describing AR and understanding the technical jargon. After three years, I have been welcomed into the community and have developed a close relationship with many people in the community. Some of this is a product of the industry itself, where there is a sense that more people need to know about this, and they need more people who understand the field and serve as advocates for the field. I found that many people were really excited about the technology and quite willing to speak candidly about their perspectives on AR (as well as other groups in the community), which made it easier to obtain rich interview data. I often used this to my benefit, as many of the questions I was asking could be perceived as ultimately benefitting the field. On a number of occasions people expressed interest in reading what I ultimately found, and one informant was quite emphatic about the need for someone to map some of these discussions. The familiarity I had with a number of key leaders also helped to establish my credibility – which was a side benefit of my initial efforts to seek out and make connections with the leaders of the community (Lofland et al., 2006). I did not interview people who were friends; rather I grew friendly with a number of people who I interviewed after seeing them at multiple conferences. Closeness is something that researchers should strive for, to acquire intimate familiarity with the social setting, as the closer we get to the phenomenon of interest the more accurate our descriptions of those meanings are likely to be (Becker, 1996). I did not come into the fieldsite with this closeness, rather mine was developed through the years

of talking to people and attending these conferences. This is one way I balance the tension of being sensitive to distance. The things I observed initially that seemed strange and new I can now explain with more subtlety and with a distinct understanding of the actors involved, while I can also still note what is strange both by considering how I might have thought about that when I first came into the community and knowing what the community finds strange.

Another consideration in the course of my fieldwork was the extent to which I would be a participant in this community. On the one hand, participation might allow for more access to certain people by building more connections and reputation within the community. On the other,, active participation in the group might start to intervene in the naturalistic setting and not allow myself the critical distance necessary in the field. Many scholars offer warnings of overidentification or ‘going native,’ where researchers might take accounts at face value or internalize the community perspective (Lofland et al., 2006). In the end, I decided the advantages of participation outweighed the disadvantages, particularly if my participation was not an effort to move the discussion one way or the other but instead to present my preliminary findings and discuss them when asked by members of the community. At ISMAR 2012 I presented a research paper about futures, though it was not an advocacy of a particular future but rather a way of discussing them in more organized ways. After that conference I was asked to be a guest on an AR podcast, where I answered questions about the paper and my research. At standards conferences I was asked to be on a task force for an AR glossary, where we were indstructed to reconcile terms that may have overlap or had two different terms that referred to the same thing. Each of these represents a strategic type of presentation for myself as a researcher to the community, and they were both how I negotiated relationships in the field and ways of understanding the field. For one, going through the ISMAR process allowed me to better

understand the processes by which that organization screens membership and values certain kinds of contributions, which was helpful in making sense of the legitimate peripheral participation that newcomers could engage in as well as some of the definitional work that was happening around AR. The presentation itself was also a way to demonstrate competence in the community, and I ultimately secured a number of interviews with ISMAR members as a result of it. The glossary task force helped me better understand some of the issues that the standards community was working on, but most importantly it gave me a reason to ask people questions about how they defined certain terms (including AR itself) and why those distinctions were important to the future of AR. I do not know how many people listened to the podcast, so I could not venture a guess as to whether that altered how people viewed me. What is clear, however, is that my participation in the community allowed me to gain access to certain people and rich data from personal interactions. I have also presented some of this work back to the community. In a move that reflected how heated and ongoing the AR definition question is, a member of the ISMAR steering committee asked me for all the responses I got to the definition question, because they were having conversations to restructure certain categories. These serve as important member checks that strengthen the internal validity of the data, but also as ‘exchange strategies’ that are not only important for access but also a means for developing and maintaining field relations (Lofland et al., 2006).

This does not come without tradeoffs, however, in terms of methodological and ethical dilemmas. With ISMAR in particular the venue could continue to be a professional outlet for me, which is something that could subtly limit my motivation to be critical of them. Additionally, while my goal is to provide a measured critique of the AR community, I found myself admiring many of the people that I observed and started to work with. I struggled a great deal with

potentially alienating people who were very kind to me and who might help me professionally. I try to be mindful of this in a few ways, both to allow me to be critical while also protecting individual people. First, the consent document that everyone agreed to protected confidentiality, so that will be protected. Second, I believe that the specific nature of my inquiry is measured in its criticism, where I might be criticizing particular visions that people have for the technology and the implications that those might have, without necessarily criticizing any individuals or organizations per se. Where organizations are an important part of the story, I include their name (not linked to any individuals who represent those organizations), but not in instances where that information is sensitive (e.g. the company that dropped out of the standards agreement) or unimportant (e.g. the fact that a retail chain is considering implementing AR is the essential part of the story, not the specific retail chain). Lastly, each of the chapters raise an issue that is contested by another group within the community, and thus the account captures both sides of the discussion in which I attempt to treat criticisms and rebuttals with equal respect. Becoming a member of this community got me close to where these conversations were happening, which is how I came to understand some of the latent meanings between actors and the relationships and goals they are trying to enact.

Even if I can be critical of individual people and visions, can I be critical of AR as a technology more generally? A ‘promise champion’ is defined as someone whose interests are intricately tied to the success of a technology (Van Lente & Rip, 1998). As a consequence of studying an emerging technology and speaking about the technology, Have I become a promise champion? My account of AR may be the first exposure that some people will have of the technology, which is something I have to keep in mind. Professionally, both amongst the people who I study and the people who I have to convince that this area is worthwhile, it could be a

benefit if the technology were to succeed because it would increase the relevance of my work. I reconcile this tension in a couple of ways. One way is to simply be honest about the potential impact that this work could have on the development of the technology, which is likely minimal. Many of the stakeholder motivations stem from strong personal beliefs and/or institutional pressures and goals, which would have a far greater impact on the technology than this account of the persuasive work that influences the technology. Another way to minimize this bias is to consider these technologies from an ethical perspective, thinking about whether I would want to live in a world where AR technologies looked a certain way and embodied certain assumptions. Particularly as a scholar who believes that these technological artifacts are embedded with certain values, it is just as important to be critical of AR where they could lead to negative outcomes. The extent to which I 'advocate' AR, then, is not an unbridled advocacy of AR technologies in all contexts, but rather an attempt to demonstrate the ways that actors in the space are cognizant (or not) of the underlying assumptions that are built into the technology and the places where that contestation is occurring. If it happens that AR does not reach the lofty technological, industrial, or commercial goals that many actors are hoping for, this analysis will have illuminated some of the struggles that pushed the technology in a particular direction and some of the factors that contributed to it. It may actually be that some of the problems that groups are raising prove to be intractable, or socially unacceptable. However AR develops, it will be these people and communities that will make many of the decisions contributing to that development. Therefore, while the account may contain many positive futures and promises of AR as a result of sampling leaders of the community, my data also captures instances where they disagree as to the appropriate path to take. The analysis explains how actors are promising the technology and the characteristics of those visions, and some of their optimism may leak in

unavoidably, but the writeup is grounded without any illusion of ‘inevitable success,’ or even the presumption that it would be desirable.

Gender and Race

These communities certainly have issues in terms of their levels of participation and inclusiveness, but part of that is due to larger structural factors, the degree to which certain people are attracted to the community, and the openness of the community itself. These are an important part of the story, particularly in considering who they consider to be relevant social groups and who they think they are designing for. I take the perspective that it is impossible to separate my own positionality from the analysis, rather the important thing is to be reflexive about it. One thing that might be striking to note is the gender representation of my interview sample (Men, n=41, Women, n=6). Unfortunately, at many of these conferences, that ratio is representative of the people participating in the AR community, which reflects many of the structural barriers for women in the technology space (Wacjman, 2000). Many scholars have noted the myriad of ways that the emerging technology industry privileges a white, male, experience and also how those structures systematically exclude women and people of color from relevant professional networks (Neff, 2012; Marwick, 2013). My account did not intend to reify this representation gap; indeed, I attempted to interview a number of women and minority leaders in the community. As with all recruitment, however, there is inevitably some non-response, which further compounds the issue of representation when there is already a disparity. What the real-time mapping does allow for, however, is an actual accounting of these perspectives in the participant observation data at these conferences. The analysis found that despite the representation gap, there are attempts in the community to factor in gender as an issue

to be considered. The ways in which they do so, however, is often without equal participation of women in the process, another form of mediation and power disparity between relevant social groups.

The other element that is important to this story is one of race, and it was something that I became conscious of both as a mapping of who participates in this community and as an experience of my participation. As with gender I did not sample based on race, rather I based it on the snowball sample. While I did not explicitly ask participants for race data, as I thought it might have seemed unusual given the line of inquiry, the vast majority of participants were white and American or European. Of the participants of color at the conference, most were Asian or South Asian. Only two of my interview participants were Hispanic. This is important both to understand the composition of the community but also to explain my own positionality in this group. As an Asian male it was not unusual for me to be at any of these conferences, and my participation was never questioned. In fact, there were times where the presumption was that I was a computer science graduate student, and I was granted a certain legitimacy in my interactions with people before I could explain my reason for being there (on numerous occasions people would read my school affiliation and begin asking about computer science professors at my university, and one time someone even began to recruit me for a job). There are times where ascriptive categories such as race can be a useful lens, given that they are important realities (albeit socially constructed ones) that need to be accounted for in planning research (Lofland et al., 2006). Negotiating these social differences is an important part of understanding the field, and I note this not as a barrier to access, rather there were definitely times when my gender and race worked to my advantage in terms of blending in and a providing a presumed authority to ask people questions. These observations, however, say something about the

community itself and the sociological/structural factors that contribute to the demographics of the emerging technology space more generally. The larger issue, however, are the ways these actors attempt to represent a more diverse set of users than the relative homogeneity of the space both in terms of gender and race.

Conclusion

The AR community is still going strong, and I am still a part of many of those groups. I think the most important decisions I made were about how to get along in the community, and I ultimately think that the decision to engage helped me better understand what was happening. First, the fact that I was recruited and asked to participate in a variety of contexts demonstrates both a curiosity about my presence and research, as well as the willingness of the community to continually socialize members into various kinds of activities. It also showed me how that process happened, as in the standards meetings I was placed on a task force with another graduate student from Spain and a university researcher from the UK. All of a sudden as I was writing about the definitions question, I was now authorized to define and merge different AR terms. While we met several times over Skype to coordinate activities, that effort is currently inactive, reflecting the difficult and ad hoc nature of voluntary standards work.

The decision to try to map the whole process of AR was also a challenge. Attending these conferences, dispersed all over the world, illustrates the degree to which following these technologies can be prohibitive to many researchers both in terms of time, money, and effort. I was fortunate in this regard, to be able to follow different entities to their local meetings and conferences, and to see the relationships and interconnectedness between these people and organizations. I once asked someone why the standards conferences had to be geographically

dispersed, and the response was ‘because standards is a global activity that affect companies internationally.’ I suspect that this globalization will become a necessary feature of studying emerging technologies, rather than the exception.

CHAPTER 3

MOBILE VERSUS HEADWORN AR

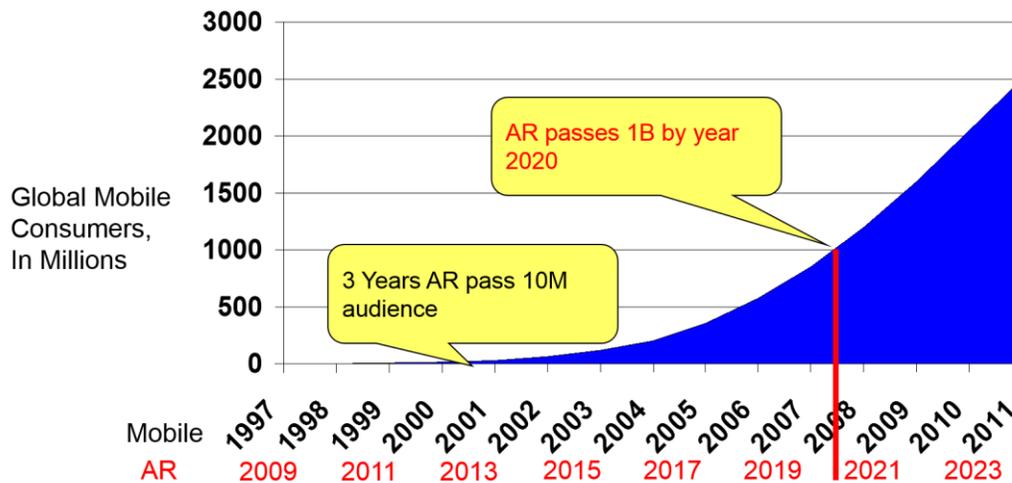
“We want people to be ‘Always On, Always Augmented’”
(InsideAR 2013, Opening Address)

“What’s our big, hairy and audacious goal (a B-HAG)? That’s right! To inspire 1 billion people to actively use of augmented reality by the year 2020.”
(AWE 2013, Opening Address)

These were the major themes at two of the largest AR conferences in the world, and they represent the unifying visions that these organizations hope to achieve and the mantras that guide their development efforts for AR. InsideAR was organized by Metaio, one of the biggest companies in the AR space based in Munich. AWE2013 was put on by AugmentedReality.org, a collective of supporters comprised of many organizations and leaders in AR. Through its organizing efforts, it is quickly becoming one of the most influential organizations driving AR. They are responsible for some of the largest AR conferences (ARE, AWE) and hosting AR meetups in dozens of cities in the US (ARNY, ARLA, etc.). Their goal of 1 billion users by 2020 is not just something that appears in the mission statement of the organization, it is a serious objective that is constantly being enacted and reinforced at conferences hosted by this organization (sometimes they start off the conferences with a quiz about their organization, throwing t-shirts into the crowd for correct answers). At AWE2013, they invited Tomi Ahonen as one of the keynote speakers, a market analyst who attempted to persuade the audience that AR was about to take off, and start on a path toward the organizations’ ambitious goal. The crux of Ahonen’s argument was that AR adoption is currently in a state similar to mobile phones in the early 2000’s, hovering around ten million users. The goal of 1 billion users would be possible if AR takes the same adoption curve that mobile phones did (see Figure 6). Ahonen’s selection as

the keynote speaker was a strategic one, as he is a well-known figure in the commercial technology space with many books and presentations about market trends.

The particular argument he was making can be understood as a plea for the feasibility of this organizational goal for AR, as well as counsel to members of the community to keep the faith and not panic. People had been projecting a big leap for a few years now, and this exhortation was intended to convince the audience that despite a current stagnation, these technologies take will experience an exponential jump within a few years. It also adds a temporal element to the charge for the community, positing that the next few years will be the critical moment where decisions that affect these futures need to be made. The goals behind these kinds of prescriptions are twofold – one is to unify the community behind the goal of supporting AR for public use, the other is to reassure and challenge people in the audience to invest their time, effort, resources, and energy into AR. In the audience were hundreds of people, who comprise the many stakeholder groups that are interested in AR.



(Figure 6 – AR Adoption Curve, Ahonen, 2013)

I open with these visions for a few reasons – (1) because it demonstrates the ambitious scope that AR advocates are aiming for, (2) because it illustrates the perceived stakes of the

contestation in the field, (3) because it motivates a sense of urgency amongst participants, and (4) because the AugmentedReality.org goal is one of the most prominent future trajectories publicly advanced within the industry and one that many people rally around. Circulating these types of communal goals is a way of organizing the participation of the entire community, old and new alike. Certainly, many stakeholders in the community would benefit were this to become a reality, as long as their interests are represented in the enactment of it. Therein lies the source of much of the contestation, then, as this graph is only a flattened expectation and numerical quantification. Missing from this vision, however, were any recommendations for how to reach that goal, or prescriptions for the structural and social relationships necessary to make that happen. It was presented more as a declaration that this was going to happen, simply by drawing an analogy between the adoption curves of two distinct technologies.

This dissertation is in large part a commentary on this graph and what is left out of it, the places where the conditions for the technology are set and negotiated, and a critique of the technological determinism that underlies these persuasive visions and how actors understand the technology. While these adoption curves and augmented futures are presented at these conferences as likelihoods or even inevitabilities, scholars of science and technology have been arguing for years that there is nothing inevitable about a technology or the trajectory that it takes (Abbate, 1999; Akrich, 1992; Bijker, 1995; Bijker & Law, 1992; Flichy, 2007; Kline & Pinch, 1996; Pinch & Bijker, 1987; Oudshoorn & Pinch, 2003). The development of a technology and its uses is the result of complex entanglements of chance, necessity, and collective human will (Flichy, 2007), as well as economic, political, and cultural pressures (Akrich, 1992; Bijker, 1995; Hughes, 1989; Latour, 1991). The attendees at these conferences represent various factions of the complex ecosystem that is the AR community, each with different interests in a variety of

technological fields and disciplines. This group of elites are speculating about this technology, and many of them have financial, political, and institutional stakes in it. It is a place where collective visions get circulated, where expectations are set, where actors are taught how to critique and defend against threats to the technology, and where actors coalesce around particular goals, all while simultaneously celebrating and legitimating the technology.

As multiple groups coalesce around different areas of AR technology, it is the resources these groups bring to bear, the problems they choose to tackle, the goals they collectively pursue, and the contours of their debates with competing coalitions that will influence the design, distribution, content, and stabilization of the technology. This chapter is an attempt to unpack this presumed linear trajectory of mass AR adoption, and make sense of how actors are positioning themselves to realize these futures. First I look at the visions that organize these coalitions, the expectations they have for certain user requirements and technological developments, and the discursive moves made to mobilize those coalitions to collective action. Even amongst those that believe in a billion-strong consumer market for AR, there are drastically different perspectives of what the future ‘should’ look like, the kind of device that will reach these goals, and the steps people should take now to secure the conditions necessary to realize that future.

A Brief History of AR Devices

The origins of AR are often traced back to Ivan Sutherland (1965) and his vision for ‘The Ultimate Display,’ which described a system that might create an ‘Alice in Wonderland’ world, with graphics that would appear to have the properties of physical matter. In pursuit of this vision, Sutherland created the first AR system, nicknamed the ‘Sword of Damocles,’ a ceiling

mounted computer that a user could strap their head in while it tracked a users' head and eye movement. It was also not a fully immersive virtual environment; users were able to see the physical space beyond the device, which is why it is considered the first AR display (Sutherland, 1968). Sutherland realized the technology was not ready to realize his vision, but he continued to work on testing the possibilities for creating 3D virtual objects for medical applications (National Research Council, 1999).

Based on Sutherland's (1968) original vision, there was substantial research in the following decades on head mounted or head worn displays (Azuma, 1997; Mann, 1997; Milgram & Kishino, 1994; Rheingold, 1993; van Krevelen & Poelman, 2010). Sutherland would later reflect that his work on head mounted displays and AR had served as an important "attention focuser" that "defined a set of problems that motivated people for a long time (National Research Council, 1999; p. 237)." While I must largely set aside the long prehistory of headworn devices (see Roland & Cakmaki, 2005), it is important to note that there was a strong academic research tradition surrounding HWD, and the development of VR also focused on headworn devices, which led to similar work being done with AR. Although many of these are still experimental, a number of academic and industrial actors in the AR community come from that tradition, and have that as their ultimate goal.

Some still refer to Sutherland's device as a head *mounted* display because the system was literally mounted and bolted to the ceiling. There has since been a recognition that these devices should not be confined to a particular space. While some of my actors continue use the phrase head mounted, for the remainder of this dissertation I will refer to these systems as 'headworn' devices (HWD). Regardless, the next step beyond getting HWD displays to work was to make them mobile in the real world (Feiner et al., 1997; 1999). This move broadened the scope of the

problem requiring the device to take the location of the user and the context they are working in into account. One of the first mobile outdoor AR systems was created in the late 90's, and required a GPS system, a wireless network, 3D graphics acceleration, and a mobile computer (see figure 7).



(Figure 7 – Mobile AR System, Feiner et al., 1997)

This system was an important first step in demonstrating how AR might be possible in an outdoor environment, but was clearly a prototype – the backpack itself was quite heavy, hooking it to a user took a long time, and the items were quite bulky (the researchers also acknowledged that they used off-the-shelf hardware in many instances rather than building their own, for cost/time/effort reasons). Recognizing that mobility and outdoor use was important for a number of AR applications (Feiner, 1999), a number of laboratories began to experiment with their own backpacks and PDAs (Papagiannakis, Singh, & Magnenat-Thalmann, 2008). But mobility was still linked to the AR experience, as seen in the picture, through head worn display.

It was only in the late 2000's, however, that many of the technological features necessary to run these applications became available on mobile smartphone devices (van Krevelen & Poelman, 2010). Smartphones suited the mobile element of AR, but offered a very different type

of display system than the headworn ones that was originally envisioned. Researchers and industrial actors began to see mobile smartphones as a viable option for AR, as they “usually come as a fully AR-capable package, [while] [...] the implementation of a back-pack setup is a difficult and non-trivial engineering task (Wagner, 2007; p. 5).” AR received a surge of attention around 2008-2009, when the first generation of smartphones with GPS, wireless, cameras, accelerometers, gyroscopes, data, and a burgeoning market for applications was released. AR companies like Layar, Wikitude, Junaio and a variety of AR games introduced applications that users could easily install, hold up, and access augmented content on top of physical space (see figure 8). Available on the first generation of smartphones, early uses were simple browsing applications, primarily for finding physical locations around the user. This represented a significant technological development, a convergence of multiple technologies necessary to support an AR experience, for everyday people, out in the world.



(Figure 8 – Metaio’s Augmented City, creator.metaio.com)

When I entered the AR community in 2011, the wave of excitement about mobile AR was strong, and had just brought a new set of actors (e.g. chip manufacturers, phone manufacturers, app developers, etc.) to the CoP. For many of the academics who had been

working on the space for a long time, however, this influx complicated the things they were trying to accomplish and their relationship with these newcomers. For one, it ran counter to Sutherland's (1968) vision of headworn displays for graphics, and almost immediately actors who supported headworn AR reacted to the challenge posed by mobile devices. Up until this point, public consumers for HWD AR were mostly hypothetical, aside from the companies who were creating eyewear as headworn TV screens, which do not interact with the real environment. With mobile, there were now apps that users could access immediately, one that had the potential to change user preferences, practices, and conceptions of the technology in an instant. It is in these moments, where new configurations and designs break through, that it is possible for certain actors to strategically align or break away to form new socio-technical regimes (Geels, 2004). Here, in this the moment of negotiation and transition, an alternative technological design and device was introduced as a new representative of AR technology, one that could become the dominant consumer device. I will argue that the introduction of mobile AR brought new relevant social groups, new problems and solutions, and new presumed users to the community, and spurred HWD AR actors to strategically defend their space and marginalize mobile devices. In particular, two alternative visions for the development of the technology are being articulated and embedded in a range of heterogeneous futures, artifacts, actor strategies and practices, including: competing research agendas, experimental technologies, and emerging industrial structures.

Mobile versus HWD

Mobile AR brought not only new technological possibilities to the AR community, but also the possibility of a new technological order. Some acknowledged that it kind of snuck up on

them: “We’d been thinking about AR as eyewear for so long. I went out and bought up every one and all of those devices, hoping to use it for research. Then all of a sudden this happens, and it is not quite what you expected (Fred, AR Eyewear Exhibition, 2012).” The companies that launched AR browsers were also getting a significant amount of media attention as well as VC funding. Some developers in the community began to create applications for mobile devices, and the companies were reporting millions of users downloading these applications (Butchart, 2011). Tech journalists were quick to proclaim that AR had finally arrived, with mobile smartphones leading the charge (Metz, 2012). This introduced a different regime for the AR community, as new relevant social groups came in with their own set of priorities, trajectories, and uses for the technology. These periods of mis-alignment and instability of socio-technical systems open up room for drastic realignment of the communities (Geels, 2004), and highlight the interpretive flexibility of these artifacts (Pinch & Bijker, 1987).

The debate between mobile and HWD has often been couched in terms of their technological differences alone (van Krevelen & Poelman, 2010), but this is only part of the story. Much of the work of setting these expectations and pushing for each form is performed at these conferences. This contestation has been ongoing for a few years now, and the result is two distinct coalitions within the AR community – mobile advocates and HWD advocates – even though from an outside perspective it might appear that everyone generally supports AR. The stakes are high, particularly in the immediate future, as the competing devices try to attract users and draw the infrastructure in line with their systems. The motivations for taking either side are obviously complex, but it is some combination of personal belief in the benefits of a particular form, professional and economic interest, and a sense of stewardship for the technology. These are also not static groups, as some might have initially supported HWD but found promise in

experimenting with mobile, while others may have come into the community because of mobile but were then drawn to HWD. The most common critique leveled by both sides is that the public will reject AR in its competing form, and subsequently categorically reject AR as a technology, irrespective of form. This was a concern voiced over and over again, a real sense that they might ‘ruin it for everybody.’ Here’s Linus, a CEO of a major AR software company, explaining their fears about mobile: “We tried to do these things with mobile, and people got excited about them, and customers weren't doing it. They started coming back to us, and literally they said ‘what else does AR do, because this has been a waste of time’ (interview).” On the other side, there are the same concerns about HWD: “It’s a possibility that the eyewear or the heads up display technology might go nowhere [...] and people might adopt the idea that AR is incredibly gimmicky and decide to stay away from it (Rodney, interview).” Others are concerned about their professional interests in the field, and their role in it, should the other device win out. Some have attempted to accommodate the new forms as a way of hedging, but the results have been mixed. For example, one browser company tried to port their content to Google Glass, only to find that the screen was too small to see the content and made for a poor user experience

. These stakeholders are also competing for developers, given that applications and content are considered critical to the success of any device. Each of these devices wants the unique applications and content, and the degree to which 3rd party developers believe that mobile or headworn will reach the most number of users could guide their development efforts. Funding is another major reason that these devices compete as separate categories, because if the problem of AR is solved by mobile, then VC funding and other grant funding supporting this research might flow toward that direction. Amongst some of the AR researchers who have been working on HWD for a long time, these practical concerns as well as a sense of stewardship are major

motivations. There's a very palpable disappointment that this meager form (mobile AR with its small screen) that emerged might ultimately derail or displace their efforts to make the device they have been dreaming about. During one of my interviews, one professor Fred pulled out his phone to give a laundry list of complaints about mobile AR applications. One of his students later explained to me that he's been working in the field for almost 30 years, and he finally believes he might see HWD AR truly succeed in his lifetime, and is worried that it might be derailed.

Just as important as the stakeholders who have aligned strongly with one form or another are those who are undecided, or who are attempting to straddle the line between these two coalitions. The organizational perspective of AugmentedReality.org attempts to take this middle ground:

Smartphones became capable of delivering augmented reality experiences to the masses and augmented reality broke out with much fanfare. Millions of users, thousands of applications, and hundreds of companies popped in the space. Users were enamored for a moment but few applications were used regularly. GPS-based floating bubbles sparked the imagination but failed to be actively used for navigation or information discovery; playing sensor-based games with camera view as backdrop, or watching videos and animations overlaid on flat surfaces without real world interaction – had an initial wow factor but provided little distinctive value. [...] Augmented Reality reached a fragile point in its life cycle. It risked becoming a short lived gimmick. A fad. Was that it? A group of people: designers, entrepreneurs, engineers – obsessed with solving real world problems with AR – thought differently. They believed that with the right mindset, AR apps on smartphones can be very useful – in the present; more importantly, advancing AR on existing platforms is a necessary step to prepare the ground (users, technologies, applications, designers) for the new way of digital interaction with the world – *once eyewear matures* (AugmentedReality.Org Story, emphasis mine).

Contained in this story are many of the fears of the community (particularly that it is fragile enough to disappear as a fad), as well as objectives and prescriptions for how to avoid those possibilities. It recalls the initial enthusiasm surrounding mobile smartphone technology, but

now with hindsight expresses disdain for those early applications and uses. This denial of its own history is common at these conferences, as a way of demonstrating experience but also as an explanation for why things may not have taken off as they had hoped. This particular version supports a short term continuation of development for mobile devices (because that's all we have), but with an explicit eye towards moving toward eyewear devices. This is the mission statement for the organization that plays a major role in organizing these conferences, advocating mobile AR as an intermediary form that will ultimately give way to HWD as the artifact that will allow AR to reach its numerical goal (1 billion) of a particular category of users (public consumers) in a given timeframe (by 2020).

On an institutional level, bridging these two coalitions is pragmatic in that it casts a wide net so as to enlist many of the actors in the space, although it does not fully satisfy both camps or conclude in favor of mobile. The actors in the community who adopt this middle ground perspective, some of whom are developers and funders, are an important group because they are the ones hearing the arguments being made on either side, and represent a major constituency that these coalitions are trying to persuade. The timeline for a proposed transition away from mobile is also vague, as there is considerable debate as to when HWD will ever stabilize or be 'mature enough,' such that there will continue to be ongoing contestation amongst advocates of either device as to whether/when to switch over. Depending on how these pragmatists interpret and engage with certain promises, they might be willing to move toward their coalition if convinced that this is the particular moment to make the switch, and ultimately be more inclined to support/advocate certain actions for a particular form.

Here I draw on the SCOT concept of relevant social groups, which are multiple entities that collectively assign their own meanings to technologies (Bijker, 1995). It is these relevant

social groups that have different interests and resources that they draw on to negotiate the proper structure and form of these artifacts (Pinch & Bijker, 1987). These stakeholder groups are split on the form debate, as actors within these groups fall on all sides of the coalition. Here is an instance where the relevant stakeholder groups were not pre-formed into categories, rather they sorted themselves into different actor categories depending on a) if they support any particular form, and b) if so, which one. The way each coalition is contesting this is by advancing futures that ultimately enlist different sets of actor groups and leave out others, with implications for how these groups are organized and the alliances that form between them (Berkhout, 2006; Brown, Rappert, & Webster, 2000; Hedgecoe & Martin, 2003). The trajectory that the technology takes is in part contingent on who makes these futures persuasive, and how actors align behind certain futures that will ultimately shape the technology. What these stakeholders are doing is an attempt to concoct winning visions of uses and users that will ultimately enlist economic support, bring key actors together, and marginalize competitors. The first part of this chapter maps out the places of contestation, in terms of speculation about technologies, future markets, and future applications.

The second half of this chapter examines some of the devices and applications that are beginning to come out, and analyzes them as solutions to and rejections of critiques from the mobile and headworn coalitions. The material design of these technologies also reveals how technologies reproduce, embody, and are influenced by professional, technical, economic, and political factors (Bijker & Law, 1992). In the design of technologies we begin to see them not only engaging in and preempting criticism from other actors within the community, but also embedding larger values into the design and presumed uses of the technology, as Bijker (1995) observed in the design of the bicycle, which embedded, raised, and weighed in on issues of

public safety, gender, and class. While AR is still very much in the early stage, these designs “represent attempts to predetermine the setting that users are asked to imagine for a particular piece of technology and the pre-scriptions (notices, contracts, advice, etc.) that accompany it (Akrich, 1992; p. 208).” Because the contexts in which AR will be used are still in question, they are actually being defined by the stakeholders who design the objects, which reorients those spaces. The design of these devices are like a set of tools for unknown users, which contain ‘technically delegated pre-scriptions’ from the designer to the user (Akrich, 1992). As some of these companies gain understanding about users’ habits and practices of engaging with these devices, and societal reaction to these devices begin to emerge, they will adjust the design and policies that regulate the device itself (Woolgar, 1991). Within the larger contestation between the coalition’s supporting different forms, the design decisions built into these machines can be understood as ways of configuring the user.

Technological Futures (Mobile Advocates)

One powerful argument that mobile AR advocates advance is that it is possible to have AR experiences with mobile handheld devices right now. Mobile device AR is accessed through smart phones and/or tablets, and this is the form in which the public is beginning to access and encounter AR. One common promise actors advance in support of mobile device AR is that it builds upon technological infrastructure that already exists and smartphones that many people already own. Dennis, an academic developing AR in the lab, explains:

I think phones will be prominent. They are there. They are mini-computers. They are doing things, they have the same processing power and sensors and technology that [researchers] were doing with a big ass computer in the mid 90's. [...] There are certain limitations like tracking and getting accurate responses from GPS that makes it necessary to attach things, but those are problems in sizes and that will get reduced more and more [...]. The pieces are there, and we're

seeing that the remaining pieces are attaching and coming together in the next 6 months, it makes sense that it should become a dominant platform in the next year to five years (interview).

Dennis's suggests that these phones are an analogue to those early backpacks of the 1990s. He also posits a version of Moore's law, that advances in computation will soon solve the problems that exist now and will continue to carry mobile technology forward. This is an example where expectations become embedded in actors' visions, and are used as a heuristic in their search for solutions (Berkhout, 2006).

At the same time, the fact that mobile AR is becoming the status quo highlights some of the challenges that confront HWD AR, and opens up space for mobile advocates to raise questions about whether HWD is even technologically feasible. With HWD there are many possible solutions, ranging from video see through (capturing the real world with video and electronically combining the digital images), optical see through (the real world is shown through semi-transparent mirrors), virtual retinal displays (the system projects the augmentation onto a users' eye), or headworn projective displays (the system projects the augmentation out in front of the user). But for each of these, there are tradeoffs and technical problems in terms of hardware (optics, processing, battery, etc.) as well as function (brightness, contrast, resolution, field of view, color, occlusion, etc.) that still need to be solved (van Krevelen & Poelman, 2010). Here is Aaron, one of the organizers of a large AR industry conference, describing some of the difficulties in designing HWD:

People think it [HWD] will never work. Even people in the field think it will never work. There are some technical limitations there that are justifiable. One is where you have a screen that completely covers your field of view and so it's a live video and overlies with graphics. The second option is to have semi-transparent glasses that only overlies the graphics and most people think that this is the way to go, but because of the lag, the computer graphics look problematic. [...] I mean it really depends on where your eye is positioned in space and how you, so it requires a lot of customization (interview).

Mobile advocates often bring up the unique technological problems that HWD must solve in order to work, notably questions of eye tracking at a close distance, resolution and field of view, transparency and opacity of the optics, and integration of digital graphics with users' psychological percepts. This issue of customization directly undercuts the mass market goals, given the wide range of people that HWD would have to accommodate. Attempting to raise the degree of difficulty (or perceived degree of difficulty) of headworn problems works to emphasize the fact that mobile AR is already here, already works. One presentation at Augmented Reality Event 2012 entitled "Retinal Overlays: Should We Even Bother?" offered a thought experiment for the technological requirements that would be necessary to make headworn retinal overlays that work with the human visual system:

What's the nirvana of perfect visual overlays? [...] How much misalignment could a human eye detect? [...] In terms of a metric for AR, [...] .419 minarc, that's what diffraction tells us the limit is for typical pupil dilation. [...] But people can actually see better than subpixel vision for straight lines, [...] studies show that people can align to an acuity of 0.13 minarc [...] I think that is the limit for AR precision. Less than 0.13 minarc or error will be perceptible. (ARE2012).

The presentation was an estimate for the limits of human processing and concluded that any head tracking or object tracking system for retinal overlays would need to achieve this level of accuracy, in order to know what part of the scene a user is gazing at. Head tracking is still insufficient, however, because our eyes can turn and throw off the alignment: "If you don't track the eye at all, the error might be 10mm in the position of the pupil, so you'll only be able to have perfect visual overlay for objects that are 343 meters away or further, so if you want to overlay things up close you're going to need that (ARE2012)." Even this requirement would be insufficient to keep up with the lag that a human eye could perceive: "The eye can saccade at a rate of 900 degrees a second. You can traverse .13 minarc in 2.4 microseconds. So if we want to

never fall behind .13 minarc, we need to track at 415 kHz. Most cameras are 30 frames per second (ARE2012).” The brain processes the overlays as part of the visual scene, so if the augmentation is misaligned with the physical world, that could lead to accidents and injury to the user. And all of this assumes that the system could know what a person is looking at – Greg, a CTO at one of the mobile browser companies, explains: “I don’t know how the glasses can solve the problem of variable focus. A users’ eye could stay exactly the same, but the system doesn’t know if they are focusing one meter in front of them or 10 meters in front of them.” Giovanni, a user experience researcher, acknowledges that this is a major problem: “The fact [is] that focus is still a problem in head mounted displays. Unless somehow we can determine what the user is focusing on. It’s quite complex (interview).”

The precise accuracy of these estimates is not important. Rather, these kinds of heuristics and rhetoric of ‘should we even bother’ are examples of elevating the problems/requirements of headworn AR (in this instance retinal displays). These arguments are about the scientific feasibility of the systems that serve to amplify the barriers to reaching a HWD future. Mobile advocates do not even necessarily need to demonstrate that it is impossible, only that the gap that needs to be traversed is still very wide and that the time to reach that is prohibitive. This argument, paired with the one made in Ahonen’s keynote (that the next few years are crucial for AR, evn though there is typically no warrant given for this, more a declarative statement of fact), may be enough to stem the transition toward HWD. These questions of technological possibility are ones that are currently not dealt with by the existing HWD on the market now (e.g. Glass, Vuzix, etc.), because those do not engage in eye-tracking or retinal overlays. On the mobile side these are technological requirements that are imposed by a particular vision of HWD AR as well

as a particular definition of AR (see chapter 4), but they play a major role in raising the bar for HWD.

Another place where mobile advocates highlight technological flaws with HWD are the input designs and how users will interact with the device. Here's Walter, a CEO of a major software company, expressing concerns about haptic devices: "It seems like haptic devices, like a haptic glove, certain things just haven't taken off. [...] Asking users to put on more things in order to use AR seems quite unlikely (interview)." Others like Rodney, note that the voice command system for Google Glass is also problematic: "Too slow and personally it's too embarrassing to have to speak out demands. [...] And it is limited in how it can interact with augmentations. That is what is going to hold it back (interview)." Here the contrast with the status quo is clear: the touch interface that users of mobile devices are familiar with means that mobile AR does not have to develop additional input methods like haptics, gestural interfaces, voice inputs, or head tilt recognition, each of which bring with them another set of technological problems and additional demands on the user. Mobile advocates also point out hardware limitations such as battery life, which was described at AR Summit as perhaps the most significant barrier for HWD. Cary laments that with her Glass: "There's not even enough battery life to take video calls only for five minutes. [...] It's mostly like a very fragile device that you can wear on your head (interview)." This hardware claim stands in contrast to the support that the mobile coalition has from hardware and chip manufacturers, implying that these powerful actors will be able to solve the technological issue for mobile devices.

All of these technological projections work to present mobile AR as the most attainable, reasonable, and logical choice for the industry moving forward, because the existence of mobile phones already brackets off many of these specific concerns. Mobile AR avoids so many of the

technological problems of headworn devices, such as the issues of display. Using mobile AR allows people to see their augments at a distance and through a screen, something that an increasing number of users have experience with. While making the alternate means of input like voice, gesture, or haptics seem difficult and ridiculous, mobile advocates highlight the touch interface as an advantage for mobile AR, that these devices already have a method for user interaction.

Technological Futures (HWD)

HWD advocates, in response, believe that these technological issues are solvable. Many of them are academics and engineers who have been working in the space for years, and have a ‘can-do’ attitude when it comes to tackling persistent design problems. One strategy is to point to companies that are releasing commercial products. Vuzix and Epson are companies that already have consumer HWD products out on the market, and there is substantial enthusiasm for devices that are in development such as Google Glass, Meta, GlassUp, and others. While, as noted above those do not necessarily address all of the criticisms that mobile advocates raise (and may not even be considered ‘AR’ – see chapter 4), at the broader level the headworn camp can hold these examples up, both to point to ongoing corporate development and to refute the arguments about ‘impossibility.’ At these conferences many of these companies have booths for displaying HWD prototypes as well as a number of panel sessions in which representatives can make their case. While they may not have the same industry hardware support as mobile, they have their own institutional allies to point to.

Another strategy is to point to academic credentials and science as a source of authority to refute claims of impossibility: “I think it’s all possible and I work with some really smart

people, we have a guy who's a Ph.D. in astrophysics. These are the kind of guys where they wouldn't accept that something is impossible (Mike, interview)." Brian, an associate professor of computer science, is optimistic that they are very close to solving these issues: "I think that most of the really crucial algorithms work, is close to being solved, tracking, detection, mapping, rendering lighting, rendering shadows, even in painting for moving objects. That stuff is pretty robust. [...] It's still at a point where it's an engineering problem, but [...] There are so many things where you really need headmounted display." Academic laboratories and centers for eyewear research are also frequently mentioned as places where these problems are being addressed.

Some of these places have also begun to conduct research to directly refute the claims of mobile advocates, in particular what the 'real' requirements of a working HWD system would be. Researchers acknowledge that headworn AR systems will need to work with our psychological percepts, but believe that these estimates for 'perfect' overlays are exaggerated:

So, the main fear will be that we cannot achieve fundamental perfection. But we need to move away from the 'what you see is what you get' approach, because our visual systems have heuristics that address these issues. Things like focus and inattention blindness, in many cases that we look around the environment and see that things are not exactly as they seem. [...] There doesn't need to be an objective perfect, as long as AR elements can direct attention to certain tasks and communicate the task better. [...] Anything is possible when it comes to AR (Liu, interview).

Liu has been working on the psychological perception problem of AR for a long time, and bristles at the suggestion that some basic extrapolations about human perception could be used as a definitive claims about the requirements of a system. Those might be prohibitive for retinal displays, but ongoing developments in other types of HWD displays might be possible, or at least be able to account for those perceptual difficulties. At many of the conferences, HWD advocates count as their allies many of the leading academic laboratories and academics in the

space, and point to examples of their work to undermine the authority of others outside of their scientific space. On the technological question, these stakeholder groups are making appeals to scientific legitimacy to consolidate and secure their access to resources and authority (Gieryn, 1999).

HWD advocates also attempt to minimize the technological capabilities of mobile AR, as devices with limited computing power. Rodney, a programmer at a major eyewear company, points to limitations in mobile hardware processing speeds for tasks like computer vision and feature recognition:

To give you an idea of what kind of speed your processor needs to run it. Take for instance any given feature recognition algorithm. For the most common feature recognition algorithm, which is fast, you need to go through every pixel of the picture twice. That's one of the biggest drawbacks and you have to make a function of those pictures and make a curve for that function. Doing that in real-time is right now so painful on the processor that once you get that done you have no time to do anything else. You can't really do any I/O calls, drawing in 3D is also CPU intensive, so once you're finished with that you really can't do anything else [...] or you lose your real time speed. So processors on mobile devices are still too slow to do real time AR applications (interview).

The small field of view on these devices are also presented as a technological limitation:

A lot of those mobile devices are similar to the things we did many years before, [...] but unfortunately the tracking is not that good. [...] I have every single one of these AR applications on my phone but I don't use it because looking through a little peephole like this and with bad tracking there's a lot of reasons why it's not ready for primetime (Fred, interview).

While bringing up technological issues seems risky given that those are areas of concern for HWD, the move that actors are trying to make is that these limitations will persist for mobile devices, because they are designed to be multi-purpose devices and are tied to certain capabilities specific to phones that have nothing to do with AR: "The phones still have to be phones first, and the manufacturers know that. [...] You're not going to get someone sane to come up with something that has higher resolution than an iPhone or comparable display at that size because

people don't care and can't see it (Fred, interview).” Fred’s argument is that while mobile phones have many of the capabilities necessary for rudimentary AR, they are not specialized AR devices, and their multi-purpose nature may work against necessary technological developments specifically for AR. The more mobile advocates try to align with hardware manufacturers, the more HWD advocates attempt to saddle them with the commitments from that industry. One element of Fred’s comment is that these actors might perceive certain technological capabilities as ‘solved,’ and stabilize around those capabilities which are insufficient for AR. The other element is one of priority: that if these capabilities ever trade off with one another, these actors will privilege the core phone capabilities. These are attempts to turn back industry advantages, suggesting that their industry alignments will prevent them from ever overcoming or prioritizing certain technological limits. The lack of control for the CoP, and the uncertainty that these outside groups will optimize for AR tries to draw people back toward the HWD frame.

The technological visions on both sides are not just discursive phenomena, they are embodied in the research agendas of their supporters and the design of experiments. A commitment to a particular vision implies support and future benefits of that form, but also plays an important role in the devices used in experimental studies, as certain actors have begun to advocate for experimentation with mobile AR to further refine the technology (Wagner, 2007). They also play an important role in the research questions that are deemed important. HWD advocates need to work on the problem spaces like optical see-through and retinal displays, eye tracking, variable focus, but mobile does not. The contestation over the technological possibilities, then, is an attempt to motivate particular areas of work and align the priorities of researchers toward those areas.

Market/Adoption Futures

The second question that mobile and HWD advocates clash over is the market factors and adoption potential for AR with their device. This area is hugely important, considering that market factors and real life users can create conditions that support a particular form and push development towards it. Thus far, however, neither side has picked up users in a way that would exclude the other (and many HWD are still not publicly available), thus this contestation is taking place in the form of speculation on a variety of factors, from hypothetical cost to future adoption projections to presumed user characteristics.

Mobile AR advocates frequently point to their infrastructural advantages in the existing marketplace. Smartphone/tablet adoption numbers are frequently cited as a warrant for supporting mobile AR, because consumers already have the devices and have demonstrated a willingness to pay for them. Recent reports suggest that 56% of U.S. cell phone owners have a smartphone (Smith, 2013) and that 22% of people in the world now own a smartphone (Heggestuen, 2013). Ahonen's adoption chart, then (see figure 4), becomes not only a possibility but also the blueprint for linking AR to that exponential growth. Beyond the adoption rates, actors point to existing infrastructure as a reason to support mobile: "You need other supporting technologies for it to make sense commercially. [...] How do I get that data to you? Now mobile phones have 3G, 4G, Wi-Fi, and people have data plans. [...] Without those things no one's going to make the data that you want to display and you couldn't deliver it to devices, you would just have a really attractive thing that you could only show a cute standalone demo (Fred, interview)." Mobile advocates also point out that beyond delivering data, things like "app stores" are already set up as places where developers can present solutions to consumers. Their argument is that the existing technological infrastructure already supports a mobile model, from

cell towers to data plan structures. These are huge advantages for mobile AR that most HWDs do not have, particularly in terms of how companies deliver those solutions and the places AR can be accessed. Without access to a data network, HWDs are limited to depending on infrastructure like Wi-Fi, which dramatically limits the possibilities for AR development, the attractiveness of their device, and the content that is ultimately available. Because many users already have these mobile devices, they do not have to purchase another piece of technology, which mobile advocates argue will be prohibitively expensive to garner any mass adoption.

Beyond technological infrastructure, the powerful organizational actors that are responsible for the infrastructure are also referenced as a reason to support mobile AR. Mobile hardware companies and chip manufacturers are frequently present at these AR companies, and are another stakeholder group that could be an active driver of mobile AR:

Especially the mobile hardware manufacturers and carriers love this. Because it's using every single feature of the phone, and you always need stronger, more feature rich phone to be able to do the latest AR. Just like you know with PC's, they were driven because of games, not because of Word or Excel, because always the next game requires you to have better processors, same thing is happening here with mobile devices that's why they're so ecstatic about this. If you look at ARE 2012 probably twenty five percent of the people are from the mobile hardware companies (Aaron, interview).

Aaron was one of the organizers of that conference, so he had access to the registration affiliation information. Arthur, a researcher in the UK, has also observed this development: “I think the hardware is well ahead of the game at the moment in mobile. I think those manufacturers are going to be pushing quite hard in areas like AR almost to give a justification to manufacturers to put their chips in the devices. I think that will be one of the key drivers (interview).” Actors point to powerful organizations and other relevant social groups as the source of promise, as (1) a potential solution to existing problems, (2) a market unto itself, and (3) a way of narrowing the scale of potential problems to solve in order to achieve AR (e.g. more

powerful mobile processors). While Fred was skeptical of their motivations, mobile advocates try to reassure people about their priorities by explaining that these groups view AR as an essential feature of the phone, and that they will continue developing faster chips and better displays. Stakeholders use these existing infrastructural advantages and organizational actors to paint mobile as the most logical option, in an effort to recruit supporters.

In addition to these affirmative promises and expectations, the mobile AR coalition draw attention to those market indicators that are negative signs for HWD. In both public addresses and private interviews, actors raised doubts about the current and future price point of HWD, which could range in the hundreds or even thousands of dollars. Mobile advocates argue that present costs notwithstanding, the costs of HWD need to come down, while simultaneously their designers must address the technological problems that face the HWD space. This both makes assumptions about what users would be willing to pay for a consumer AR product, while also pointing out that the cost would exclude many potential users. This is held up in contrast with the massive and growing market for mobile devices.

Besides price, mobile advocates point to a number of social issues that are prohibitive for HWD. Bran explains: “There’s going to be I think a nerdiness factor that deters people from wearing digital eyewear, that they are going to be clunky and ugly and not all that functional (interview).” This nerdiness claim is one that is often made, and it has very specific connotations and issues of identity associated with it. By arguing that digital eyewear projects a ‘nerdy’ image, these represent attempts to limit the user base of HWD by casting them as ‘nerds’ in the pejorative sense, while also positing that as a barrier to ‘broader’ adoption. Sometimes this is done tacitly through visuals and the choice of visual representations, where in presentations AR advocates will poke fun at the type of people who would wear these devices (Figure 9)



(Figure 9 – Image of AR User, www.augmentedrealitytrends.com)

In particular these visual representations and comments explicitly attempt to link HWD to certain user attributes, specifically someone who is white, male, and socially inept. These representations draw upon longstanding cultural depictions of ‘nerds,’ especially the anti-sociality and negative meanings associated with computer users (Kendall, 1999). The implicit argument here is that certain users will only engage with computing in ways that are societally acceptable and normalized (e.g. mobile).

Another concern that mobile advocates like to hammer home is the social signal that wearing HWD sends out in terms of class. Here’s Murray: “If you wear Glass, you’re saying that you have enough hubris to wear a \$1,200 device that goes beyond hipster, and you wear it as a piece of jewelry to show that they are hipster [...] not the capabilities that it delivers to them (interview).” These concerns about social presentation surrounding consumer electronics are not new. Mobile devices are also fashion statements in their own right (Katz & Sugiyama, 2005), but what opponents of HWD are highlighting is how obtrusive HWDs are in displaying this identity. Here’s Murray again: “Because you see that status characteristic right next to their eyes, right in your line of vision when you’re interacting with them. That has not happened with any status characteristic. [...] They’re forced to look at it.” Some HWD opponents have started making claims about how this is something ingrained in our species’ development: “Human beings are

extremely sensitive to faces. There's even an area of your brain that lights up when we recognize faces. So when there's something there, we are troubled by this (Dennis, interview)."¹ Mobile advocates extrapolate from this argument about human perception in an attempt to raise the bar for the HWD design to be something either unnoticeable, or a device that people already wear.

Lastly, in addition to linking the technology to certain values of the user, mobile advocates also raise aesthetic, ergonomic, and user interface issues that could affect the adoption of HWD devices. First, some simply believe that it is a mistake to place all of these features on someone's head:

Eyewear focuses on putting the processor up where the display is, up by your face. [...] It does all of the processing on a little dongle that comes down to a wire from their glasses. I really don't know [...] if the processors and graphics cards are able to be headworn and still be robust. You've got to consider heat issues, power dissipation, how much hardware can you possibly fit on a headmounted or a headworn apparatus? Are you going to have a 3G antenna on that? It's very much in question as to how much hardware can people squeeze on your head until people are satisfied with the augmented reality those devices can do (Rodney, interview).

Frequently I heard mobile advocates talking about the tradeoffs between form and function with HWD, where the more functionality it might have, the heavier, bulkier, and uglier that might be. At AWE NY, a straw poll by a speaker found that only 25% of people in the audience would actually wear a Glass device. His point was to demonstrate to the audience the lack of public adoption potential for HWD, given the low numbers even in a room full of AR enthusiasts who are already enthusiastic about the new technology. The takeaway message from that presentation was to 'never underestimate the power of vanity,' as a warrant for turning away from the consumer HWD AR market. These attempts to broaden the discussion to other issues of identity

¹ What Dennis is referring to is the Fusiform Face Area (FFA), which has been shown through functional magnetic resonance imaging (fMRI) studies to become more active when faces are shown (Kanwisher, McDermott, & Chun, 1997).

is an example of an anticipatory debate, to gauge the extent to which a particular technological option will meet popular resistance (Hedgecoe & Martin, 2003).

In these debates, the move to declare that HWD AR must take a specific form ('regular glasses') represents an interesting tactic, as rarely has the design of a material technology or device been so predefined in its ideal state, with little room for modification, and with even opponents adamant that this design is necessary for adoption. Contrary to what one might expect, the ideal form of glasses is brought up not necessarily to simplify the design: in focusing on glasses, mobile supporters are attempting to redefine and heighten the problem space for HWD and prelimit the interpretive flexibility possibilities for HWD. Here we see an instance where supporters of one form are not only using that vision to guide the research agendas of supporters, but also to define the design and research requirements of their opponents. By linking the social psychological perception arguments against HWD to questions of aesthetics, mobile advocates make these seem intractable, which allow them to criticize all HWD designs that come out in relation to a common pair of (sun)glasses, and point out all the physical indicators of difference that appear in the prototype. To the extent that HWD advocates accept these particular framings and assumptions about users, their designs are structured around this goal of making HWD look the same as 'regular' glasses.

Market/Adoption Futures (HWD)

HWD advocates vehemently refute these claims, in particular the cost arguments:

The arguments are ones that can be met with answers. Like you say things are going to be too expensive. Unless it's something where you can convince me that we'll never be able [to lower it], like gold for example, [...] so if I say this is going to require a couple pounds of gold, it's going to cost a lot of money [...] So if you tell me that eyewear will always be expensive. I'm sorry I'm not sure they will always be THAT expensive (Fred, interview).

Cary sees this as an issue of implementation, which will inevitably be solved: “Now the whole point of it is that these are produced on a mass scale enough that more people can wear them and the retail price will be \$200 to \$400. It has to be affordable.” For them, cost is a concern that is presentist, and not something that will be an issue as long as they can demonstrate a need for HWD, which they would need to do anyway. With that ubiquity, they are hoping that some of those social identity issues will go away, much in the same way it has for mobile: “social norms change. In the 90s if you saw someone walking down the street with their head [down] in a phone, you would have thought that was crazy. [...] Norms around eyewear can change too.” For them, they do not accept these predefined reactions to technology, instead arguing that technology can change those.

HWD supporters present themselves as unconcerned about the design questions, bracketing that as a question of industrial design:

As long as it doesn't look too dorky where that's a moving target, then it gets interesting. Now you got people who are talented designers who might be able to take a number of different designs that folks have created, [...] you could take some of those designs and turn that into something that not only worked with a reasonable field of view, but looked appealing [...] That doesn't take a technologist to do, it takes a really good industrial designer (Fred, interview).

Actors also point to a number of prototype designs that might be appealing to prospective users (see Figure 10). Implicit in these, however, are signs that HWD advocates accept the constraint imposed by mobile advocates, that HWD devices must be analogous to sunglasses, and instead try to claim that HWD can mirror objects that people have already shown a willingness to wear (e.g. sunglasses). At ARE 2012, one of the eyewear companies boasted that they were consulting a brand-name glasses company to aid them in their design, which was both an example of how expertise from another field could solve potential problems as well as an admission that they are

trying to base the technical design around the aesthetics as a constraint. While some of the early prototypes like Glass have veered away from these models, the acceptance of this frame by HWD advocates could ultimately push designs toward this direction.



(Figure 10 – GlassUp, www.indiegogo.com)

Actors often bring up market reports that suggest potential for adoption of HWD. Walter, a CEO of a major AR company, cited in our interview a Forrester Research survey that 12% of Americans, or 21.6 million people would buy a pair of Google Glasses if they were available today (Epps et al., 2013). As a counter, while HWD advocates do not dispute the rising numbers of smartphone adopters, they instead challenge the presumed link between smartphone adoption and AR adoption. They argue that a mobile device fundamentally limits its usability for AR applications for a few reasons. The first is having to hold up the phone. Walter, the founder and CEO of a major HWD hardware company, attributes that as the reason why mobile AR is not being fully embraced:

It's our feeling that it's just not realistic to walk around a city or wherever with your cell phone 15 inches in front of your face. So [...] the kind of applications they have now, if you want to find your nearest Starbucks you can have your mobile device, pick a layer, and see the buildings and the streets through your viewfinder and it'll say go up one block and turn left [...]. Single use applications or something that you pull out and [...] do something with, it's not really intended for [...] continual use. It's mildly interesting, [...] there's a wow factor but it

seems to fizzle awfully quick. It's just kind of a lack of something there (interview).

The second is that many uses of AR involve navigation or finding things in your surroundings, so users would have to move around holding their phone up (see Figure 11).



(Figure 11 – How a User Has to Experience Mobile AR, www.autoslash.com)

Because it might be difficult to move around while holding your phone up like this, HWD advocates argue that mobile AR usage is relatively stationary. At every conference, invariably someone takes a phone out, hunches over, and takes a couple of awkward steps forward as a way of portraying the unlikeliness that users will act like this. This happens quite frequently, and is a strategic attempt to project unrealistic user behavior with mobile devices. The underlying message is that while mobile AR may be mobile in the sense that it can be accessed in a variety of places, it does not allow users to be mobile in the sense of mobility in the act of using AR. These types of rituals demonstrate the importance of these meetings and conferences, as a place to demonstrate and proliferate tactics for promising and (in this instance) discrediting certain forms. Through these physical performances, HWD advocates are strategically attempting to construct this ‘unrealistic user’ behavior, with the goal of breaking the connection between the adoption numbers of smartphones with adoption of mobile AR.

They also point to the last five years as evidence that users do not wish to engage in this type of behavior, given that mobile AR adoption has not reached its lofty goals. Here's Will, a COO of a major AR company:

Right now, augmented reality has the highest attrition rate [...] for any category of apps. It's only one in ten users that download it ever actually return to the application. [...] So that's one reason why augmented reality is getting such a bad name. It's because it seems more like a gimmick and not as applicable to everyday uses (interview).

While there may be many factors that explain the lack of user uptake in the last few years (content, knowledge, awareness, etc.), HWD advocates place the responsibility squarely on the mobile device itself and the way users have to access AR as the defining reason people are not adopting AR. These represent a strategic reading of the past in order to set expectations for the future, specifically to counter the argument that there should be a transition from mobile to HWD. Rather, these futures are efforts to persuade the community that the development of HWD is the prerequisite to mass adoption.

Normative Future Visions (Mobile)

Mobile and HWD advocates also contest each others' vision of what the future will look like with these devices in them. Often, mobile advocates will talk about not just a particular application, but how mobile will be the catalyst for normalizing an array of AR practices. Here's Gerald, co-founder of a mobile browser company: "This is how we start with the masses, how my mother-in-law in Orange County will start using it. Because in "Good Housekeeping" she will see an ad, or Oprah Magazine she will see an ad. She can download the book Oprah's recommended by just holding her phone over it (interview)." These types of anticipated future uses are quite common, with user characteristics embedded in those futures to convey that

mobile AR will be so ubiquitous that ‘even x person’ will know how to use it. These show up in many different places, with the most common visions for mobile AR being one in which any person could point their phones at something, and see something either touristy, artistic, historical, or educational.

While the futures presented for mobile AR are relatively simple and generic, probably to emphasize their accessibility, mobile advocates are simultaneously attempting to load HWD with specific contexts and future visions of AR that might have negative societal consequences. Bran, a lawyer working on issues of AR, brings up the private viewing condition inherent to a HWD, highlights the possibility for social misunderstanding:

Imagine everybody around you has a different digital experience and they’re all playing a game, pointing and clicking and shooting in different directions and running into each other it can get hairy. I think of what happened socially when Bluetooth headsets first became popular, where people were really thrown by seeing people apparently talking to themselves without realizing they have the headset on, not knowing to look for it, causing all sorts of social awkwardness. Yeah it will be the same thing when people are interacting with digital information (interview).

Walter’s point is that with mobile AR, there are still social cues that might orient others that someone is augmenting. With HWD, however, the form might allow for a future where AR technology affects social and societal norms in negative ways. This private viewing condition also might allow people to view content surreptitiously, and a number of actors brought up the concern that HWDs would allow people to watch pornography undetected. Lastly, the private viewing condition could magnify power asymmetries, as people could use HWD in ways that employ facial recognition or record without their knowledge. These particular depictions bring in specific relevant stakeholder groups, in particular women who may be subjected to stalking, general creepiness, or worse, enabled by the use of AR (see Figure 12).

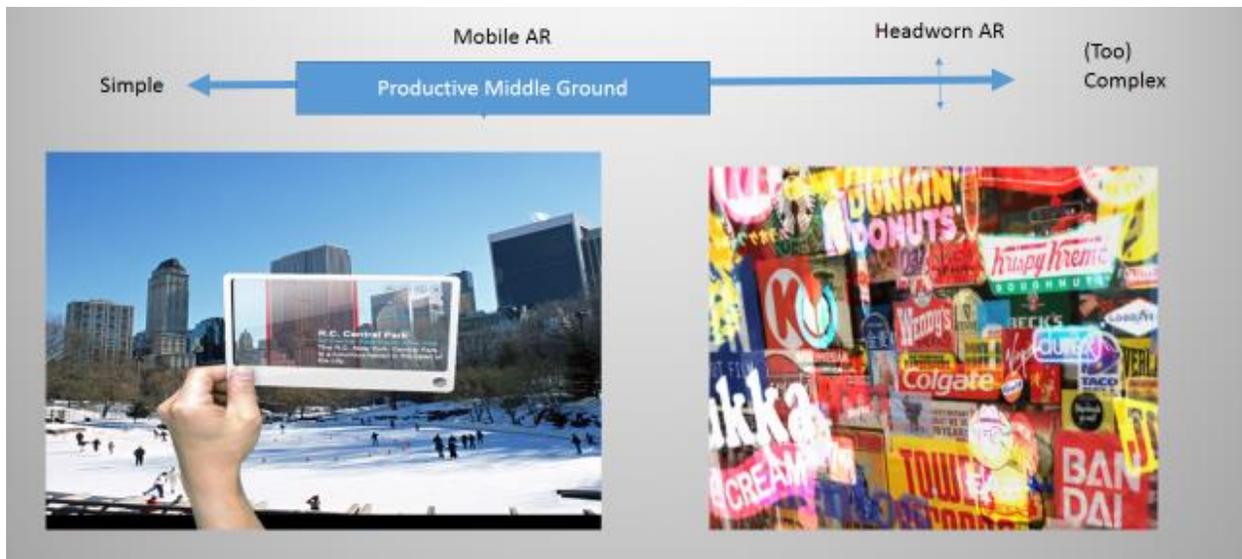


(Figure 12 – Facial Recognition and Social Networking. www.engadget.com)

Other depictions of the future try to turn the ‘immersiveness’ of HWD from an advantage to a disadvantage, by highlighting the possibility of overload and distraction. Because HWD claims to have a wider field of view, the question of mental overload and ‘how much is too much’ is more often deployed against HWD. Mobile AR advocates argue that people only have finite perceptual capacity: “Our perceptual systems are set by hundreds of years of evolution. That’s not going to change. At some point we’re going to keep adding things to the visual scene and we won’t be able to handle it (Liu, interview).” Coupled with cognitive overload are the future distraction possibilities of HWD AR. Mobile AR advocates point to dystopic visions of the future where AR is littered with advertising or people get injured wearing HWD. This is a discursive move to redefine a perceived problem with AR, specifically the lack of immersiveness with their screen display, and turn that into an advantage by arguing that more immersiveness may ultimately be undesirable. Finally, one future deployed by mobile AR advocates appeals to the concept of humanness, where developments in HWD might somehow cross a line in our

relationship to technology: “If we are able to wear these devices to see in infrared, ultraviolet light spectrums, then, are we still human? Or are we humanoid android things? [...] It’s certainly something that I think will become part of the human dialogue in the next two decades” (Dennis, interview).

These discursive attempts tap into historical, powerful fears about technology, specifically notions of information overload and our relationship to technology (Carr; 2010, Eppler & Mengis, 2004; Sturken, Thomas, & Ball-Rokeach, 2004; Turkle, 2011). It is somewhat surprising that even advocates who ostensibly support AR would heighten fears about a particular form, whether it is unintentional or strategic, but it happens quite frequently at these conferences. This type of modulation between warning and reassuring people about the technology is a specific type of work that is a unique characteristic of emerging fields. Marvin (1988) observed early electricians having to balance selling the danger of electricity such that those outside their community could/should not experiment with it, without overselling the danger in a way that would scare the general public. Mobile AR advocates are attempting to strike a similar balance between scaring people about the particular form of HWD and scaring them too much about AR generally, while also making the case that these potential negative futures could be bracketed or avoided by supporting the mobile form. Depicted visually, I find that while future visions of mobile AR are relatively simple, they are still battling over a productive middle ground while pushing future visions of HWD that are complex and potentially scary (see figure 13)



(Figure 13 – Spectrum of Mobile AR Futures. Created by Author)

Within this contestation, different stakeholder groups are attempting to coalesce around certain structural forces that suggest promise in a particular form. It is not any singular device or any one product they are attempting to advocate, but there is a larger coalition to advance mobile as the preferred form for AR. These members include those who are already enlisted in support of the mobile form by their organizations, as well as other 3rd party developers, enthusiasts, and newcomers speculating on the form. The goal is to align people toward the form and teach people how to make these promises, through a particular construction of technological futures, marketing/adoption trends, and future visions.

While this technological frame for pursuing mobile is used to structure the interactions of members, it cannot completely structure the community and various actors have differing levels of inclusion in the frame (Bijker, 1995). These arguments are not held universally across mobile advocates, rather they collectively map the different domains that this coalition of advocates contest to move development toward mobile AR. Some actors begin with the technological difficulties of HWD, because those are the prerequisite to any HWD support. Some actors might

grant the technological feasibility of HWD, they could still raise the bar for HWD in terms of design, aesthetics, and adoption. Others could concede both the feasibility and aesthetics of HWD, but express negative expectations for the supporting infrastructure behind the technology. Finally, one could grant that all of the technological, aesthetic, and infrastructural barriers associated with HWD could be overcome, but still express concerns for the future applications, privacy implications, and fears for what HWD represents as reasons to eschew that reality. Together these claims and promises could be woven together as an entire frame or selection of “even-if” arguments that ultimately support mobile device AR. These are specifically directed toward the undecided members of the community, either people who believe these forms can co-exist or the pragmatic people who believe that mobile will transition to headworn. Persuading them on any one of these domains might pull them toward advocating mobile or at least not initiating the transition.

Normative Future Visions (HWD)

HWD advocates present dramatically different versions of the future. First, they refute the claims that an HWD future would lead to overload/distraction, arguing that users will be able to adapt to AR displays and that the technology will ultimately make things less chaotic:

That’s something we will adapt to. [...] I think that we evolve and we consistently evolve as a species as a human race as technology evolves we evolve with it. [...] Long term its actually going to make people think less [...] because the computer is remembering for you so [...] it’s going to be less of an overload and more of a you know, piggyback mechanism (Rodney, interview).

These are lay theories for how people will adapt to AR, but these kinds of substantive visions that presume some fundamental change either in people or society in the future (Michael, 2000) represent one method of minimizing critiques of HWD that draw on the pervasive narrative of

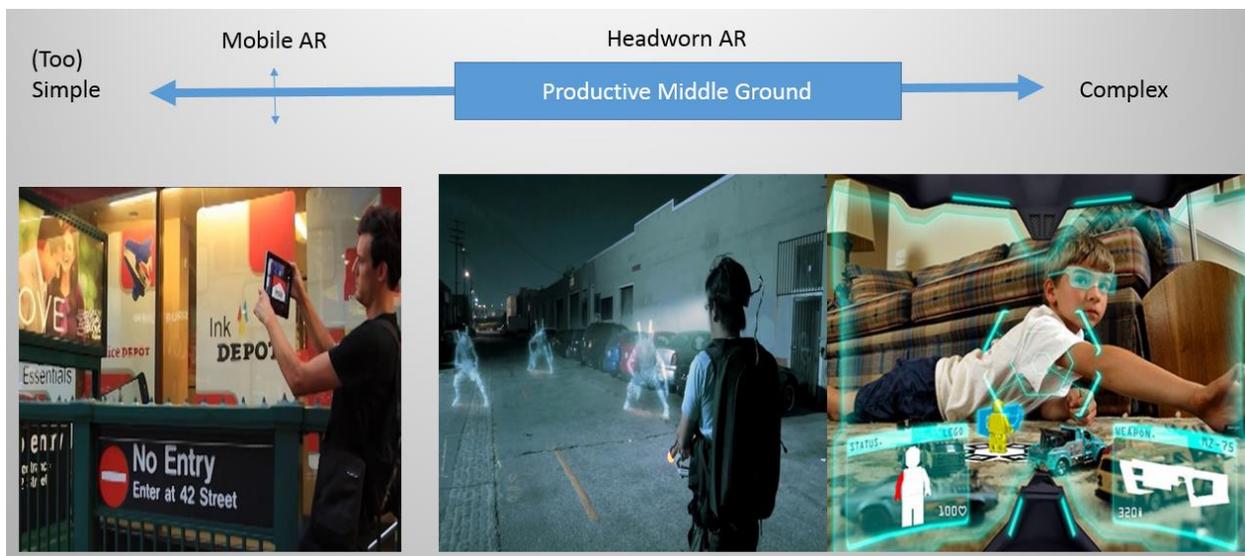
technology causing us to lose control, lose our autonomy over our surroundings, and lose a sense of humanness. Others turn the overload and distraction argument against mobile AR, pointing out that having a headworn system that is unobtrusive is likely less distracting than walking, holding a phone in front of you, and focusing on a mobile device. At InsideAR 2013, one speaker showed an image of someone walking into a padded street corner to illustrate this point, that we might actually be more engaged with the world with HWD. These are all part of the effort to redefine the problem space and stabilize the meanings surrounding the device.

In terms of the content of the futures, HWD supporters argue that their future subsumes all of the mobile applications while offering much more. The idea that HWD makes AR more immersive is commonly noted as a critical advantage of HWD. For HWD as consumer products, the visions of immersive gaming are prevalent, with videos being shown of AR creating elaborate real-time game environments. Contained in these visions are usually scenarios where the use of hands in interacting with the augmentation in a game environment is paramount, which is a strategic way of highlighting the lack of mobility and need for hands in using mobile AR. The ‘private viewing condition’ of HWD is also shown as an advantage for attracting users, highlighting how cool it is that these devices can provide the user more information about the environment. This becomes an area where the same technological design holds radically different future meanings for particular relevant social groups. Bijker’s (1995) example of this was the Ordinary bicycle that both posed a safety problem for non-users but attracted young men for the same reason, to demonstrate their daring and skill. The reasons that mobile advocates worry about the private viewing condition of HWD is an example of the same kind of interpretive flexibility across relevant social groups; the fact that a non-user might be concerned what a HWD user is looking at may be precise the reason that certain users want the device.

Lastly, HWD are unmerciful about the desirability of a mobile AR future. ‘Gimmicky’ is a common catchphrase that HWD advocates use when describing mobile AR. The fear is not only that these futures are mundane and boring, but that they are actively damaging to the cause of AR technology. Rodney, a programmer at an eyewear company, worries about how mobile AR represents the technology:

One claim that I often have to fend off is that AR is gimmicky. AR is only just smoke and mirrors, it’s just tricks. [...] A lot of what’s going on currently is like that. Worst case scenario that leads to the entire eyewear industry being scrapped. If people have adopted the idea that augmented reality is incredibly gimmicky and decide to stay away from it. That would be the ultimate worst case scenario (interview).

Rodney’s quote emphasizes how many actors are thinking about the stakes of this debate, as a zero sum battle that threatens the existential success of the technology. In the eyes of these advocates, a boring future (one that mobile AR advances) is the most threatening thing for generating excitement, enthusiasm, and tangible support for the technology. This is why their presentation of futures attempts to push mobile AR toward the mundane, while attempting to highlight productive yet complex uses for AR in gaming and entertainment (see Figure 13).



(Figure 14 – Spectrum of Headworn AR Futures, Created by Author)

Unlike mobile AR advocates, who can pick and choose which domain to contest and strategically concede with ‘even-if’ arguments, HWD advocates have a much more difficult task of having to rigorously defend their form. They cannot concede that HWDs are not technologically feasible, otherwise the whole point is moot. They also cannot concede the aesthetic or social barriers to adoption, because the crux of the HWD argument is that the form is what will be the critical enabler of adoption. Lastly, they cannot concede that futures with HWD are ultimately undesirable, because those are essential to rallying support for the development of these devices. They dispute mobile in all three domains that they are defending against, first by claiming that mobile AR is insufficient for meeting the needs of what AR ‘should be’ because of the technological display limitations, as well as casting doubt on the urgencies of the companies making the devices. The second contested domain is market/adoption, where these actors work to sever the link between device adoption and AR adoption, by constructing an ‘unrealistic user’ of the technology that is viewed as unrealistic. Lastly they express disdain for the aspirational visions of mobile AR as boring and mundane, while advancing fully immersive and ‘cool’ visions where the free use of hands is required.

These strategies do not just emerge in an ad hoc way, rather there are specific places where these tactics get discussed and honed amongst members of the coalition. After the AWE2013 conference was over, there was an impromptu session organized by two leaders in the HWD space to discuss ways to move the form forward. More than 40 people were in the room, and the session went late into the night. It was a clear example of how these conferences are critical places for organizing work in the community and developing the technology. This meeting offered one coalition of stakeholders the opportunity to not only define the critical problems they need to work on, but also a place to plan and circulate their rebuttals to critiques.

They divided up into three groups – one was in charge of tackling the psychological issues, others the technological issues, and lastly the societal issues they believe will confront the technology. When the groups reported back on the issues they came up with, several heated discussions broke out about the problems, how they might address them, and the degree to which those are serious concerns. One example came when a CEO of a major eyewear company outlined his concerns of the surveillance possibility of HWD. Other advocates immediately began to make the case that other devices (including mobile) did the same thing, therefore his claim was not unique to HWD, and that this could be a way of refuting those claims. This argument was an attempt to minimize the problem, and emerged at this workshop as one of the points of rebuttal that HWD advocates could make. It continues to appear in presentations in defense of HWD, as almost a year later at AWE NY a representative from a HWD company enacted this rebuttal on stage, telling a story about someone in the audience who was ‘going off on a rant about privacy,’ and how he turned on a recording device on his mobile phone to capture this persons’ objections to marginalize his complaint.

It was at these places where the goals and strategies for advancing and defending HWD were made clear, first to continually promise the possibility of HWD, a tactic of delaying demonstration and protecting the space for research to happen – which happens frequently during the promissory phase of emerging technologies (Bazerman, 1999; Fortun, 2008; Hedgecoe & Martin, 2002; Van Lente, 2000). In effect, their claims and appeals to scientific authority are promissory enactments of the dictum that ‘once the machine works people will be convinced’ – the flipside of ‘the machine will work when all the relevant people are convinced’ (Latour, 1987). These discursive acts are attempts to enlist support for the conditions necessary to make HWD possible, e.g. attracting investors, enrolling necessary allies, and preventing

people from supporting alternate forms. Frequently at conferences people will promise artificial deadlines of three, six, twelve months for certain technological problems to be solved or indicators to be met, all part of the ongoing promise of HWD. While they buy time for HWD, they have to openly criticize mobile as a viable option in the interim. Their strategy is to criticize mobile and continually promise HWD within all of these domains, and hope to keep supporters from moving toward mobile.

Much of the work in socializing newcomers to the community happens as both mobile and headworn supporters make their case within these three domains. It is not uncommon for speakers to start their presentations by explicitly acknowledging their position, where they might say ‘I’m not a believer in mobile AR (or HWD), because of x technological, market, or normative specific reason.’ Through these demonstrations, newcomers learn how to engage in and enact these contestations, by seeing how the community leaders walk, talk, and work.

Other more tangible instances where instruction happens are where leaders show people how to create their own AR experiences and applications through their developer tools, giving them access to the tools necessary to create AR while directing newcomers to a particular form of the technology (mobile). What they are contesting is how someone in the community promises the technology, by legitimating those claims as the masters (depending on which coalition they represent). Simultaneously there are attempts to teach newcomers the practices and tools for creating AR, show them what is possible, and present futures that are desirable and beneficial. Because the future of the community depends on not only the leaders but those who will become leaders, the contestation over the form and how newcomers are socialized into this debate is one that has implications for the direction and identity of the community itself.

Materiality, Design, and Ongoing Contestation

Within the broader context of these social and discursive moves, as the debate goes on, there are AR technologies which emerge that are worth examining. Because this has been going on for several years, this mapping of different domains of contestations is beginning to show up in the material development of the technology, and becomes a way of understanding the artifacts themselves as well as how the arrival those artifacts changes the contours of the debate for both coalitions. The awareness of the relationship between the defined problem spaces of the relevant stakeholder groups and the technology is one that many scholars recognize, but the specific and contingent ways that this happens highlights both the interpretive flexibility of these artifacts as well as how the broader strategic discourse shapes certain material features and design decisions (Bijker, 1995; Pinch & Bijker, 1987). They enter and alter the broader contestation over mobile or headworn in terms of their technological capabilities and design as well as in the ways that actors begin to point to these physical artifacts as material signs of promise.

In the ongoing development of the mobile devices themselves, designers have attempted to overcome particular purported technological limitations and market factors specific to AR. Mobile advocates argue that these are examples of how AR drives the technology as well as how new developments can drive mobile AR, particularly in how it gets integrated into their promissory discourse about the mobile form: “Phones today are getting really good, and at some point there’s going to be a diminishing return on how many features people need. Phone companies are looking at AR as the next thing that is the reason consumers will want to upgrade from an iPhone 10 to an iPhone 11 (Aaron, interview).” With these powerful stakeholder groups behind them and mobile chip manufacturers like Qualcomm and Intel prominently represented at these conferences, mobile advocates point to the latest developments in mobile technology to

make the case that they will solve the technological limitations that HWD advocates try to assign their devices. One of the most recent projects held up is Google's Project Tango, a mobile device that attempts to create real-time mappings of the space around them using computer vision, depth sensing, and robotics (Figure 15).



(Figure 15 – Google's Project Tango Video, <https://www.google.com/atap/projecttango/#project>)

In Google's video of Tango we see some attempts to engage in this broader discussion and rebut critiques of mobile. First, in the image on the left, we see an example of how Tango moves through the room from a first person perspective – but the device is free floating and omits hands from the image. This omission is part of the larger contestation about user behavior, as HWD advocates often contrast their hands-free capabilities with cumbersome depictions of mobile AR. In this instance, they simply removed it from the demonstration. The free floating mobile device is a common characteristic of promotional material (Metaio also has this as their promotional marker that enables an AR experience), so as to not call attention to the issue.

The image on the right is a diagram of the technological capabilities of the device, from the resolution of the camera to its multiple other sensor capabilities, all of which advance the technological promise of mobile AR. Mobile advocates argue that these are reasons to support mobile, in particular hoping to motivate people to develop on their device. Riley, a computer science professor, explained that the device got them excited and motivated him to start working

with the device: “We focus on the handheld tablet-based games when we were doing the game stuff because it's kind of a well-defined experience. When Tango came out [...] we got one of them. I want to do room-based versions of this. I think that for games, especially in constrained spaces, also has some interesting possibilities (interview).” This is also an attempt to co-opt the gaming and entertainment space that HWD is trying to claim, by broadening the scope of possibility for mobile. These kinds of advances spurred another round of speculation – at the 10th International AR Standards Community meeting, one of the participants Sid, an app developer, began excitedly talking about how AR will drive the design of mobile devices: “The age of the single camera phone is done. I think we’ll have multiple types of cameras and sensors for depth and mapping and maybe even infrared, which will allow for all kinds of AR applications.” This can be understood as a continuation of the battle for technological promise (soon AR will push technological features in x way), market/adoption futures (people will continue to upgrade phones with these capabilities anyway), and normative visions (imagine all the AR applications you could develop with these capabilities in place). It is simply the latest local development in the technology that intervenes in this debate, one of the many small cycles of innovation that drive a technology (Geels, 2004). In this case the artifacts are depicted in certain ways and stand in for arguments in the larger contestation over mobile and HWD, as actors hold them up to rally support, minimize critiques, and exclude rivals.

Design of HWD to Coopt Advantages

Beyond experimental devices, HWD AR companies have also made strides in creating artifacts that materially realize the possibility for AR, and their booths feature devices and prototypes that they hope will become the next consumer device. At ARE2012 the big draw was

that four eyewear companies were willing to have a panel session on a stage. By AWE2013 there were enough eyewear booths to claim a whole section of the exhibition hall. Through some of these early prototypes and in how HWD advocates are pushing them, we begin to see how they are embedded with features and design properties that strategically engage in the larger discourse of mobile versus headworn. In particular they stand in as arguments against some of the critiques from mobile advocates, as actors continue to work on not only the design, technological capabilities, and policies surrounding early HWD prototypes but also the work to make the form persuasive.

SeeBright

Two of the major advantages that mobile advocates have over HWD are the existing adoption numbers for smartphone devices and the accompanying data infrastructure behind mobile. The data infrastructure is not necessarily impossible to extend to HWD, but currently most HWD companies (except Google) do not have the partnerships with mobile carriers that allow them access to their 3G and 4G networks. The startup companies like Meta, GlassUp, just to name a few, are most disadvantaged, because they are too small to negotiate with those entities. Even some of the larger companies like Epson and Vuzix do not have deals with mobile carriers. These devices can only offer data through Wi-Fi, which limits their use to places where users have Wi-Fi access. This infrastructural limit is a disincentive for developers of applications, as well as potential buyers of HWD. It also shapes the scope of what is possible with these devices. With Wi-Fi alone, many of the public AR browsing applications would be impossible, instead the applications would need to be confined to particular spaces (e.g. gaming). To overcome these infrastructural limitations, one design solution comes from a startup called

SeeBright. At AWE2013, they were demoing a device that ran using a mobile device, but through a HWD (see Figure 16)



(Figure 16 – Seebright, www.technonauta.com)

This hybrid solution takes the two largest advantages for mobile advocates on the market side, adoption rates and data infrastructure, and simply coopts them to operate a HWD and solve the data transfer problems. This design concedes mobile's market advantages, but places that device into a HWD to access the benefits of the form (i.e. hands-free, immersive, etc.). This type of design demonstrates how the capabilities and power asymmetries of actors in a relevant social group shape the design of technologies (Klein & Kleinman, 2002). Here a smaller company in the space had to rely on existing technologies to access certain structural advantages. It also changed how they could promise the technology, and the uses that were possible. Because it required another piece of technology, one of their executives explained to me that the idea is not to have this be used all the time. There are also other limitations to relying on another device, however, most notably that it introduces multiple barriers to usage. First is that a user has to have a particular brand of mobile device (Seebright runs off an Apple iPhone). Second it adds another step to using a HWD, notably inserting a phone into it, which also has the effect of

limiting the users' access to their phone for the duration of use. Lastly, this particular deployment of the phone does not solve the input question, or change the technological requirements for eyetracking and other issues related to HWD.

This design is, however, a tangible instance where actors are trying to coopt advantages of mobile devices in their HWD design. SCOT explains that each stakeholder group has their own set of problems, as well as problems they want to project onto opposing forms. This design takes one problem (data transfer) and solves it in the most direct way possible, by using another device. This solution may not be acceptable to other HWD creators, however, because one of their arguments against mobile is that AR requires specific computing and graphic capabilities. The actors who attempt to cast aspersions on mobile technology and processing power would not concede that this device could deliver AR, particular what is required for HWD. It also does not necessarily solve the aesthetic issues with HWD, it simply limits the purported applications to the home.

What Seebright illustrates, however, is one possible way that a transition to HWD could happen: not with HWD replacing mobile but as another display option integrated with a mobile device. The design of this device was the result of certain system and infrastructural disadvantages, but with a HWD form. This also illustrates how these coalitions may be fluid, as the co-option could cut both ways. Some mobile advocates could "defect" and start using their devices to power HWD, to broaden their AR capabilities and capture a market. Other possible hybrid designs might involve using the mobile device to control and provide inputs to a HWD, which would be one way of bridging the interaction problem for HWD. These would further complicate the broader mobile versus HWD discussion, and could result in a solution that enlisted actors from both sides. At the same time, Seebright offers one example of how there

may be tradeoffs to this approach. They could heighten the learning curve required to use the devices, further narrowing the possible user population for AR exclusively to smartphone users, and limit usability to situations where users are carrying both devices.

Design and Regulation of HWD to Rebut Critiques

Google's Project Glass

Google's announcement of Project Glass in April 2012 was one of the first consumer HWD products to receive significant public attention, so it is no surprise that they have been a bellwether of sorts for the issues that might confront HWD. This was an important moment for the HWD coalition, as it represented interest from a major technology company and another step toward a consumer product. When explorer editions were released, the material capabilities of Glass not only changed the debate, but seemed to respond to these broader discourses and concerns. Glass was one of the first HWD that made possible certain heads up features and applications, but in doing so also elevated certain fears and critiques of HWD, which they attempt to deal with through not only design but also policy regulations.

One capability of Glass is that it allows people to take photos and record with the device. This enables certain uses and applications, but has also heightened controversy over the device. Beyond concerns that Glass could be used to record people without their consent, it also allows people to record by directly looking at someone and through a private viewing condition. In a tangible instance where design decisions are made to address larger discourses and critiques, Google designed in safeguards, such that the wearer had to either verbally instruct Glass to take a photo or physically tap the device. They also built in a light that turns on when the wearer is recording with the device. There is no functional reason for this light for the wearer, rather its

purpose is to indicate to other people that the Google Glass is being used by an individual. These design elements specifically address those concerns about surreptitious actors and the lack of social augmentation cues inherent to a private viewing condition, raised by mobile advocates and other opponents to HWD. We can understand this design as a technical “delegate,” a technical feature that compensates for the moral deficiencies of users (Latour, 1988; Pfaffenberger, 1992).

This particular delegate, however, has already been criticized as insufficient, as it a) does not address the fears that augmentations are private, b) replicates a social cue with one that might not be well understood, and c) can be circumvented. The circumvention of the intended design has surfaced as a counterdiscourse to Google Glass. First, an application called Winky was released that allows users take photos by winking, instead of the built in voice command or physical tapping. A number of developers have created workarounds to the light – one example, a company called Google Glass SunShade, demonstrates that it is possible to cover up the light. Others have found ways to circumvent the red light, either by altering the hardware or disabling it through software, further breaking down attempts to design away these critiques of the private viewing condition. These are examples of how users are attempting to ‘antiprogram’ the technology (Akrich & Latour, 1992), to reverse the intended actions of a design. These design moves illustrate how actors on both sides are attempting to situate the technology within these larger discourses, and are considering certain relevant social groups (in this case non-users of Glass) in the initial design and modification of HWD. These counterdelegation efforts could heighten certain fears that people have of the technology, and serve to “challenge the mythos of regularization (Pfaffenberger, 1992; p. 303).” With an emerging and disruptive technology like wearable devices, this will be an ongoing contestation that could greatly influence the normalization/adoption of the device. Glass demonstrates that although Google considered these

social problems and tried to minimize them, the problem cannot be stabilized through technological solutions alone.

A second design decision that Google made was to place the display out of the center of a users' vision, instead having the screen located above the eye and to the right. Greg, a Google Glass Community Leader, explained that a major reason that they placed the display up and to the right was because, in pilot tests, they found that a centralized display was creating headaches and people did not like it. This was the practical reason for the design, but it now changes how the device and its functionality are promised and marketed to people: "It's there when you need it, out of the way when you don't (Greg, AWE NY 2014)." This intentionally scales back on some of the promises of immersive HWD, instead positioning it as an unobtrusive device that the user has the choice to engage with. In doing so, this design is also being promised as a way to bracket some of the criticisms of overload and distraction.

This particular display has also begun to yield some interesting usage patterns, some intended and some not. Because Glass has been released to a select group of users through an 'Explorer' program, it is one of the first HWD to have actual users' feedback in the co-construction of HWD technology. These types of usability tests attempt to understand both the capability of the device and the user, in particular how the machine takes on meaning from its relationship to the user and vice versa (Woolgar, 1991). While the demonstration video seemed to suggest that Glass would be more or less constantly overlaid, one unexpected development from these tests was the way that users have been engaging with Glass. Greg, one of the Glass Community leaders, reports that at a recent Explorer meeting he asked if people had used the device for 30 minutes, 10 minutes, 5 minutes, 1-2 minutes, or 1 minute. The responses were that almost no one had used it for a period of time beyond 1 minute, and those who did only rarely.

Most of the interactions are about 5 seconds, making it more of a ‘glanceable interface’ for ‘microinteractions.’ These usability tests reconfigured the device through how users engaged with it, and ultimately its intended purpose. Greg, in his role as a Glass community leader, has been giving presentations to encourage developers to rethink the kinds of applications they need to create for Glass, because many of the ones that require intense engagement (e.g. video) as well as most reading/writing applications appear to be unacceptable to current users. At the same time, they have begun to raise questions about the kinds of data that would be necessary to design ‘glanceable’ interactions that are relevant to the user in real-time. The example given at AWENY 2014 was that if Glass wants to tell you once you get to the airport where exactly your gate is at, that might require a script to data mine your email to know that you have booked a flight, link that to time and location, and sync that with real-time airport data. The privacy implications of these applications, as well as the degree to which people want to give that up or would find those creepy, is something that is currently a major problem for Glass. This is an instance where the full spectrum of interrelationships between technology is important (Boszkowski & Siles, 2014), because it is the combination of the production of the device, the materiality of its design, and the early user engagement that has shifted the content now being considered for the technology. The design of the artifact has spawned certain types of user behavior, which in turn shapes the development space, which opens up questions about data, how Google wants to leverage that data, and issues of privacy.

Thus far, the goal has not been to encourage different or more lengthy engagements with Glass; Google is even attempting to solidify these uses by warning Explorers about extended use: “If you find yourself staring off into the prism for long periods of time you’re probably looking pretty weird to the people around you. [...] So don’t read War and Peace on Glass (Explorers

Do's and Don'ts).” In this instance the usability test revealed a particular use, which caused the company to shift its focus and categorize more intensive and persistent uses as bizarre or unacceptable. This not only tries to address fears that people might ‘Glass Out’ (i.e. zone out with the device) and send negative social signals, it draws on the feedback from these usability tests to inform application designers and what is possible with the device.

These design decisions have sparked much criticism of Glass in the AR community, as some denigrate it as a mere ‘notification device,’ not something that delivers an AR experience. At the same time, how these HWD will actually be received by outside people and institutions is also contested, with a variety of factors (e.g. the context it is used in, the behavior of other bystanders, and the extent to which people perceive recording) possibly contributing to bystander acceptance of HWD recording (Denning, Dehlawi, & Kohno, 2014). The specific capabilities and designs of HWD begin to entangle multiple relevant stakeholder groups: users, non-users, property owners, and private and public policymakers. Not only do users matter in how a technology develops; non-users and active rejectors of the technology matter as well (Kline, 2003; Wyatt, 2004).

Like the anti-automobile groups that destroyed roads to make them unpassable and the counties that banned automobiles altogether (Kline & Pinch, 1996), groups are now forming in opposition to HWD like Glass. Many object to the recording functions of these technologies, most notably bars (for social reasons/protection of non-users), casinos and banks (for security reasons), movie theatres (for copyright reasons), and schools/hospitals (for privacy reasons). In these instances it is a struggle over the property rights and interests of those establishments which motivate their opposition. These relevant social groups are choosing to exert their

authority over place to create conditions that are unwelcoming for the technology (and perhaps to deter potential users of the technology).

Besides the recording functions, the ability for Glass to add things to the visual scene contributes to the longstanding conversation about computing and distraction. A California woman was arrested for driving while wearing Google Glass under the statute that prohibited the operation of a motor vehicle:

If a television receiver, a video monitor, or a television or video screen, or any other similar means of visually displaying a television broadcast or video signal that produces entertainment or business applications, is operating and is located in the motor vehicle at a point forward of the back of the driver's seat, or is operating and the monitor, screen, or display is visible to the driver while driving the motor vehicle. (California Vehicle Code Section 27602).

The case was thrown out in court because there was insufficient evidence to prove that the device was operational, but it demonstrated that it could be illegal if the device were operational. This frames the question as one of classification: in the policy's terms, are Glass or other HWD a video screen? These types of legal distinctions and controversies that arise when the technology is deployed in everyday life play an important role in the trajectory of a technology.

As an exemplar of a failed technology, some have even started to draw direct comparisons to Glass (Pachal, 2013). Glass is the first HWD device to test these public boundaries, and the public policies and laws allowing (or banning) the device in certain settings and how Google chooses to deal with those stakeholders will be an important ongoing contestation to monitor. It also shapes the types of applications that are possible, e.g. any application while driving would be illegal. Right now, Google has offloaded some of that responsibility onto the Explorers themselves, by asking them to wear it in these places but advising them to act in a particular way: "If you're asked to turn your phone off, turn Glass off as well. Breaking the rules or being rude will not get businesses excited about Glass and will ruin it

for other Explorers (Glass Do's and Don'ts).” In these statements they both acknowledge that these select users will have the responsibility to negotiate with these stakeholders, and the idea that they opposition from these stakeholders could ‘ruin’ it for other potential users suggests they are aware that these larger policy and legal bans could be a structural threat to the future of the technology. The Explorer program, then, is both a way of testing out the technological design issues of the device (e.g. types of applications, how people are using it) as well as a program to help expose and educate the public about the device to quell unfounded fears. In doing so, the program also helps Google with testing, preempting, and anticipating some of the larger institutional, legal, and societal issues that the device raises.

Google is having Explorers pave the way, but are being quite cautious about the program. For the first couple years, they maintained tight control over the users of the Explorer program, as a list of people that Google chose to invite. At the same time, they are using marketing to construct a particular image of who a Glass user might be by showing representations of a diverse set of Glass users (see Figure 17).



(Figure 17 – Google Project Glass Press Photos, www.theverge.com)

Several visual characteristics of these advertisements of HWD stand out, particularly as an attempt to use marketing to rebut criticism from mobile advocates. The choice of users represent more than just attempts to create a market for HWD. In particular they can be understood in terms of the representation and associations they are trying to avoid with the technology. Part of these advertisements may be to refute the ‘nerd’ criticism, as certainly all of them are portrayed as hip, attractive, well dressed, models. Just as notable are the ways that these advertisements strategically depict women and non-white users of Glass. At a time where the technology industry in Silicon Valley has been criticized for its severe gender imbalance and white, male, party culture (Marwick, 2013; Miller, 2014), mobile advocates have attempted to confine Glass to a particular user group, most explicitly by citing and showing images from the blog “White Men Wearing Google Glass”² and linking these identities with asocial uses such as viewing pornography and spying on people. In these visions, Glass (and perhaps by extension other HWD) becomes a way for men to control space through surveillance and potentially abuse that technology, similar to fears of surveillance videos and operators being linked to fears of sexual harassment and voyeuristic monitoring of space (Koskela, 2002).

These advertisements, then, are not just efforts to publicize Glass, but are attempts to socially “unconstruct” a presumed user by replacing him with distinct others, with different attendant characteristics, motives, and aesthetics (Lindsay, 2003). These visions have gender, race, and class components, and are an explicit attempt to convey a particular user identity and particular uses that have very different connotations (e.g. using Glass to take video recordings of your child, as opposed to the stranger sitting next to you on the bus). This also is something that

² whitemenwearinggoogleglass.tumblr.com/

the AR community is generally trying to negotiate, both in terms of their own participation as well as who might also be interested in AR. The contestation over Glass is an important one for HWD AR, and important to prevent certain user identities or ones associated with the term 'glasshole' from stabilizing.

Beyond these marketing efforts, Google has taken tangible steps to preempt these critiques through policy decisions that prohibit developers from creating sexually explicit material for Glass applications. Whether they can successfully enforce that policy remains to be seen, but demonstrates how the form of the device gives rise to certain anxieties and dystopic futures that advocates of HWD feel the need to counter, either through user identity construction, design, or policy. Another policy that Google announced is that developers are banned from advertising on Glass. For the time being, developers are prohibited from serving or including any advertisements. This is an instance where a policy restriction is imposed not on the user, but on the stakeholders thinking about uses for the device. This is a radical departure from their business model, as Google's earning reports show that they make 96% of their revenue from advertising. Prohibiting advertising applications also limits the types of companies that might be interested in developing for the Glass platform. Putting this policy decision in the context of the larger discussion over HWD, we can understand this move as an effort to counter negative visions of the device that worry about unsolicited locational advertising that might be annoying and distracting. Some of the first parody videos of Google's promotional video showed unwanted advertising popping up on the screen, causing users to run into things. The question here is whether this policy is a long term decision, in which case Google would have to find an alternate economic model for Glass based on sales or data, or simply a short term move to insulate HWD from negative critiques in the early stage of development, but that could change as

the technology stabilizes and succeeds commercially. Either way these policies have implications for the content on HWD AR, the perception and discussion surrounding HWD, and the users of HWD.

Finally, Google has a policy banning facial recognition from applications. This is another restriction on the early group of Google Explorers who serve as the testers/developers/users of the technology. In this instance the policy restricts something that is technologically possible, again with the non-user of Glass in mind. It is intended to quell fears about Google Glass and the ‘creepy factor.’ At AWE NY 2014, Greg explained that this was a conscious decision on Google’s part: “Google has forbidden facial recognition, specifically because they are concerned how other people will feel. Also European law probably wouldn’t allow for it.” This reveals that the primary concern is to preempt fears of nonusers, but a secondary issue which an international company like Google has to account for are the laws of various regions. The larger conditions (e.g. legal, societal) in which users would be situated shape what the company will encourage and makes possible on the device. This is another instance where a possible technological promise of AR is removed through policy, specifically because the HWD could trigger critiques that associate the device with power asymmetry, privacy violations, lack of consent, and surveillance (Liao, 2012).

Despite all of these policies, there are still actors who continue to object to the device, what the device represents, or what they think the device is doing, and are motivated to remove/destroy them. Again, similar to instances when early groups opposed to automobiles physically assaulted automobile drivers (Kline & Pinch, 1996), there have been numerous instances where HWD users were physically assaulted or had their devices destroyed.³ These

³ <http://www.businessinsider.com/i-was-assaulted-for-wearing-google-glass-2014-4>
http://www.huffingtonpost.com/2014/02/25/woman-attacked-google-glass_n_4854442.html

incidents (and the media coverage of them) paint HWD are not just socially awkward, but artifacts that could open up the user to physical harm from other people. These incidents become part of the larger discussion surrounding HWD, and may alter the acceptance of HWD generally as well as people's willingness to wear them.

MetaPro

HWD advocates collectively need to fend off mobile AR, but they are still competing against one another to be the best consumer product. Meta is one such competitor to Glass, which recently launched a Kickstarter campaign for their AR glasses prototype and plan to launch a new product in the summer of 2014 (see Figure 18). Unlike Glass, the Meta advertisement takes a different tack to constructing the user, by simply omitting a representation of them. Instead it isolates the device against a black and white background. The video emphasizes the sportiness and sleekness of the design, while also avoiding having to choose a particular person to embody the user. The video is also shot in point of view style, which has the same effect. The design is also similar to a pair of eyeglasses, another example of a HWD company that is accepting the mobile frame for what these products need to look like and basing their designs on a pre-existing consumer product.



(Figure 18 – MetaPro, www.spaceglasses.com)

Meta also makes a strong appeal to scientific authority as a selling point and promises that they can overcome key technological barriers. They enlisted two leading AR academics, Professor Mann from the University of Toronto and Professor Feiner from Columbia University as their chief scientist and lead advisor, respectively. The other executives are students from their laboratories. The scientific authority and credentials of the founders is a prominent part of the discourse surrounding Meta, in particular to raise funds on Kickstarter, a platform for funding projects like emerging technologies. On the Meta website, their CTO explains his academic training in this way: “Steve [Mann] has a motto “demo or die” according to which his doctoral candidates need to produce a new iteration of their wearable computer every day or else they are out!” Drawing on the academic authority through these types of stories is a longstanding part of the promise for how the technological barriers of HWD will be overcome. In this instance it was a successful strategy for fundraising, as their CEO explained at an ARNY meeting that they were already well over their Kickstarter goal in less time than they had allotted. This was also part of the performance for the community itself, to persuade people that their company was a serious contender in the HWD space

There are also important distinguishing features of the MetaPro compared to Google Glass, which illustrates how even amongst HWD supporters there are design decisions made to determine who can better overcome the limits claimed by mobile advocates. It is in these differences that we see how these devices are attempting to move the conversation. One important design difference between the MetaPro prototype and other HWD such as Google Glass is that MetaPro offers a binocular display. This allows it to show 3D objects rather than just 2D objects, while also promising a wider FOV. In a direct comparison with Glass, they

advertise having 15 times the screen area. This design decision allows them to theoretically show different types of AR objects in stereoscopic 3D, which also taps into the vision of immersiveness and complexity for HWD. The design also allows for a different type of user interaction than Glass, through gesture recognition rather than voice, head tilt, or touch. This is an attempt to capture the promise of interactivity that is an important element of HWD AR visions, the ability to integrate the use of hands in applications and move augmentations to accomplish specific tasks. Besides the direct comparison in specifications, supporters of Meta have also been critical of the design of Google Glass. Mann (2013) writes:

I worry that Google and certain other companies are neglecting some important lessons. Their design decisions could make it hard for many folks to use these systems. Worse, poorly configured products might even damage some people's eyesight and set the movement back years. [...] The very first wearable computer system I put together showed me real-time video on a helmet-mounted display. The camera was situated close to one eye, but it didn't have quite the same viewpoint. The slight misalignment seemed unimportant at the time, but it produced some strange and unpleasant results. And those troubling effects persisted long after I took the gear off. That's because my brain had adjusted to an unnatural view, so it took a while to readjust to normal vision. [...] The current prototypes of Google Glass position the camera well to the right side of the wearer's right eye. Were that system to overlay live video imagery from the camera on top of the user's view, the very same problems would surely crop up.

In this quote we start to see criticisms between HWD actors, and attempts to link certain designs to larger health issues like eyesight damage and distorted vision. Mann also impacts this by saying that it could 'set the movement back years.' The implication here is that resolving these issues within HWD design and supporting the 'correct' design is critical to preventing those larger health problems often brought up by mobile advocates from becoming broadly associated with HWD. The MetaPro is not available yet, so it is still unknown how well they have actually address these concerns, but the design decisions already being made (e.g. binocular lenses,

gesture interface) are attempts to create a device that more fully realizes a particular vision of HWD AR, while trying to link their competitors' designs to physiological issues.

MetaPro also raises the recording issue, as a device that scans the virtual environment surrounding the user. Instead of accepting the frame that private recordings would be creepy, Mann has been arguing for years that these personal recording devices could be used for *sousveillance*, or inverse surveillance against authorities (Mann, Nolan, & Wellman, 2003). The argument here is that these devices *must* be embedded with these capabilities to provide people the necessary check against powerful actors. At an eyewear workshop for HWD at AWE 2013, Mann was adamant that the only way to prevent some of the social harms of 'unwanted surveillance and recording' is to possess the power to simultaneously record through HWD. Unlike Glass, then, which is attempting to limit and signal the recording possibilities, the supporters of Meta are embracing it as an asset. This is an attempt to redefine the problem space, not as a world in which this device adds another means of recording, rather that we are recorded all the time and the only way to respond is to watch the watchers. It remains to be seen how they will address some of the other criticisms that Google Glass is encountering, specific to the private viewing condition, the content issues, and other negative societal outcomes (privacy, overload/distraction, etc.).

At the AWE 2013 eyewear workshop, the case was also made for HWD advocates to embrace *sousveillance* as a principle to deploy when defending and inverting criticism of HWD recording. Framing recording devices as a check on power and authority might make the technology more desirable, even noble. When during the discussion one of the participants raised questions about whether *sousveillance* was really applicable to HWD, one of the leaders of the forum pressed him in return for viable alternatives – both to the problem he was outlining, but

also a viable response to this criticism that HWD supporters could make. When another person raised the problem of having a recording get taken out of context, the response was that HWD now allowed you to also record and prove that an action was taken out of context – the rebuttal to this claim would be to make recording about protection. There was ultimately no formal resolution to this question, but as an informal gathering of the HWD coalition, this meeting served a political purpose: to internally organize and coordinate their actions as well as circulate and debate possible rebuttals. This meeting was the most tangible form of coordination of these discursive responses on either side, but demonstrated the social, persuasive, and strategic ways that these groups are trying to influence and affect perception of AR technology.

These arguments are strategic and offer two possible ways to diffuse conflicts for the technology, but the way they are deployed is often as an oversimplification. First, the way they are attempting to minimize focuses on the technological act of recording, but what is actually qualitatively different about HWD is the ease, speed, covertness, and type of recording that is possible. In the example above where a HWD advocate covertly recorded someone on a mobile device, that was possible only because the person had set up their device to do so, and only captured audio not video. The ‘point of view’ video recording is something that HWD allows that mobile does not. Additionally, HWD may allow for sousveillance that could limit intrusions of privacy from the powerful, but could also open up new social groups to additional surveillance by unknown persons and entities. Practical questions of what these recordings might make possible and how these should be managed may also affect the enactment of these principles, like how these recordings would be stored, how long they could be kept for/retrieved, and who has access to these recordings. While these are questions that eventually have to be answered, for

now HWD advocates are simply trying to fend off fears of recording and hosting these groups to disseminate their rebuttals.

GlassUp

Two of the reasons mobile advocates claim that HWD will not find a mass market is the bulkiness of the object and how it works with things that people already wear on their heads. At AWE2013 I tried on a pair of prototype glasses from GlassUp, and two problems arose (see figure 10). The first was the width of the glasses and the inflexibility of the material, which made it difficult to put on. The second was that this design did not accommodate a person who already wears glasses, so I had to take off my existing glasses to put on the GlassUp prototype. This is a major problem that HWD design has to account for, as an increasing number of people in the world rely on corrective lenses. The representative at the booth suggested that perhaps a user could wear contact lenses instead. This solution acknowledges that their design excludes people who need corrective lenses but do not have contact lenses. For companies who want to include users that already wear glasses, one option might be to design a device that could be worn over glasses, which would minimize the problem but could potentially look quite bulky. The other option might be to fit the AR lenses with corrective prescriptions, but that might pose another barrier for adoption by being more costly to design, requiring users to update their devices, running into other systems and laws for medical providers, and being limited to ‘common prescriptions’ that would still exclude users with more specialized needs. GlassUp, as well as other devices that attempt to look like sunglasses, have these practical design and ergonomic issues to deal with, as well as how they will integrate with existing user needs and technologies (e.g. glasses, contact lenses). Devices that do not attempt to replicate existing products, however,

have more options for alternate designs. Google Glass, because it is off-centered, monocular, and a small screen, could potentially circumvent this issue. At AWE NY 2014, Greg explained that Google is partnering with existing frame manufacturers to figure out how to integrate Glass with those, as well as exploring an alternate design of a product that would clip to a pair of corrective lenses.

This is an example of how the particular problem space of HWD is constructed and defined by certain stakeholders, and how they are taking steps to solve it. Mobile advocates make the case that HWD need to replicate sunglasses in an attempt to raise the technological bar and bring in problems related to user customization. The HWD companies that accept that definition then must think of ways to either embed the customization into the device's optics itself, or offload that responsibility to the user to wear another type of corrective lens. The HWD companies that are attempting to create designs that do not mimic sunglasses have more options in terms of integrating with existing objects or as attachments to existing objects, but have the problem of falling afoul of concerns about the obtrusiveness and social awkwardness of these devices.

Conclusion

This chapter demonstrates how the 'trajectory of AR' is being contested broadly at a device level, with many structural, institutional, economic, and societal factors embedded into the design process and the debate surrounding the technology. At the design level, we see how certain visions and uses become embedded into the material form of particular devices. The gesture recognition input promised by MetaPro is an attempt to realize a vision of true immersiveness and interactivity with augmented objects, which might be necessary when using

AR to do certain tasks; Google, on the other hand, urges their users not to immerse, to merely glance. We also see design attempts that try to coopt the technological, infrastructural, and economic advantages of other devices, as in the case of SeeBright running a HWD using a mobile device. Mobile advocates attempt to predefine the HWD form, and certain actors attempt to meet that form but struggle with the customization. Lastly certain designs are made to rebut negative visions, like the Google Glass recording light, while other stakeholders regulate these concerns away through policies that ban pornography, facial recognition, and advertising. At the discursive level, all of these design elements are part of the larger coalescence of claims that HWD advocates advance and coordinate, to protect their development and interests as well as push back against alternate forms. The organizing spaces at these conferences are in part to allow stakeholders to gather, debate, and decide on strategies for rebutting these criticisms. These conferences become formal and informal places to recruit allies, circulate claims, and internally strategize about problem areas and how to address them.

Mobile advocates are similarly strategic in their work to advance their form. Some argue that HWD is technologically impossible. Others point to the ways that mobile usage is normalized, while simultaneously linking HWD to a particular user identity. Still others try to link HWD AR to negative societal outcomes and dystopic futures, specific to the private viewing condition and the immersiveness of the form. This is somewhat surprising, that even amongst people who nominally support AR technology, they are willing to raise fears about the technology to push it toward their preferred form. The ongoing material developments of the mobile form are explained as being driven by AR, as well as enablers of more complex and desirable futures.

New stakeholder groups and early adopters are emerging to adopt, contest, and tinker with the design and possible uses for the technology (e.g. hacking the recording light, circumventing the content policies). These actions, along with stakeholder groups who are preemptively banning HWD, can be read as part of larger societal contestation over certain users and uses of HWD. This battle is multi-sided, however, as HWD designs have to address not only these societal concerns but also distinguish themselves against competing forms of the technology, competing designs, and competing developer platforms. The specific design and user practices surrounding Glass also demonstrate the relationship between artifacts, social practices, and content (Lievrouw & Livingstone, 2006), how design privileges certain practices, which shapes the content and raises concerns for how that content gets created. All of this could have further implications for the use of the devices as well as how we understand the technology. This chapter is an early attempt to map this contestation, through which we begin to see how “the emerging character of a new technology [...] will depend on the constantly shifting relationship of actors and structures in both these domains [design and domestication] (Silverstone & Haddon, 1996; p. 46-47).” This chapter foregrounds issues that will continue to be contested in the upcoming stages of technological development.

CHAPTER 4 –

BUT IS IT AR?

“Perhaps the classic story in this vein was about the student who, asked to define electricity, said he used to know but had forgotten. “How sad,” replied his weary professor, “the only man who knew what electricity was has forgotten.”

(Carolyn Marvin, *When Old Technologies Were New*, 1988).

“What is AR? I know it when I see it. [...] It would be very hard for me to create a definition.”

(Dmitry, interview)

Introduction

Defining technologies is an incredibly difficult task, but one that has tremendous implications. Carolyn Marvin (above) illustrates that this is not only a common problem with emerging technologies, but also one that is fluid and constantly being drawn and redrawn. The AR community is struggling with a similar lack of clarity on this question, as people have very different ideas for ‘what AR is,’ even though AR is the common domain in which they all work. At one of the first meetings I ever attended, someone expressed an opinion, quite forcefully, that a particular application was absolutely not AR. This made a strong impression on me and was something that I worried about a great deal early on in my fieldwork, fearing that if I did not know the ‘right’ answer to the definitional question or called something AR that wasn’t, I would be discredited and found out as an imposter. Such is the challenge of attempting to enter a community organized around a technology, and learning how to speak about that technology in a way that demonstrated legitimacy (Lave & Wenger, 1991). It was only after many years in the AR community that I realized that this was not something you could have a “right” answer to – rather the debate itself, the intensity of it, and the perspectives that were advanced were indicative of a larger struggle over the technology, the future of the technology, and who was in charge of its characterization.

As mentioned in the previous chapter, Sutherland (1968) was amongst the first to describe and build systems that would later be categorized as augmented reality experiences. While he is sometimes referred to as the ‘Father of Augmented Reality,’ in actuality he did not refer to his systems as such. The term ‘augmented reality’ was coined by Caudell and Mizell (1992) to describe a system at Boeing that superimposed virtual graphics over physical space, highlighting cables so workers would know how to build certain systems. That particular usage was more of a description than a definition, in that it referred to an outcome and possible application of the technology without outlining criteria for what it entailed.

The significance of the term became broader than its origins, in part for how it resonated and got taken up by a group of people to describe and define their work and activities. The term ‘augmented reality’ has united a group of people who design, market, test, and utilize these technologies, much like the organizations that commit to protecting the cultural environment of an emerging technology space (Boyle, 2007), which gives it strength as an industry while also providing value as an organizing/marketing tool. Today, over 20 years later, a burgeoning industry that unites under the domain of AR is still struggling with the meaning of the term, while simultaneously creating the artifacts that enable its existence and industrial/commercial success.

While sometimes these are referred to as ‘augmented reality technologies,’ this is somewhat misleading. In actuality there is no device that is inherently an ‘augmented reality technology.’ For example, a mobile smartphone is not necessarily an AR device, but it has certain capabilities that might allow it to be used for an AR experience. What ‘augmented reality’ actually describes is the combination and convergence of many enabling technologies

(GPS, camera, compass, graphics, accelerometer, computational processes, etc.) that may allow for AR features, applications, and experiences (depending on how one defines AR).

In an effort to understand this definitional question (and as a practical matter, to avoid confusion later on in interviews), I began to ask people how they defined AR. While no one was flippant enough to say they once knew but had since forgotten, the answers covered a wide gamut of perspectives and approaches to answering the question. Some people bring up Hollywood depictions of AR as in the films *Iron Man* and *The Terminator*. Others might talk about example AR applications like Layar or Wikitude, or mention specific devices like Google Glass, GlassUp, Meta, or Vuzix. Still others might give you technical criteria for what makes something AR. Lastly, some will challenge the premise of the question, arguing that before we can answer ‘what is AR’ we had to first ask the question ‘what is reality?’

These answers reveal a larger dispute within the AR community: there are competing interpretations on this question, and the very definition of AR is hotly contested by different stakeholders in an attempt to reach a consensus. While there have been many linear histories of the term from one definition to another (Grover, 2014), they overlook the context and goals of those particular definitions as well as how actors are taking up different perspectives to rigorously enact a vision of AR. This chapter reveals how actors advance and defend different interpretations of AR, each of which attempts to move AR in a particular direction. I examine how the definitional debate can simultaneously be understood as a critical act of drawing boundaries around the community, a strategic act to embrace and distance AR from certain other technologies and devices, a way of guiding work and defining problems, a way of demonstrating expertise and membership in the community, and an attempt to persuade others to adopt their

definition as a standard. In other words, like all standards, the definitional question was politics by another means (Bowker & Star, 1999).

Almost as important as the question ‘what is AR?’ for understanding precisely where people draw the boundaries is the follow up, ‘what is not AR?’. The examples listed above fall within certain interpretations of AR but outside of others. As actors contest their inclusion, what they are really doing is legitimating certain applications, devices, practices, and the people who can be a part of the AR community. Many people working on AR believe the stakes of the debate are existential to the success of the community. For all of them, it determines their inclusion in the community. What they are really battling over is who has authority over the technology, the commercial/public success and perception of the technology, the resources they are able to attract and marshal toward the development and realization of their vision, and what the technology will ultimately look like. In the practices of scientific work (in this instance a subfield of computer science) actors must represent the domain in which they study and demarcate its boundaries, as a form of boundary work, a process that helps assert who has the authority and expertise to speak for that domain (Gieryn, 1983; 1999). One place this work contributes to our understanding of boundary work is as a mapping of a contemporary domain, where the questions are continuing and ongoing contests for credibility, as well as boundary work not just between scientists but with other industries. The dispute over ‘what is AR’ has tangible implications for motivating work and determining what technologies ‘count’, standardizing a language for talking about certain technologies, and classifying the technologies in ways that privilege certain perspectives (Bowker & Star, 1999).

The organizations that comprise the AR community often take positions regarding the definition. To understand this I draw from an organizational perspective, where there has been a

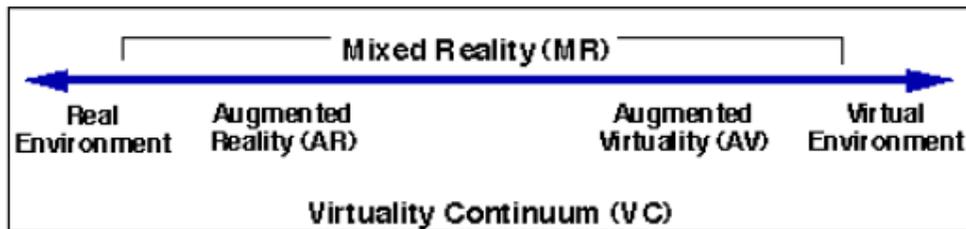
growing recognition that communication is what constitutes the organization itself (McPhee & Zaug, 2000, Putnam & Nicotera, 2009). This case examines an instance where the definitional question becomes all of those things, particularly as entire organizations form around certain interpretations and discursive constructions of the technology. This case demonstrates how actors are specifically contesting these definitional issues and communicating about them in ways that are constitutive of their respective organizations (McPhee & Zaug, 2000; Putnam & Nicotera, 2009; Taylor & Van Every, 2000). More importantly, however, these organizational perspectives help us understand some of the motivations and implications for the definitional positions that people and organizations take, and the implications for how that coordinates activities, structures the organization, negotiates membership, and positions the institution to outside actors (McPhee & Zaug, 2000).

The first half of this chapter maps the debate over the definition, while the second half examines how certain classification schemas are advanced to draw those boundaries around the field. Across a variety of stakeholder groups, this debate over the definitional parameters has implications for the CoP in terms of the organizations who work on the technology, the people who choose to engage, the problems they choose to tackle, and their ability to access resources, all of which will ultimately affect the technology.

Defining AR away from Virtual Reality

The first discussions about the AR definition occurred among academics, the earliest stakeholder group to research AR. Milgram and Kishino (1994) defined AR on a continuum with virtual reality, with reality on one end and virtual reality on the other. In between are two classifications of ‘mixed reality’, one which is ‘augmented reality’ – defined as physical reality

with virtual objects overlaying the scene, and the other is ‘augmented virtuality’ – which is a predominantly virtual environment with real objects in the scene (see figure 19).



(Figure 19 – Mixed Reality Spectrum, from Milgram & Kishino, 1994)

This definition was an important first step toward defining AR as a mutually exclusive subset of virtual reality (VR here referred to 100% virtual environments). It essentially re-drew the problem space of VR, with AR isolated as a separate area of inquiry. The distinction made here between each category is the proportion of real and virtual. This account does not delve into the specific developments in VR (see Rheingold, 1991), but instead focuses on how AR actors explained their actions in relation to the VR discipline.

AR became a subset of VR, but one that actively resisted the larger field. The definition of AR became one way to legitimate and distinguish AR as its own separate line of research apart from VR. Fred, a long time CS professor who has been working in AR for decades, retells the historical trajectory:

VR was becoming academically respectable in the early 90's. [...] However, shortly after that there was a big rush into commercial VR, some of it due to overhype by researchers and a number of companies that invested heavily in VR. The technology wasn't ready and some really subpar devices were pushed out to meet a quarterly business cycle. People didn't accept it, development slowed, and investment slowed as marketing managers felt burned. It no longer felt like it was cool to work on or the next big thing. [...] Soon many researchers turned to AR to strategically move away from VR, and it was being done by enough people that there could be a symposium workshop and conference in the late 90s (interview).

The conference that Fred is talking about was the ‘International Workshop on Augmented Reality’ (IWAR 1999), and the founders were explicit that they wanted to be distinct from VR.

The stated motivation and goal of ISAR was “for the international AR community to get together and concentrate on topics of AR research and applications without the ‘overhead’ of events where AR is only a related topic.”

The perceived decline in public popularity for VR was another important motivation for making a clear definitional move. Liu, another professor who has been in the field for many years explains: “Previously AR was more or less like a subset of VR. Gradually you see more and more people actually focusing on AR as AR becomes more popular (interview).” There is very much a sense of ‘jumping off the VR ship’ mentality with AR, both as a technology and as a discipline. Fred and Liu’s quotes represent the present day narrative about why AR researchers would want to distance themselves from VR (commercial failure that led to lack of support across domains), but they also highlight the community’s sensitivity to public perception and what stands as the exemplars of the technology. These concerns will resurface throughout this account.

Milgram and Kishino’s definition was the first step to clarifying AR’s independence. This definitional move setting AR in opposition to other domains of computer science can also be understood as an effort for institutional positioning (McPhee & Zaug, 2000). This “institutional positioning is vital to secure resources, support income, and legitimacy for the organization (McPhee & Iverson, 2009; p. 83).” Structurally and organizationally, then, the definition was used to create separate organizations and conferences that were only relevant to this particular area. ISAR eventually merged with another conference, and added the ‘mixed’ from Milgram and Kishino’s (1994) definition into its name and became ISMAR. This historical narrative also sets prescriptions for how actors should behave, as certain actors were deemed ‘responsible’ for the decline of VR, their umbrella field. Particularly in Fred’s explanation,

marketers, industry, and the venture capital ‘rush’ of funding into an emerging technology were to blame, because it rushed the technology onto a particularly compressed timeline. Also at fault were the researchers who ‘overhyped’ the technology, promising to deliver certain things to attract funding that they were then unable to produce. This perspective continues to guide many of the institutional actors in the organization, as a kind of origin story often told by its leaders. At ISMAR2013, this was reinforced again by the organizers: “ISMAR [...] doesn’t want to go through the VR experience. Researchers if anything have been toning down the hype to make sure that doesn’t happen.”

Not going through the VR experience also means not being associated with VR, and that original definitional move is still socially enforced. A student poster presentation at ISMAR2012 was questioned by a professor as such: “Where is the AR in this? Why isn’t this just VR?” Whether something is ‘just VR’ is an a priori question that academics hoping to enter the community must answer, to preserve their status as part of this special subset of the mixed reality spectrum. This kind of questioning by the leaders of the organizations is a demonstration of the goals that newcomers are expected to meet as well as enforce in turn. Jim, another professor of computer science, puts it succinctly: “you can publish AR research at VR conferences, but you can’t publish VR research at AR conferences.” Clearly, the definition affects the membership negotiation flow of the organization (McPhee & Zaug, 2000), as incoming members are socialized by long-standing members to proclaim their difference with VR. The narrative element is also important here, where contemporary decisions and justifications must appear consistent with how members make sense of past activities and decisions (Boden, 1994). Because the organization articulated its genesis as a move away from VR, many of the original participants feel that they need to consistently reinforce that distinction, and newcomers who

come into this organization are taught to enact that as a definitional value as well. This is how one segment of the CoP teaches and enforces a particularly definition and distinction of AR.

Azuma's Criteria Definition of AR

While the mixed reality spectrum was useful in drawing the boundaries in relation to VR, it was originally intended to be a taxonomy of the field, where the “reality” condition was simply one criteria. The categories offered were more about the proportion, and did not include clear criteria for defining AR. Additionally it came out of the VR tradition, which focused primarily on the technology available to VR at the time (e.g. video see through and head-mounted displays), and so did not account for a rise in mobile technologies, and inherited some of the assumptions about computing specific to VR (Barba, Mynatt, & MacIntyre, 2012). Further clarification was necessary to more clearly draw the definitional boundaries around AR.

In an effort to further clarify criteria for AR, Azuma's (1997) survey of AR technologies offered this definition of AR: “as systems that have the following characteristics: combines real and the virtual, interactive in real-time, registered in 3D (p. 356).” This definition of AR was also intended to accommodate a multitude of devices and sensory modalities, which makes it more expansive than any single technology and broadens it from Milgram & Kishino's (1994) focus on head mounted devices and visual AR.

It is difficult to overstate the impact of Azuma's definition and its criteria – by one measure of citations, Google scholar currently lists 3936 citations of the paper. On an organizational level, the academic community has coalesced around this definition. Sam, one of the members of the steering committee stated at ISMAR 2013: “the shadowy hand behind ISMAR, the steering committee, have been basing our definition of AR on Azuma's dissertation

for years.” For ease of application and understanding the debates, it helps to separate out Azuma’s (1997) definition as such: 1) Combines real and virtual, 2a) interactive, 2b) real-time, and 3) registered in 3-Dimensions.

Applying Azuma’s Definition to 1st and Ten

For researchers of systems or applications, tailoring their work to Azuma’s definition has become a way of orienting into the academic community. Underlying that, however, is a sentiment expressed in Dmitry’s quote at the start of this chapter, where people ‘know AR when they see it,’ which leaves more ambiguity in considering how to apply the criteria. There are instances, however, where this type of subjectivity simply will not do, and actors will apply Azuma’s criteria to real world examples and applications to determine if they are, in fact, AR. Within these debates, a number of different interpretations arise that lead to further segmentation of the definition. One example of this was the debate over the 1st and Ten system, which shows virtual lines on a televised broadcast of American Football (see figure 20).



(Figure 20 – 1st and Ten, Sportvision)

1st and Ten was one of the first applications associated with AR to gain widespread public recognition, as popular press depictions often reported it as such:

That yellow first-down line shown on televised football games isn't really on the field. But to viewers, it appears to be there, just like the turf and the players. The technology [...] is an early commercial example of a field of computer science called augmented reality (Berlin, 2009).

The system debuted in televised broadcasts of the National Football League in 1998 to approximate the yardage the offense needed to gain a first down. It soon became a fixture of all professional and college football broadcasts, and has been experimented with in other sports like hockey and some Olympic events.

Sportvision, the company that created the 1st and Ten technology, are happy to make the association with AR. Merton, an executive at Sportvision, explained:

I got there just after they put AR on television. [...] I think the first set of impacts has to do with providing augmented reality that naturally meshes with [...] understanding the game, understanding the rules of the games, understanding the progress of the game, and just basic enjoyment. [...] When I go and talk to another technology partner here, [...] *we'll talk about AR explicitly* and they have a very high level of awareness of [...] where it (AR) is right now (interview, emphasis added).

This association with AR, however, whether attributed by the company in an attempt to enter the community or the media, is contested by some members of the AR community. In one debate that took place on a listserv prior to an AR standards meeting in March of 2013, focused specifically on the 1st and Ten system, the participants began to apply Azuma's criteria to the system. Boris, an industry leader working on standards, argued that it does in fact combine the real and the virtual [Azuma #1], because the yellow line is virtually mixed with the reality of the football game. The result is a composite of virtual objects and the real environment. Riley, a professor of computer science, argued that "real" has to be the actual matter that is in a users' physical reality (e.g. the field itself if someone were at the stadium). Even within the prongs of the definition, then, those defining terms still need to be interpreted and contested. In his view,

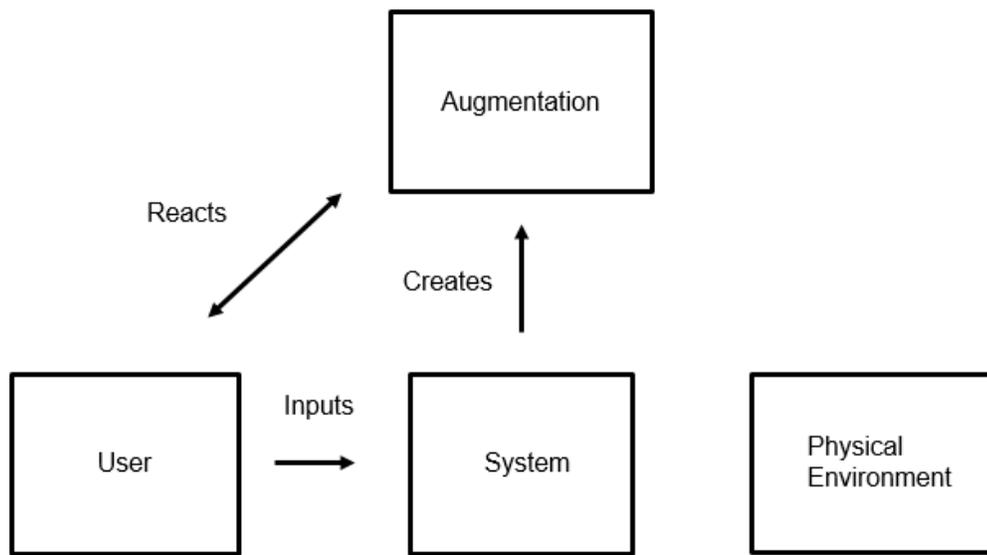
the 1st and Ten system shows a televised representation of a field with graphics on top, which they argue does not combine ‘real’ and ‘virtual.’ Under this interpretation, for it to meet criteria 1, the user would need to be physically at the game using a device to see a first down marker.

Some began to contest the restriction that it had to be physical reality. Samantha, a doctoral student, argued that excluding graphics as ‘real’ might be too restrictive, or at least arbitrary, since we are perceiving that thing as our reality. She counters with an argument that ‘real’ could be interpreted as whatever our eyes physically register. Riley rebuts that this expands the question of what is ‘real,’ broadening this to larger metaphysical questions that philosophers have been struggling with for centuries. This is deemed out of scope, and many argue that they should be bracketed from the conversation of AR. Others believe that defining ‘reality’ is a prerequisite to determining what the criteria of ‘mixing real and virtual’ means. This back and forth did not reach any consensus on these issues, rather it illustrates the range of interpretations one might have for ‘real,’ and how that might include or exclude certain technologies as ‘augmented reality’.

Azuma’s criterion #2a stipulates that AR needs to be interactive. The example given by Azuma (1997) is that 3D films (e.g. Jurassic Park) do not qualify as AR because they are not interactive for the user. Also these graphical overlays are not interacting with the real world, rather they are projected out of the screen. This definition is one of exclusion (e.g. what is not interactive), however, and is less clear on what something would need to be to constitute interactivity, and whether interactivity is a binary or a continuum.

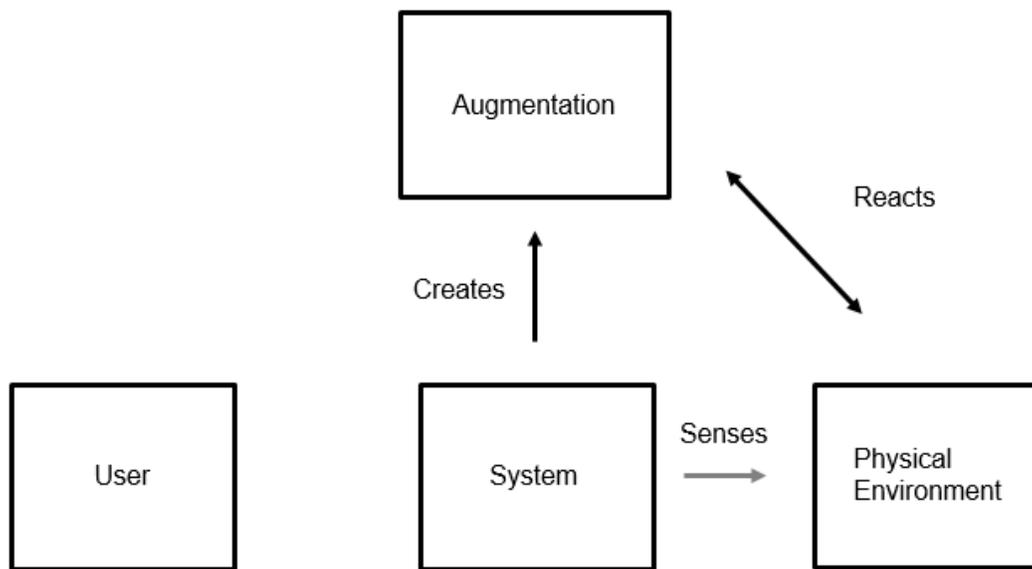
Once again, different interpretations were proffered. Riley argued that the 1st down line does not meet the criterion: “Ron [Azuma] (and I) mean interactive for the person for whom the experience was designed for. 1st&10 is aimed at the viewers of the event. The consumer of the

event has zero control over it. It's not "interactive" by any rational definition of interactive.” Here is a user-centric definition, one that makes an appeal to the author’s intent of the definition and argues that the augmentation has to react to the user input (see figure 21). The argument is that like a 3D movie, users do not control the 1st and Ten system.



(Figure 21 – User Centric Interpretation of Interactivity, Created by Author)

Others challenged this interpretation, arguing that it should not be about the user. The fact that the system produces virtual objects that react to the physical environment should be sufficient to meet the condition of interactivity. With the 1st and Ten system, the lines are programmed to react to players’ bodies that cross the line so as not to occlude live action. Because of this, proponents argue that it is ‘interactive’ with the real world. Here is a different interpretation of interactivity, that happens at the source level with the environment, not at the user level (see figure 22). Further, they argue that 1st and Ten is able to change camera angles, respond to instructions, thus meeting the user interactivity criteria for the operator of the 1st and 10 system.



(Figure 22 – System Centric Interpretation of Interactivity, Created by Author)

Some argued that a source based interpretation is better because user interactivity is subjective. Samantha makes this point by questioning how much user interactivity is enough, and whether simply activating the experience would count. Others argued via hypotheticals, pointing out that people are passive viewers of the 1st and Ten experience not by choice, but by the limitations of the television medium. If people could someday manipulate their own television views and the 1st and Ten, then that the system would meet the user interactivity criterion. Each of these different interpretations of interactivity attempts to draw limits around a range of technologies (real and hypothetical, current and future) that would/should be classified as AR. Alternatively, one might argue that user and system interactivity are both necessary to consider a device AR, which would further limit the categorical definition of the technology.

Azuma’s criterion #2b, ‘real-time,’ becomes another point of contention, as it is a construct of computing that refers to the speed of calculations that a system must make. Because any calculation necessary, however minimal, has an asymptotic relationship to the actual passage of time, the question is in fact always ‘to what degree’ constitutes ‘real-time.’ Boris argues that

1st and Ten met the real-time criteria, because there is no perceivable delay in the augmentation given that it adjusts to events happening on the field (e.g. players/objects moving over the line). This interpretation of real-time as ‘without perceivable delay’ makes it a user-centric question, dependent more on the limits of human perception than computation. Riley disagreed, arguing that the first down line is not real-time in that it is delayed from the real time events: the television broadcast is slightly delayed so the augmentation can be placed atop it. He argued that because it could be viewed days later without any change in the experience, it cannot be real-time. This interpretation argues that real-time should be a source-based criterion, and argues that independent of perception, the computing has to be linked to real world events as they are happening. What these examples do is highlight not only the ways actors are contesting these technologies, but they also reveal disputes within the criteria of the definition itself.

While on one level these definitional questions are theoretical, the fact that Sportvision readily called their product AR means that they could potentially redefine AR by their inclusion. The definitional question could result in concrete actions if the organizations and members conclude that they do not believe a particular system is AR (e.g. by attempting to certify products, publicly questioning the link to AR, or simply persuading more people to support and enforce a particular definition). To date the debate over 1st and Ten has remained within the community, but the many brief histories and books mention the system in passing and the public association with AR remains a tension, raising the stakes from a theoretical discussion to one with real-world consequences. Some people are happy to adopt a broader definition to widen the CoP and have AR associated with what is a popular, widely known, and a well-executed application. For them, not meeting each prong of Azuma’s criteria is insufficient to exclude it outright; the system is a close approximation and some of those criteria might be met with future

developments. Other members prefer more rigid boundaries that exclude the 1st and Ten system. The failure to meet any one prong is a reason to reject it, because Azuma clearly states that the systems must possess all three characteristics. The popularity of the application is irrelevant, rather it is the integrity and defensibility of Azuma's criteria that is paramount.

The discussion above is just one example of the dispute, which springs from a much larger ongoing discussion over what AR is and who falls under the umbrella of AR. It also represents a form of boundary work, much more about retaining their scientific authority to declare things AR, rather than acceding that right to the media or to individual developers. One professor's appeal to ownership of the definition itself (e.g. "Ron and I") and to rationality (implying that one would have to be insane to disagree) are attempts to monopolize professional authority and resources to exclude others. Gieryn (1983) explains that these strategies are particularly effective for "ideological demarcation of disciplines, specialties, or theoretical orientations within science" (p. 792).

The reason this debate gets raised at every AR conference is because the argument itself is crucial to members' learning, reinforcement, and negotiation of the identity of the organization. Only through having these arguments can people "know what they think," which informs other decisions about the institutional positioning of the organization (Weick, 1995). These arguments take place in formal Q&A sessions (e.g. "where is the AR in your work?"), physical meetings (steering committees), online listserv discussions, and informally in conference hallways and at social events. Professor Azuma is on the ISMAR steering committee and present at many of these meetings and conferences, which further enforces the need to acknowledge at least an understanding of the definition. When I approached him about his definition, Azuma responded: "Oh God, do I really I have to go on record about that?" His

response was suggestive of someone who has both had too many of these definitional questions and been consulted to settle too many disputes, but also reflected a larger concern that any further interpretation he made would then be used as a declarative authority as people continue to engage in the messy, ongoing, definitional debate.

All of these arguments become the practical structures that constitute an organization: “Through multiple layers of ordinary talk, people in organizations actually discover, as a deeply collaborative and contingent matter, their shared goals, many agendas, environmental uncertainties, potential coalitions, and areas of conflict (Boden, 1994; p. 8).” It also illustrates the dynamic nature of constitutive communication, where ISMAR is not a fully formed organization with discursive entities but instead an organization that is continually forming over time through discourse and a process of ongoing communication, sensemaking, and negotiation. It is through these debates that epistemic authority is enacted, such that actors decide where to locate the legitimate jurisdiction over definitional facts, in people and in structures like steering committees (Gieryn, 1999).

Within ISMAR, Azuma’s definition is taken so seriously that actors need to negotiate with it as a membership requirement to enter the community. One example comes from Brian, a computer science professor who presented at both ISMAR2012 and 2013, and how he came up against this definition in the peer review process:

If you look at Azuma’s classic definition as visual, 3D, real-time. One of the comments I got back was ‘is this really AR because when you look at the [application] interface it’s a 2-dimensional representation.’ I argued that it was because the app is tracking the 3-dimensional tags and the alignment has to be specific to those tags otherwise the app doesn’t work. There was some debate about ‘Is it AR?’, so I had to make the case in my rebuttal (interview).

These interpretations and disagreements are a form of membership negotiation, which eventually becomes hardened into ways of self-structuring as the leaders explicitly build this requirement

into the screening process (McPhee & Zaug, 2000). ‘Suitability’ for ISMAR is a criteria for reviewers, and becomes a place where anonymous arbiters can raise definitional objections to incoming work. One major problem, however, as we can see through the 1st and Ten® example, is that these interpretations are not static. Individual members have differing perspectives, which means different interpretations are enforced. Greta, a former co-chair of ISMAR who looked at many reviews, explains:

I found that the people within the organization had different ideas that weren’t always articulated but became very clear in reviews. We’d see often times where reviewers [...] would say ‘this isn’t AR’ and give a reason that [to me] wouldn’t be a reason to disqualify it, so there’s a lot of unarticulated things or ideas of ‘community’ that people give to this.

Her comment reveals latent divergences in how actors define the scope of problems, the aspirational goals for AR, the extent to which certain systems are valued, and the conditions that have to be met to be within the boundary of AR. This inclusion and rejection based on definitional criteria is an instance where ‘boundary work’ translates into practical action as an important means of social control and sanction within the academic community (Gieryn, 1999). The leaders of the organization solidify the definition hierarchically by forming the steering committee, an example of how organizational structure flows from the communicative practices (McPhee & Zaug, 2000). They then formally grant authority to less experienced members to enforce the definition through the review process and suitability criteria, but maintain programming authority. Peripheral members who want to become full members then have an incentive to model their practices after the expectations of the steering committee, or at least what they presume them to be, which one member of the steering committee admitted was highly dependent on Azuma’s definition. These multiple layers of definitional checks are designed to preserve a particular type of AR, which in turn structures the participation of newcomers to learn

this value (Lave & Wenger, 1991). While this demonstrates how it internally structures one organization, we begin to see some members taking more tangible actions to enforce the definition with other applications, devices, and members across the community.

Smartphones, Tablets, and broadening the definition of AR

Ongoing technological developments and devices spark new questions, forcing the community to revisit the definitional question, given these new entries into the space. One such instance occurred when several companies announced that they were offering ‘AR’ browser applications on smartphone devices. As mentioned in the previous chapter, this move sparked a split in the community around what kind of device should deliver AR. It also had the effect of bringing a more diverse set of actors into the AR community, and a sizable portion of whom wanted a broader definition than a strict application of Azuma. Even some who had previously relied on Azuma’s definition were now questioning whether it should include or exclude mobile applications and whether it was too ambiguous in its criteria. Amy argues that these real world developments pushed for a broader definition:

One thing that has happened, and is one of the reasons why someone like myself could be identified as working in AR, is through consumer devices, through iPhone, camera, GPS, devices that what would house AR experiences. That also opens up exploring in a way. [...] When technologies become more widely used, when that happens the definition of AR always changes (interview).

Will, a member of the ISMAR steering committee, explains that his perspective has been changing: “Historically, I was much more a strict interpretive person. I’ve loosened up a little bit about it. I still think that there needs to be overlay on the physical world. [...] But how tightly coupled and how tightly registered and how real time it is, I’ve certainly loosened up on that (interview).”

Oliver, a professor of computer science, laments that even though applications like mobile AR browsers represent a significant breakthrough in technology, they are too often excluded or disparaged as ‘pseudo-AR’ by those with a rigid definition of AR. Rather than getting bogged down in the criteria of certain degrees (e.g. what is real-time enough, interactive enough), he argues that the ability to overlay text and graphics on top of a physical location specific to the user should be sufficient. Some feel like there should not be a strict requirement to meet each prong of Azuma’s definition, and instead suggest a reasonableness standard. Oliver offers a broader definition of AR as ‘any context-aware, or location-based computing experience that overlays content on the physical environment.’ Industry members of the AR CoP favor a broader definition, as many companies make the association – Layar, besides including ‘AR’ in its company name, defines AR as technology that “allows for a digitally enhanced view of the real world, connecting you with more meaningful content in your everyday life. [...] AR adds layers of digital information – videos, photos, sounds – directly on top of items in the world around us.” Wikitude and Junaio, other major mobile browsers, similarly market their products as AR, as do other mobile applications.

However, some pushed back against this expansion of the domain of AR and potentially bringing new participants in the space. Some worried that the meaning of AR was being ‘co-opted’ by these newcomers to the community, and again engaged in boundary work by defending a more narrow application of Azuma’s definition to try to exclude these applications.

Here’s Dennis, a graduate student working in a laboratory researching AR:

A lot of things are being called augmented reality even though they’re not. [...] Anything that’s on the phones right now, because it has to feel like it is part of the world, in my opinion. Just having things hovering around, based on location but not the actual real world environment, that’s not AR, because it’s not interactive. It’s just a nice little thing that you’re putting there (interview).

Another way actors definitionally exclude mobile applications is by adding an implicit criteria of immersiveness, which was not a prong in Azuma's definition but is very much a part of the strict vision for AR. Here's Sven, an AR artist:

I think at this moment I can cope with waving around my iPhone for a couple years, but eventually I'm looking forward to this more immersive experience. [...] It comes back to this definition of AR. I think AR really happens when you make this mix of something in this physical space and something in the virtual space. So if you combine data and the material and the experience as a whole then something new happens (interview).

The arguments applying Azuma's criteria to mobile browsers reprise similar arguments about 1st and Ten. People argue that many mobile AR browsers do not 'combine' real and the virtual because it is location based, and simply adds a floating object that does not recognize its surroundings. On the question of interactivity, Dennis excludes mobile on the basis that AR must be interactive on a source-based level (see Figure 22). Oliver counters that this is an arbitrary distinction, arguing that mobile browsers are interactive as long as they meet the user-centric interpretation (see Figure 21). Sven wants to add an immersiveness criteria to the definition of 'when AR *really* happens,' which functionally excludes smartphones.

Companies that produce AR experiences on mobile devices have attempted to enter the community, by sponsoring events, presenting at industry conferences, participating in standards meetings, and hosting exhibits at the academic conferences. After one of their sessions, however, a few times I have overheard grumbling about how that session about mobile AR was not really AR. Certain members of the community started to make a hierarchical distinction between "hard/strong AR" and "soft/weak AR." Here's a keynote speaker at ARE2011 making that exact distinction as a call to action:

There's been a lot of 'so-called augmented reality' on mobile devices over the past couple years, but most of it really sucks. Most of it is what I would call 'weak AR' based on the compass, GPS, and vague sense of how stuff out there in

the world might relate to your device based on those rather crude sensors. ‘Strong AR’ is when you, when some little gremlin is actually looking through the viewfinder at what you’re seeing, and it’s saying ah yeah that’s, this is that, that’s that and that’s the other and everything is stable and visual, that’s ‘strong AR’. Of course the technical requirements are so much greater (ARE2011)”

Here, “hard/strong AR” refers to an idealized state of AR in a heads up form, with visual results that indistinguishably interact with elements of the real world, and that require algorithm-intensive visual computing. “Soft/weak AR” refers to anything handheld, where the mixing of the real and virtual is imperfect, and uses only geolocational data (Abrash, 2012). This hierarchy grudgingly accepts mobile applications as AR, but emphasizes that there is a much bigger goal and that these early applications are merely temporary analogues. Beyond the explicit expression of value in ‘hard AR’ and disdain for ‘soft AR’, this represents how actors can integrate new technologies into AR while still upholding a higher ideal that is aspirational and difficult. In the AR space, the hard/soft dichotomy is one way actors in AR have integrated mobile AR into the CoP, which simultaneously allows them to talk about the possibility of AR, while still pointing to a different vision and need for ongoing research. This work of creating hierarchies and applying scientific definitions reflects the limited nature of epistemic authority, where it “exists only to the extent that it is claimed by some (typically in the name of science), but denied to others (which is exactly what boundary-work does) (Gieryn, 1999; p. 14).” The longstanding members of the CoP, faced with a set of newcomers, began deploying discursive practices that reproduce certain power structures at the exclusion of alternate configurations. Maintaining this hierarchy allows them to still secure investments toward a particular future of AR, while also positioning themselves as the authority and experts of that future.

Representatives from mobile AR companies vehemently objected to this attempt to marginalize them, arguing that this two-tiered categorization has done the community a

disservice by framing the issue in a dismissive, divisive way. One representative from a mobile AR browser company posted a strong rebuttal on their blog. To them, this ‘hard’ versus ‘soft’ dichotomy represents a need to preserve an idealized form of AR as ‘true AR,’ and serves as a thinly veiled attempt to segment out mobile browsers as ‘so-called AR’, a discursive tactic to delegitimize their credibility similar to calling something ‘pseudo-science’ (Gieryn, 1983).

Outside of the influence of the academic contingent, companies are choosing to call their applications AR anyway and attempting to create a de-facto definition of AR that is broader than Azuma. However, the question of self-identification begins to reveal tensions over what Wenger (2000) refers to as the negotiability of boundaries, and whose perspectives are accommodated in the definition. While everyone ostensibly supports AR, these debates represents sensemaking work that make obvious the multiple stakeholder groups involved and reveal their differing agendas for organizing (Weick, 1995).

This persistent debate has a noticeable effect on longtime members of ISMAR as well as newcomers to the field. Here is Greta again, explaining how she came to the community late but is now very conscious of the issue, having learned to conform to the definitions:

If I’m at ISMAR and we’re talking or listening to computer scientists or engineers I’ve learned they might have a stricter definition, like Azuma, or that they have particular papers and researchers in mind who have strongly identified AR for their community (interview).

Learning the vocabulary within ISMAR takes place through direct observation of research presentations within the conference and observing formal/informal debates on the subject, and becomes a way to demonstrate belonging, competence, and expertise (Lave & Wenger, 1991). Will mentioned that his perspective has changed over time, and he has loosened up on the definition. This demonstrates the fluidity of the debates but also how, over time, these shared practices connect communities and offer learning opportunities in their own right, where would-

be practitioners not only learn ‘about’ the community but also ‘how’ to be a member (Duguid, 2005).

Defining AR in Arts, Media, and Humanities

ISMAR originated as a conference focused on the science and technology (S+T) development of AR. That was broadened in ISMAR 2009, when a separate track for arts, media, and humanities (AMH) was added to the conference call. One of the co-founders, Stan, explains his motivation:

What I'll do is just a little other thing called AMH, which would wrap around the S+T track with not only arts and humanities but also industry startups, and workshops with applications for the military, medicine, and entertainment coming on board (interview).

The call for papers for ISMAR now included topics like art installations and performances, interdisciplinary and transdisciplinary studies, media studies, and critical design. Several other members of the technical steering committee embraced the move, and eventually decided to keep AMH as a continuing part of the ISMAR program for future conferences after 2009. At ISMAR 2013, one of the members of the steering committee explained: “The reason we added AMH to ISMAR was because we found there was something lacking. There was a lack of real critical social discourse in this area of technological innovation – cultural ramifications are overlooked quite often particularly with new technologies, especially given a rush to commercialize.” Here’s Will, a member of the steering committee and former co-chair for AMH: “There were several reasons we started thinking about AMH. One of them was just that people were already doing things that were interesting to the community [...]. I think it's also been an effort to try and expand the community.” ISMAR as an organization decided to change their institutional structure to reflect a broadened organizational focus.

Starting AMH was one thing, but the implementation and integration of AMH remains unresolved and has proven to be a major challenge for ISMAR and its members. “AMH has a complex relationship with ISMAR in general. Part of it is that researching augmented reality, in computer science, you can do research and never have a human being involved. It just doesn't matter (Will, interview).” Within the organization, definitional questions have been raised regarding the AMH track: namely, how rigidly to enforce the definition of AR when evaluating AMH submissions.

For example, several artists have taken up AR as a prompt/platform to ask interesting questions, with little regard for the sanctity of the academic definitions: “The goal of [our art] is to challenge the audience to think about how we perceive these works and [...] to work with other disciplines as a collaboration of ideas (Julian, ISMAR 2013).” At ISMAR 2013, the AMH program featured an art exhibition held by the Mixed and Augmented Reality Arts Research Organization (MARart). Comprised of artists from all over the world, one of the stated goals of MARart is to explicitly challenge those definitions:

The Mixed and Augmented Reality Arts Research Organisation (MARart.org) seeks to develop new dialogues in regards to high-end research methodologies, cultural inquiry and representation in the increasingly immersive and pervasive field of mixed and augmented reality art. [...] MARart is pervasive mixed reality by default. It is physical and virtual hypersurfaced content. And more recently it is networked, and this is the catalyst for a new, broader definition of mixed and augmented reality art to be attempted (MARart.org).

The MARart exhibition at ISMAR 2013 featured many diverse works. One of them was Damian Hills's *assimilate: enactive topologies* project, using a projected floor space to explore how people engage in collaborative narratives. Another was *#PRISOM* by Mez Breeze and Andy Campbell, a head mounted display with a built 3D environment where users navigate and engage with objects, while simultaneously being confronted by issues of surveillance and privacy in this

virtual environment. One last example is Jorge Ramirez's *Cymatics Disport*, where he makes sound visible through a physical object.

The MARart exhibit was well attended by many ISMAR participants in attendance. The chairs of ISMAR2013 were instrumental in organizing the event. The event, however, and some of the art examples listed, revealed some of the key differences in perspective and motivation. For one, none of the exhibits were offering demonstrations of any new or novel technologies. Some did not even use a display of any kind: Ramirez's piece used sound waves to change the physical properties of an object (placing corn starch over a speaker to see what form it would take), but did not 'combine the real and the virtual' through a display and was not interactive for the user (although it was interactive with the sound waves). Hills's piece was a screen on a floor which was real-time and interactive in that it responded when people stepped on it, but was not 3-D. Breeze and Campbell utilized a heads up display to create a 3-D, real-time, and interactive world, but it was a fully immersive virtual reality environment that arguably did not 'combine the real and the virtual. For those rigidly applying Azuma's definition, many of these works would have been excluded. The artists countered that the technology is less important than the question of what the technology allows for and what to do with the technology, as well as what the technology represents and what it could potentially mean. To think about these pieces in terms of definitional criteria would be to miss the point and the message of these works, though technically could have justified excluding them from ISMAR. Greta explains that this is a reaction to people starting to push boundaries:

One problem [...] that AR has is that if you want to push it, and that often happens when you're talking outside the more narrow definition of AR, [...] there's a sense of 'Where are the limits? What's not AR?' [...] I think the response among some people in the community has been to go out of retreat to a more technical definition of AR, in order to safeguard from this complete unlimiting (interview).

Enforcing the definitional boundaries, then, might exclude certain ideas at a prima facie level, even if they are valuable for the community to think about. For instance, because this question of VR is a barrier point of entry for ISMAR, it might seem antithetical to allow the #*PRISOM* virtual reality game into the space. These particular pieces came into ISMAR as part of the art exhibit and were not reviewed as papers, but they illustrate the kinds of difficult nuances and distinctions that AMH raises.

Many members of the steering committee are quite emphatic about the need for the AMH program, as repeatedly they have said that they “need the arts, the social sciences, to really drive this industry forward (Liu, interview).” The idea is that these applications will motivate a broader and more critical approach to thinking about these technologies. The steering committee, by and large, has shown a tremendous willingness to enact the AMH program, even placing the conference on a single track program so all participants could see the sessions across thematic tracks. How rigidly to enforce Azuma’s definition for AMH continues to be negotiated within ISMAR. If the goal is to bring perspectives from other disciplines into the AR space, an easing of the definitional requirement may be necessary, particularly because people from other disciplines may not have the same rigorous adherence to Azuma’s definition.

For other members, however, enforcing a narrow definition is not a matter of choice, but central to the mission of ISMAR. With other segments of the community attempting to broaden the definition of AR away from Azuma’s criteria, there’s a sense of righteousness for those who defend it – that if nothing else they must uphold the definition within their own domain. For others, there are more self-serving reasons to enforce a narrow definition for AMH, as it opens up another venue for publication. Some members of the steering community have described

efforts to simply pull the AMH track toward S+T by submitting S+T papers and posters to the AMH track. Here's Stan again:

What they've been using Arts and Humanities for is Science and Technology lite, so for all the doctoral students who need a paper, but who can't get into Science and Technology. And they can't get in because they have this whole program [...] were they want to keep x percent accepted (interview).

Greta, a co-chair for AMH, also noted this occurrence: "We would have papers submitted to the AMH track that clearly should be in the S+T track for various reasons, [...] by design to see if they could get in easier or something (interview)." This is one way of changing the definitional requirements of AMH, which would have the effect of pulling AMH closer to S+T. These definitional questions can be a way to exclude contributions from other fields or impose another barrier where only definitional insiders (e.g. engineers) are able to enter.

These are not just boundary skirmishes within a discipline, but over the related disciplines and the extent to which other areas have a role in pushing AR forward. While the steering committee has concluded that other areas should have a voice in their venue, other actors are deeply ambivalent about the work from other fields, and wonder about the value of AMH being represented at ISMAR at all. For many computer science professors and students, they submit to ISMAR and travel to the conference to see the latest developments in technological research and to see what is now possible with new technology. Eric, a professor at ISMAR 2012, explained it like this: "I come to ISMAR to find out if there's a better tracking algorithm, better camera and computer vision algorithms, better ways of displaying augmented content (interview)." For these actors, the definitional enforcement may be one way of remedying that, and thus a reason to eliminate AMH.

The larger conversation over how much to ease the definitional rigor when broadening the domain speaks to an internal struggle within organizations, as these boundaries are negotiated

by the participants and embody their vision for certain institutional structures. The leaders of ISMAR want to invite outside perspectives, but with the AMH example we can see how these definitional debates about ‘what is AR’ end up being politics by other means. One person may have concerns about what the definition means for the organization, while another simply has a strong belief in the definition and is adamant that it should be enforced in their domain of influence. Some may be motivated by self-interest to secure more publications, or position their work as more central to the AR question, or they may be trying to eliminate the AMH program because they do not see value in it. It may also be some combination of those motivations, but the public action of enforcing a strict definition for AMH is the same.

Are these Academic Questions Counterproductive?

Calling something an ‘academic question’ in the pejorative sense suggests that the question is of little practical value. Actors in the AR community sometimes use the phrase ‘academic question’ to describe both who is asking the question ‘what is AR,’ and the extent to which a segment of the community thinks those questions are important. While the definitional question is one that occupies a great deal of time at academic and standards conferences (with whole sessions and task forces devoted to the question), it is less of a priority on the industry side. A marketing manager from a mobile AR company explains:

We need to stop debating about the jargon and definition of AR and instead focus on the experience and how it will help people get through their day. We need to take it back to [...] how it will allow people to digitize things they couldn’t engage with before and allow people to [...] discover content (interview).

Todd, another marketing manager, points out that “users do not care whether something is AR in a definitional way, so we should stop worrying about it (interview).” For those who come from an industry orientation, it is not a productive exercise to apply a priori criteria about the source of

the augmentation and the conditions under which it is produced. The perspectives above worry that doing so might also have negative consequences, while being narrowly focused on an ideal state of AR can exclude valuable perspectives, ignore real world applications, and confuse people about the term AR. Another more subtle motivation might be that their products may not fit under certain definitions of AR, so industry actors are not keen on being too exacting for other applications. The underlying perspective that is most prevalent at these industry conferences is that the simplest approach would be to take a colloquial approach to defining AR, not as a term of art but by simply combining the definition of ‘augment’ and ‘reality.’ This ensures that their products can be classified as AR, while also giving them significant leeway to define the term. This is an instance where corporate and marketing interests attempt to expand their epistemic authority into the scientific domain, to redefine and diminish the importance of academic definitions (Gieryn, 1999).

Unsurprisingly, the academic and standards communities have reacted strongly to this move to loosen the definition of AR, as many of their practices and the way they teach their members are built upon the definition. First, members of the academic community bristle at this definition as amorphous and without a clear line of demarcation. Brian, a professor of computer science, is emphatic that this definition confuses certain enabling technologies with AR:

A map that knows your current position is not augmented reality, it’s not mediating your vision. Those virtual objects don’t line up with the world in any meaningful way. A hearing aid is not augmented reality, all its doing is amplifying it’s not producing some virtual sound that interacts with reality in a meaningful way. [...] A system where you have a camera, you press a button, and it doesn’t take a picture until your pet looks at you. That’s automation, but it’s not augmented reality (interview).

Within Brian’s answer, we also see alternative language that he uses to categorize those other devices (e.g. amplification, automation), which further preserves AR as a unique category of

research. Riley argues that multiple groups are responsible for this shift, and makes a slippery slope argument about the expansive potential of combining the meanings of ‘augment’ and ‘reality’:

The term ‘augmented reality’ is being pushed away from its academic meaning by a few different groups. [1] Those who are excited about the idea, and think whatever they happen to be doing is ‘augmenting’ ‘reality’ and thus they call it augmented reality. My issue with the ‘phrase is defined by composing the definitions of the two words’ is that it's meaningless. Coffee augments my reality (‘I'M MORE AWAKE NOW!!’). LSD clearly augments my reality. Pollen augments my reality. Where do we stop? [2] Those who argue for a slightly broader definition, up to and including any ‘context aware, or location based, computer experience’ [...] [3] Those who have a business and want to benefit from the hype by calling whatever they do AR so they can generate some buzz. In all cases, however, I think it misses the point. All of these examples simply focus on the word ‘augment’ and then say ‘XXX is augmenting reality’. Semantically, it's hard to argue with. But, practically, it avoids the real issue: what are these different experiences, what do they offer the user (or developer) in terms of capabilities, and are there meaningful differences between them. [...] If everything is AR, then nothing is. [...] So I will keep calling those kinds of experiences (and not other things) AR, until a better collection of definitions is put forth, since my community has been calling it AR for > 20 years.

In this post, we see multiple levels of boundary-work being enacted. There is a claim that the academic meaning has the epistemic authority, and that the academic community has had it for over 20 years, but that the meaning is being ‘pushed away.’ Second, specific groups are held responsible for this, namely lay-people who are confused by or uninformed of the academic meaning, which creates an unlimited definition. Also responsible are people who understand the academic definition but selfishly want to broaden AR to include mobile location devices. These do not strictly meet the Azuma definition, and the arbitrariness of their proposed counter definitions is fed by public confusion over the term itself: “Because at least [in academia] I can say my definition is this, and here is the argument that this meets that definition. If you don't have a definition, then it comes down to does it have the flavor, no one can explain why one

thing is accepted over the other (Brian, interview).” Both sides worry about the public ‘confusion’ over AR, but they disagree about the cause of this confusion. Industry claims that it is the debate over AR itself that is confusing, which recommends a loose definition, whereas the academic community claims that it is the myriad of examples that get called AR that is the source of the confusion, which recommends a rigorous definition.

A more aggressive type of boundary work is being enacted in response to the third group, the commercial actors who simply call things AR for marketing purposes to generate hype. Because these companies are proactively trying to claim the authority to define AR through the release of their products, people in the academic community are especially critical of the implications of their definition, their motivations, and their place in the community. First, they charge that those definitions are limitless, but they also criticize the specific uses that are being called ‘AR.’ The implications of such a perspective is that it could harm the field of AR more generally by giving users negative impressions of AR both in terms of content and execution. Here’s Will, a COO of a major AR company:

They’re technically not even an augmented reality app – they use 2D augmented material that does not understand or obfuscate objects, doesn’t understand the scene. So that’s one reason why augmented reality is getting such a bad name (interview).

This statement attempts to link certain definitions to fears that those uses and mis-associations are discrediting the larger work of the community.

Actors also begin to cast aspersions on the commercial motivations behind this broader definition. These are concerns over what Wenger (2000) would call negotiated imagination, or the extent to which two different groups see themselves as being members of the same community with common interests and needs. Within the academic community, some worry about the extent to which certain businesses ‘care’ about AR as a technology, and express

concerns that they are simply using it as just another marketing tool that presents a poor user experience. Concern about the profit motive of industry actors is another way that scientists attempted to distinguish their boundary from practitioners and sellers of the technology (e.g. industrial craftsmen). Gieryn (1983) argues that one demarcation is the argument that “Scientists seek discovery of facts as ends to themselves; mechanics seek inventions for personal profit” (p. 786). In this instance, casting aspersions about the profit motives of other actors discredits them as responsible stewards of the technology, while simultaneously affirming the academic authority to make the determination. This distrust of marketing motivations is strong amongst the AR community, and some are strategically attempting to exclude these industries through a narrow definition.

These newcomers play an important role in the CoP, and in many ways the newcomers and longstanding members are dependent on one another – the newcomers to learn and eventually become the leaders that carry on the CoP (Lave & Wenger, 1991). Some of these newcomers, however, in their rejection of existing definitions, are attempting to displace those practices and replace the process of learning. In the process they are attempting to expand their authority, however, and continue calling their products AR. For them the more important question is who can successfully change media representations and consumer understandings to align with their broader definition, with the goal of benefiting their companies and field financially. This reflects the zero-sum nature of epistemic authority, that the “legitimate right to have one’s reality claims accepted as valid or marginally useful is no plum at all if everyone enjoys it all the time (Gieryn, 1999; p14).” This comes through most explicitly in Riley’s claim that ‘If everything is AR, nothing is,’ where accepting a definition without boundaries threatens the whole enterprise.

Is Google's Project Glass AR?

Google's Project Glass is just the latest artifact to revive the definitional question, and has been the subject of much speculation as a high profile instance of a large company investing in and releasing a commercial HWD. Unlike 1st and Ten and the mobile applications that actively use the term AR, Google has been conspicuously absent from the AR community. The first demonstration video of Glass titled 'One Day' was released a month before ARE 2012, one of the largest annual AR industry conferences in the world. Google did not show up, even though the event took place in Santa Clara, CA, only a short drive away from Google Headquarters. This was not a one-time omission, as Google has not attended in any official capacity any of the subsequent AR industry conferences I attended (AWE2013, InsideAR2013, AR Summit, AWENY 2014). Bill, a computer science professor who used to work at Google explains that this disavowal is a conscious move on their part: "When I was working at Google we were told not to refer to Glass as augmented reality or an augmented reality device." The sense for Bill was that Google did not want to get into

Despite the lack of explicit association, the AR community quickly engaged in definitional boundary work around Project Glass. Almost immediately after Google's announcement, people began to argue that the examples shown through Glass were not AR. Here's an excerpt of a blog post from Riley, addressing the media, applying some of Azuma's criteria to Glass:

To all the press: this [Glass] is a heads-up display, it's not "augmented reality." AR is about putting content out in the world, virtually attaching it to the objects, people and places around you. You could not do AR with a display like this (the small field of view, and placement off the side, would result in an experience where the content is rarely on the display and hard to discover and interact with),

but it's a fine size and structure for a small HUD. The video application concepts are all screen-fixed ("heads up" instead of "in the world") for this reason.

Here we see a reprisal of the user and source interactivity question surrounding 1st and Ten, as well as the implicit 'immersiveness' criteria that was raised about smartphone and tablet AR. Since the Project Glass announcement, many members of the community have adopted this boundary perspective – at all the conferences I attended, it almost became a rote expression where any reference to Project Glass was followed by the caveat “and I know it's not AR, but...” This type of move represents not just definitional boundary work, but also as a reinforcement of a community held norm and an active signifier with which to demonstrate belonging in the CoP (Lave & Wenger, 1991). It also represents an instance where there were enough people within the community to adopt this position, such that it had to be acknowledged

While Glass might not purport to be AR and some members want to keep it that way, the relationship is more complicated than just mutual disavowal. Members of the AR community still talk about Project Glass a great deal, and in a way that reflects their hopes and fears. Brian explains that this definitional ambivalence is a product of disappointment in Glass and concern over it becoming an exemplar: “I've been frustrated by Google Glass. Is it the AR headset I've been dreaming of? No. I want a full field of view with high resolution. Not making AR apps using it where they'll just be limited to the tiny bit of screen space and you have to hold your head at this weird angle.” Frank, a developer for an AR company, also explained to me that there's an element of ego and hubris in the community: “The fact that Google didn't come says something, and people in the community are competitive, maybe a little bit arrogant. They think they can make something better than Google (interview).” Many members of the community worry that Glass will fail as a commercial product and be held up for ridicule, which will then be used as a warrant against other AR concepts. Some draw analogues to the cautionary tale about

the rush to commercialize VR. That story serves as an important part of the shared repertoire of communal resources (Wenger, 2000), as a consensual historical narrative of the technology that is shared and circulated amongst its members to encourage and justify skepticism about forces outside the community.

Without embracing Glass as AR, however, other members have begun expressing optimism about what it represents. Some AR companies like Layar have been developing their applications to work on Glass. AR conferences have used Glass strategically as an attraction, at AWENY 2014 there was a developer session specifically for Glass and a Google Glass Community leader as an invited speaker (although not an official Google representative). That speaker, Greg, came into the event very conscious of the community's attitudes toward Glass, even acknowledging them explicitly: "I know Glass is somewhat of a bad word in the AR community. [...] Glass isn't AR, but it is not trying to be." This clarification was another gesture to the community's definition, even while the conference was capitalizing on the tangible excitement about Google's presence.

One hope is that Google will help demonstrate that there is significant commercial interest in wearable devices. Teresa, an academic researcher, expressed optimism that Glass would make it socially acceptable: "A hat tip to Google for putting it onto popular culture. [...] The important thing that Google Glass does is it doesn't just talk about the technology, it focuses on the lifestyle and what you can do with it in everyday life." Bran, a lawyer working in the digital media space, was also hopeful that Glass would also pave the way for resolving many legal battles he anticipates for HWD: "I think the biggest thing is the scope of privacy which is a huge issue and because they're so consumer focused, they'll be well equipped to deal with that. Google has the money and resources to fight these battles over negligence, content, copyright,

and privacy and clear the way for the industry to grow.” These moves by members of the AR community to exclude Glass through definitions while embracing it functionally represents calculated boundary work. It leaves open the option of appropriating Glass should it succeed, while rigidly enacting barriers if it fails.

Classifying the AR Space

These definitional debates are not just standalone conversations that stakeholders use to structure their organizations and enforce the boundaries of the community, they also get embedded into formal systems of classification. Just as there are multiple interpretations of the definition of AR, there are many taxonomies circulating that attempt to map the AR space. These classification systems are not only organizational tools, they also exclude applications and privilege certain types of work based on how they are sorted (Bowker & Star, 1999). The desire to improve/reconcile these classification schemas is an ongoing endeavor, which is a part of how maps of science are endlessly remade (Gieryn, 1999). At ISMAR2012, I attended a preconference devoted to ‘Classifying the AR Space,’ an all-day session outlining the current gaps in classification and how to create a better system. The goal was specifically to address the problem of parsing the various applications, devices, components of the AR field, but also to either bypass or settle the definitional question. Another example is the glossary taskforce at the IARSC Meeting, which was organized to reconcile different definitions of AR terms. These efforts to standardize language have been noted as an important part of the creation of professional communities, where there is considerable frustration and contestation over the expanding variety and complexity of language (Marvin, 1988). The ongoing development and debate over the classification schemes is part of the invention and development of a technical

language, and represent sites of contestation that help us understand the priorities of the AR community.

Technique Centered Classification

Milgram and Kishino's (1994) definition was described above for its move away from VR, but within that they were also attempting to create a taxonomy for different types of visual displays. The way that they propose drawing distinctions between displays is through three criteria: "Extent of World Knowledge (EWK), Reproduction Fidelity (RF) and Extent of Presence Metaphor (EPM)." EWK refers to how much the system knows about the real environment, RF refers to quality of the virtual environment (on a spectrum from simple edges to complex 3D objects), and EPM refers to how much the user feels present (Milgram & Kishino, 1994). That particular taxonomy was focused on devices in 1994, which were primarily virtual reality headset techniques (Barba, Mynatt, & MacIntyre, 2012). Attempts to apply this taxonomy to explain current system would not account for mobile techniques that are now common to the AR space (Normand, Servières, & Moreau, 2012). The taxonomy also makes assumptions about the effects of technique on presence that may not be true (e.g. presence would be classified as low when the display is monoscopic and high with stereoscopic see-through 3D graphics). Scholars have argued that presence is multidimensional, and in fact there are multiple operational definitions of presence (social richness, realism, transportation, immersion, etc.), with varying levels of cause and effect for how people perceive presence (Lombard and Ditton, 1997). This technique centered classification presumes a relationship between augmentation and presence, while also operationalizing presence only as immersion. This would classify devices based on something that would require empirical investigation (e.g. how much do people perceive

presence using this device). Supporting this taxonomy, however, would be one way of excluding mobile devices from the AR space and making a particular form of immersion part of the criteria.

User-Centered Classification

Rather than classifying based on techniques, Hugues and colleagues (2011) propose a user-centric taxonomy. This eliminates the system requirements for sorting things, but instead categorizes systems based on how a user perceives the augmentations. First they propose a binary classification of whether the user has an augmented perception of something real, or if they are using AR to create extend reality in ways that are not physically possible. They propose multiple levels of categorization: augmented documentation, reality with augmented perception or understanding, perceptual association of the real and virtual, behavioral association of the real and virtual, substitution of the real by the virtual or vice versa. One advantage of a user based classification is that it is easy to grasp (although would require expertise to sort into subcategories), but the question of what other criteria to sort by is still up in the air. One presenter at this preconference proposed distance as another user criteria, which is how far from the user the system attempts to understand the scene around them. Based on a continuum, the four categories they propose are based on the region surrounding the user: localized to the head area, body (arm's length), distance (out of arm's length), and remote (out of sight) (Hoang, Smith, & Thomas, 2012). Another presenter wondered if the perceptual categories should be more closely tied to psychological perceptions of how people actually see the world (Lu & Duh, 2012). These user-based classifications only deal with perception, and do not sort using interactivity as a standard.

While easy to understand, this type of taxonomy raises its own set of issues. One is that it excludes the discussion of what is real: “The question of what is real, existence and properties, is not examined here (Hugues et al., 2011; p. 2).” Bracketing this question a priori is an instance where the classification system deems it unimportant, or at least secondary to the issue of classifying what ‘counts’ in AR. Another limitation is that this taxonomy might not be sufficiently indiscriminate. Normand and colleagues (2012) observe that very different methods like mobile devices and projector devices would be lumped in the same category. Different devices might also traverse each of these perceptual categories depending on the application. The distance standard proposed at the preconference is also admittedly arbitrary, as “intermediate” was intended to capture the size of a room, but “long” could be anywhere from 30 feet to as far as the eye could see. Relying on psychology to define the user criteria would not necessarily improve the categories, rather it might just port the perspectives (and conflicts) of another discipline into the AR classification. It is also unclear how the taxonomy would recognize differences across users in how they perceive augmentations. Lastly, the categories only assume augmentation of one sensory modality, vision, which may not account for how augmentations might merge or bridge different senses.

Information/Presentation Centered Classifications

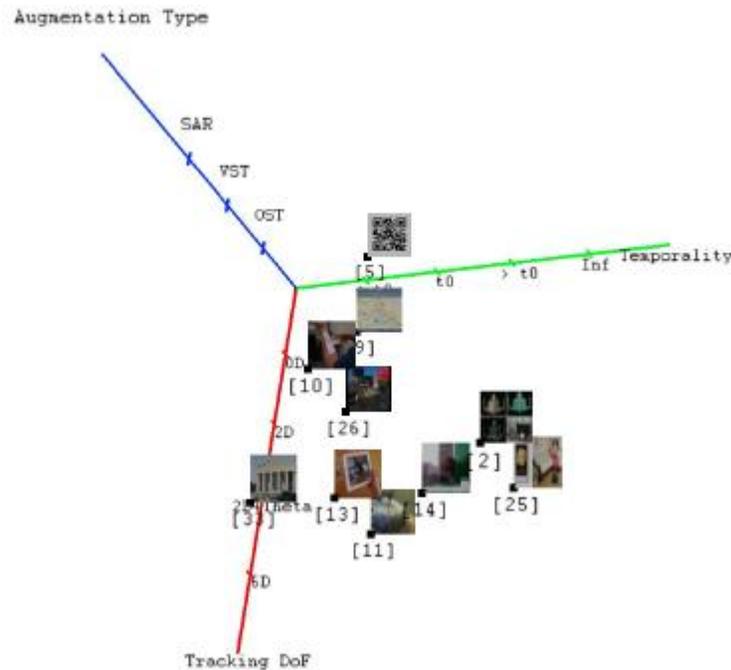
Suomela and Lehtikoinen (2004) proposed a taxonomy for dealing with location based information, with one factor being the environmental model used (0D to 3D) and the viewpoint used (1st person or 3rd person perspective). This classification system only deals with location, which means it only maps one particular type of mobile application. A more specific classification based on the AR presentation space is proposed by Tönnis and Plecher (2011),

which is based on six criteria: (1) temporality: whether the device continuously scans the environment or is discrete; (2) dimensionality: how the augmentations appear (2D, 2.5D, or 3D); (3) registration: how the real world is understood by the device; (4) referencing: the relationship between augmented objects in the scene; and (6) mounting: whether the device is hand-mounted, head-mounted, mounted to another real-world object, etc.. This particular classification is quite extensive, but only maps ISMAR papers onto the schema, which makes it less clear how it might accommodate real world examples, or some of the contested examples above. At the preconference, they explained it was primarily for steering academic research and defining the problem spaces that are understudied.

Degree of Freedom Classification

Normand, Servières, and Moreau (2012) propose a taxonomy that combines many of these perspectives. Their classification is based on four factors: degrees of freedom in tracking, the relationship between the user/device/sensor/real world, the type of tracking, and other rendering modalities. This is the first classification schema to consider the possibility of other sensory modalities besides visual, and is intended to be more adaptive to future technologies. This classification scheme also accommodates technologies like QR codes and Google Maps that are often not considered AR. However, the classification also places normative value judgments on particular technologies. When talking about QR codes, the authors caveat its inclusion and suggest it is less important to work on by saying: “it is questionable whether these kinds of applications can be considered AR,” On the other end of the spectrum, it notes that the 6D category is what has been “traditionally called augmented reality by computer vision scientists (Normand, Sevieres, and Moreau, 2012).” The classification does allow for the possibility that certain applications could fall along many of the coordinate lines, and it demonstrates how it

might classify existing real world technologies (See figure 23). It is quite a complicated criteria system, which may be helpful for parsing out fine distinctions, but could also make it difficult to sort and require significant expertise to make those distinctions.

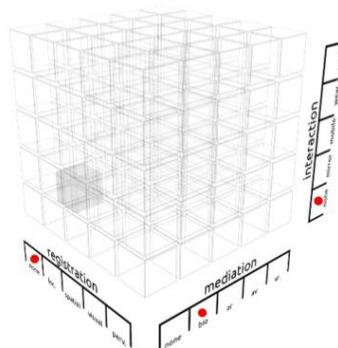


(Figure 23 – Classifications of AR applications, Normand, Servières, & Moreau, 2012)

#RealityCube Classification

Rob Manson, the CEO of BuildAR and an active participant in the AR standards discussions, proposes his own framework for understanding the space. For Manson, the weakness of these previous classification schemas is that they are pre-limited to certain types of systems and do not account for the whole spectrum of theoretical positions. For example, the application of Azuma (1997) would often spark the question of what is ‘real’ (see 1st and Tenth example, above), whether graphics might count as real, and whether that is a deeper metaphysical question. In other classification schemes, real was already operationalized as physical matter, which bracketed that conversation. Some classification schemes that focus on user perception as a criteria, or primarily on the manner of presentation, often ignore the question

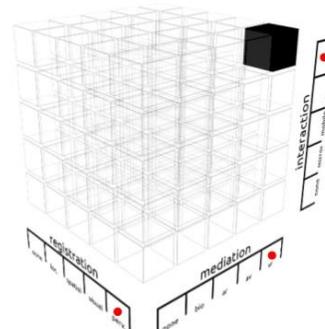
of interactivity. Manson’s AR Cube framework attempts to remedy this with three different axis – 1) mediation, 2) registration, and 3) interactivity. The mediation axis builds on Milgram & Kishino’s (1994) definition, and how your reality is being altered ranging from: none, biological, augmented reality, augmented virtuality, virtual reality. Registration speaks to how the augmentation is triggered, with categories for none, location, spatial, visual, and pervasive. Lastly the interaction axis refers to the users’ relationship with the device, whether it is none, mirror, mobile, wearable, or pervasive. In effect, the cube attempts to place some of Azuma’s (1997) on a continuum, while distinguishing between certain applications and devices. One strength of this approach is that it does not determine a priori philosophical positions like the ‘Brain in a Vat’ question of how we know the world around us is really there and not a figment of our imagination, and could potentially evolve to accommodate hypotheticals (see figure 24). In that way it is much broader than other classification schemas that focus primarily on AR.



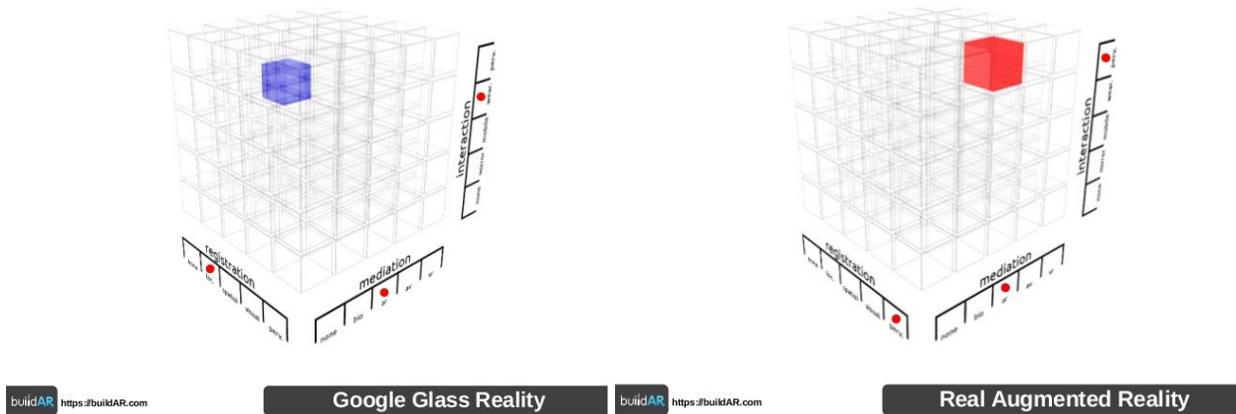
buildAR <https://buildAR.com>

Everyday (Real) Reality

buildAR <https://buildAR.com>



“Brain in a jar” Reality



(Figure 24 - #RealityCube. www.slideshare.net/robman/the-futureis-unreal)

However, Manson’s cube, like all mappings, runs into trouble where the lines are drawn. By simply adding Milgram and Kishino’s (1994) mixed reality spectrum as an axis, it includes augmented reality as a classification category, which we have already seen is heavily contested. While in the context of the cube it is intended simply to denote a proportion of real and virtual, it uses ‘AR’ as both a category unto itself, which is problematic. Second, the categories are discrete but are subject to interpretation, given that many of these are spectrums. One issue would be parsing out the proportion for what would be ‘augmented reality’ and ‘augmented virtuality.’ The other issue would be determining which constitutes ‘pervasiveness’ on both the registration and the interaction axis. Lastly, the cube still labels things as ‘real’ AR, which projects a normative vision of what AR should be.

But is it AR?

All communities engage in boundary work, but the specific ways these boundaries form and are enforced are not inevitable. Understanding these processes requires studying the people engaged in these discussions and the practical politics of how boundaries get defined and defended: “Someone, somewhere, must decide and argue over the minutiae of classifying and

standardizing” (Bowker & Star, 1999; p. 44). The work of classifying through definitions or defining through classifications has significant power, and is a site of political contestation for those who create those standards as well as those who are governed, bound, and excluded by them (Bowker & Star, 1999; Busch, 2011; Lampland & Star, 2009). Over the contested edges of science, those “cultural maps are the interpretive means through which struggles for powerful ends are fought out (Gieryn, 1999; p. 15).”

First, this categorical work and how it is enforced plays a critical role in managing membership and naturalization in the CoP (McPhee & Zaugg, 2000; Wenger, 2000). We see this most clearly in how these definitions are applied and used to challenge specific academic projects looking for entry, but also in more subtle moments of the tacit demonstration of competence and expertise within the AR CoP. They also become a “means of social control, as the borders get placed and policed, ‘scientists’ learn where they may not roam without transgressing the boundaries of legitimacy, and ‘science’ displays its ability to maintain monopoly over preferred forms of conduct (Gieryn, 1999; p. 16). This rigorous control of language is essential to constructing an expert identity, in spoken word and in written text. This is particularly noticeable when scientists attempt to assign other terms for what is not AR, like amplification and automation. This makes “correct technical language correctly used essential to the expert’s claim to professional authority (Marvin, 1988; p. 46).” These classification systems, which further reify these technical criteria, can also serve to exclude people who misunderstand these fine distinctions or are not taught these systems, while reinforcing how elite their domain is (Marvin, 1988).

Second, these classifications carry with them normative value judgments, specifically about what types of work/uses are important and what the technology ‘should’ look like (Bowker

& Star, 1999). This is apparent when certain hierarchies are deployed to distinguish certain applications as ‘strong/hard’ or ‘weak/soft’ AR, and in the classification schemes that declare some as ‘real’ AR, include some that “may not be AR,” and order applications spatially. These mappings can be understood as “arrangements of spaces [that] define logical relations among sets of things: nested, overlapping, adjacent, separated. At the same time, some leave out certain applications or bracket certain discussions due to the criteria they use for their classification schema. Coordinates place things in multidimensional space, making it possible to know the direction and distance between two things (Gieryn, 1999; p. 11).” These spatial mappings render certain things as closer to ‘true’ AR or ‘immersive AR,’ while other objects are moved further away. They also reveal people’s motivations for supporting/opposing/declaring certain applications/devices as exemplars of AR.

Finally, this chapter demonstrates that these boundaries are fluid and are contested by different groups within the community, in particular the boundaries that people enact to defend their domain (Gieryn, 1983). The move by ISMAR to expand the venue to include work from the arts and humanities challenged the understanding of the domain and boundary work that was enacted. Some were trying in earnest to solicit other disciplines, while others were attempting to secure advantages for themselves by using it as another avenue to present their claims. Still others used the ‘scientific’ definition to deny privileges of the space to others who did not adopt that position.

AR is an emerging technology, which imposes a unique set of conditions and pressures. Particularly in technological innovation communities, the positive and negative connotations attached to exemplars have important collective social meanings for community itself as well as important implications for the successful adoption of the technologies. Those within the

community are attempting to negotiate this issue, amongst others, with their definitional claims. This process is not static, faceless, or an inevitable outcome of organizing. Organizational commitments and historical origins have shaped the way certain actors approach the problem, even as ongoing technological developments have sparked contestation over their interpretation of AR criteria. An organizational attempt to broaden the focus of ISMAR elicited many responses, but the definitional battle is the place where contestation is occurring most strongly, revealing the motivations of many different members of the community. Other organizations are basing their marketing strategies around the term AR and advancing looser definitions that include themselves, staking a claim to the epistemic authority to speak for AR. These are met with strong objections and more boundary work by longstanding members of the CoP. Unpacking the specific strategies that people employ to constitute the CoP builds our understanding of theory and organizations.

Through this case, we also see how organizations are constituted by these communicative processes, practices, and debates about technologies. We also see how people engage in these communicative activities for purposes of institutional positioning, individual sensemaking, membership negotiation, and activity structuring. These then get solidified into organizational structure, hierarchy, and systems of classification. Through this, we develop a way of understanding not only how these organizations are constituted, but also how their constitution might ultimately affect the development, trajectory, uses, and users of an emerging technology.

CHAPTER 5

MARKETING, RETAIL, AND ADVERTISING

“If people ever found out what we talked about at these conferences, my guess is that 90% of them would be totally freaked out”

(Justin, Marketing Analyst, Interview)

In its emerging commercial forms, AR is typically accessed through mobile smartphones and tablets. Much of the early promise of AR was that it would allow for unique mobile, locative, and personalized experiences, and content creators could deliver digital overlays on physical space. AR does not exist in isolation, however, in reality it enters a complex ecosystem of existing media where there are multiple obligations, economic models, and systems of convergence (Bolter & Grusin, 1999; Castells, 2001; Jenkins, 2006; Livingstone & Lievrouw, 2006). The possible uses of this technology, however, as well as what content will populate the form, what those augmentations look like, the practical applications for using AR, and the ways it is integrated with existing media, are still being negotiated. This is one of the important functions of these industry conferences, where the companies and institutions working on a technology get together to push its development, modify it, and mold it toward their own goals. It is important, then, to understand who is at these conferences, why they are interested and investing in AR, and the direction they are trying to take the technology.

One emerging relationship that bears watching is how actors in the AR space engage with the marketing/advertising industry. The need for resources from the advertising industry as a driver for media content is not new, rather it is a continuation of a relationship from printed advertisements since the first handbills could be printed (Schramm, 1988). A full embrace of the advertising model were the penny papers that gained widespread distribution through low subscription costs and high income from advertisers. Similarly, free television broadcast is made

possible by the networks themselves, who garner resources from advertisers who want to air commercials during popular shows (Williams, 1990). These business models assume that there is a finite space to reach people and a limited amount of time consumers have, therefore the way to do so is to support and fund industries that produce content people are trying to access. More recently, digital advertising has become a dominant part of the 21st century media landscape (McStay, 2010; Spurgeon, 2008), as agencies continue to look for new opportunities to bring products and services to the public's attention. In doing so, they have also played an important role in the shaping of all modern media, in terms of their growth, trajectory, content, uses, and regulations (Jenkins, 2006; Turow; 2011; Williams, 1990; Wu, 2010).

“the business model for most media industries is underscored by advertisers' demand to reach prospective customers (Anderson & Gabszewicz, 2006; p. 567).” AR is simply the latest technology to struggle with the role of advertising and sponsorship of the media, particularly how it comes to fund and mediate selling content to users, selling the platform to advertisers, and hosting the material content that is produced by marketers. This chapter begins by examining some of the AR browser companies that have become early leaders in the AR space, as intermediary platforms offering services that focus on the connection between advertisers and consumers.

There has been an increasing recognition that a confluence of factors in the emerging technology environment have introduced new organizing logics, stakeholder groups, global economic pressures, and work practices (Powell, 2007). Venture capitalists are also a driving force in this nascent industry, as investors push companies to get big fast (Crain, 2014; Kirsch, Goldfarb, & Gera, 2009; Zook, 2005). This space is characterized by a winner-take-all ethos (Frank & Cook, 1995; Neff, 2012; Powell, 2007), where the need to position an organization and

set the market is amplified. Interested stakeholders are desperate to determine where they think their market will ultimately go, which affects the areas they choose to compete in. One of the things that technology scholars need to address is the interaction between the wider economic, market, and social forces that push technological development (Klein & Kleinman, 2002). These economic pressures and realities are ones that give certain social groups more resources and capabilities, which shapes the relationships in the field. For the past few years, certain AR companies have been attempting to grow and sustain a business model by licensing software development toolkits and building partnerships with the marketing, advertising, and retail industries.

One economic model that these browser companies can embrace is the “two-sided market,” where the readers/users are the loss leaders/subsidized segment that ultimately attract advertisers (Rochet & Tirole, 2003). Users do not pay for the application or any of the content or infrastructure that supports the application, rather the revenue comes from advertisers who wish to reach those users. Some of these browser companies also create products and software development tools for developers to use, with the goal of attracting more content creation to their particular platform. Lastly, some AR companies directly provide the content for AR-based advertising, creating the experiences that advertising divisions/agencies want for users. While these models are typically examined in terms of economic viability and strategies (Anderson & Gabszewicz, 2006; Rochet & Tirole, 2003), what is overlooked is how these pressures and stakeholders are materially shaping a nascent industry as it enters this complex media environment. This relationship is one of mutual interest, where AR advocates are looking for resources but marketing and advertising firms also have interest in entering the space.

One motivating factor for those who invest in AR is keeping up with the latest developments in media technologies and exploring new possibilities for advertising. Mobile devices generally have opened up new channels for advertising that can reach consumers where they are and where they are going (Spurgeon, 2008; Tsang, Ho, & Liang, 2004). As mobile smartphones have become prominent, new capabilities such as mobile internet, audio and video, and near field communication systems like QR codes are now tools for advertising (McStay, 2010). Locative advertising offers the potential for marketers to engage in more targeted and personalized advertising closer to the moment of sale (McStay, 2010).

These emerging technologies complicate the media landscape, as some have argued that media is converging across different forms, contributing to new patterns of media consumption, and reshaping the relationship between media audiences, producers, and advertisers (Jenkins, 2006; Turow, 2011). One consequence of this is that there has been an increase in advertiser interest in content production, as advertisers are always looking for new ways to reach distracted, distrustful, and disinterested consumers (Spurgeon, 2008). From the marketing perspective, AR seems to offer another means of getting users to engage in a new way, and another way to animate their brand through mobile devices using unique visualization capabilities that were previously unavailable.

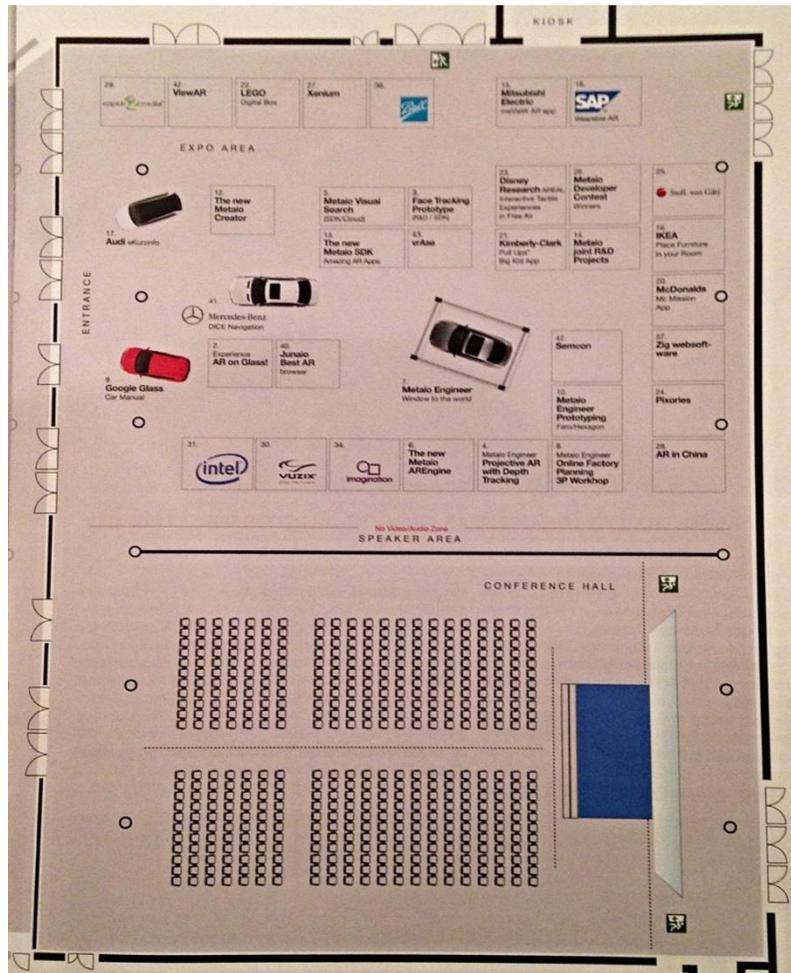
These economic partnerships and pressures already play a dominant role in the structure of the nascent AR industry. The early moments, as actors are making decisions regarding their business strategies and development priorities, can shape the formal properties and distributive mechanisms of the technology in ways that facilitate and pave the way for advertising. First, this chapter examines how the early business models of certain AR companies are increasing centered around the resources provided by marketers, and how they are actively competing for

those clients. Second it looks at how actors are making material and technical choices about the formal properties of AR based on this this orientation towards marketing/advertising. This is not a linear story of an industry co-opting AR technology and deploying it for their own ends; these stakeholder groups are actively shaping and driving the technology. Lastly, it examines how these developments are not inevitable, and how certain groups within the AR community are attempting to minimize the marketing relationship and distance the technology from those industries. These reveal a larger contestation over the future of the technology itself, as well as another rift in the community that demonstrates the importance of economic power within these groups. In doing so, we can begin to understand how advocates and stakeholders are attempting to make AR an integral part of the media environment, how their decisions are affected by economic, social, and cultural factors, and the range of potential social outcomes that result.

Inside the AR Conferences

InsideAR 2013 was a major AR conference in Munich hosted by Metaio, an early industry leader in the AR community. The conference, like many industry conferences, was billed as a place where leaders in the AR industry could come together and show their wares to potential clients and competitors. It was held in the Small Olympic Hall at the old Olympic site in Munich, and there were two main areas of the conference. On the one side was the stage for speakers; on the other side, separated by a curtain, was the exhibition area for product demonstrations and information booths (See Figure 25). AR Summit 2013, a conference held in London, was billed as the premier AR conference in Europe. They had the same divided setup, with a booth exhibition area just behind the speaker stage. Augmented World Expo 2013 (hosted by AugmentedReality.org), had the largest exhibition area that took up the entire hall in the

Santa Clara convention center. There were so many booths that they had to be grouped by large banners into vendor categories – eyewear, hardware, software, and gaming. ISMAR 2012 and 2013 also had demo/exhibit areas where product and hardware demonstrations were occurring, with awards going to the best demonstration.



(Figure 25 – InsideAR Map, Image Captured by Author)

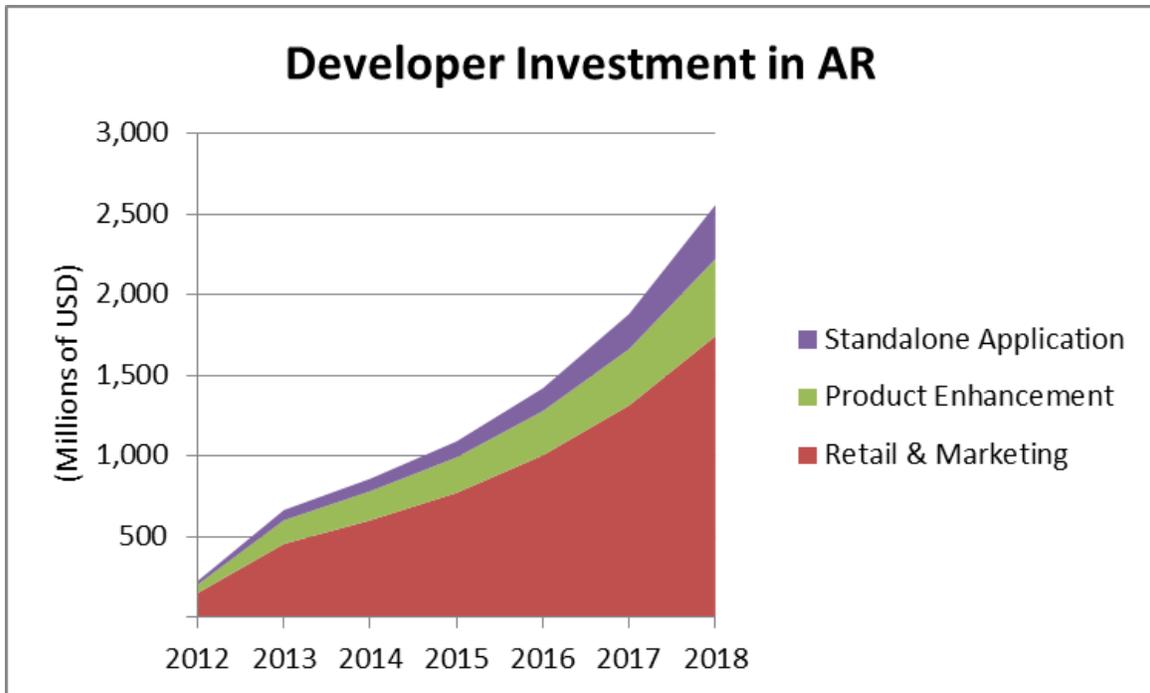
For many, these Expos are the main attraction at these conferences, with all the flashy showmanship of a consumer electronics show, but focused solely on AR. With multiple booths laid out in rows, these types of exhibitions are common in the technology industry. Indeed, spectacular demonstration has been a common tactic throughout history of modern technology, from Edison’s lightbulb (Bazerman, 1999) to early electricity displays (Marvin, 1988), to

Engelbart's 'Mother of All Demos' showing off the personal computer and the mouse (Bardini, 2000). What stood out at the AR exhibitions, however, was how many of the booths were displaying an advertisement or product campaign integrating AR as a component. Walking through the exhibition area, a number of times I was stopped by a smiling booth representative in a sharp suit, politely asking if they could show me their application. Sometimes it was a representative from the marketing department of a large company, showing an AR advertisement for one of their products. At InsideAR 2013, for example, there was a booth showing a Lego box being augmented with a 3D graphic rendering of the completed set; the representative from Lego boasted that it was one of the earliest and most well-known marketing campaigns involving AR. McDonalds was also there, demoing their latest campaign with AR mini-games to educate people about their sustainability practices. There were demo booths from Ball Manufacturing, Kimberly Clark, Audi, as well as many Metaio booths demoing their various marketing projects.

Other times it was someone from an AR design company who actually created the campaign, who was showing off their work to hopefully attract other clients. Patched Reality is one such agency, which works with multiple SDK's to create AR solutions for products like Pepsi Max and Ben and Jerry's ice cream. Explore Engage highlighted their work with Paramount pictures to promote upcoming films and with automotive companies like Kia and Audi to show augmented 3D models of cars. These AR companies are hoping to display their expertise with AR hardware, AR software, and 3D graphics, to sell the technology and their services to potential clients. Another common demonstration was retail uses for AR. Several booths at ARE2012 and AWE2013 were showing applications that used AR as a 'magic mirror' to overlay products on a potential shopper, so they could see how a product might look on them. One such booth was a demo by Total Immersion, a virtual fitting room solution for retail

purposes, where AR would overlay eyeglasses onto a users' face. While one side of the exhibition hall promises AR glasses, another side uses AR to sell glasses.

These partnerships included a variety industries, from entertainment to household items to food and drink, and gave the industry conferences a distinct corporate feel. The prevalence of the marketing and retail I saw in the exhibition halls is supported by reports from the AR companies themselves. Many AR companies are eager to boast about their role in successful AR marketing campaigns, often reporting these numbers in their promotional literature. In the InsideAR conference program, Metaio claimed that after they implemented a point-of-sale in-store AR experience in every Lego retail store in 2010, sales worldwide rose 17% in 2011, in what was considered a sluggish toy market. While this is only a correlation, the article was attempting to highlight AR as the reason for the increase. Client industries are also heralding the possibilities for AR; as an IKEA representative explained, they have distributed 210 million copies of their catalogue, and last year they had 600,000 people use the AR feature in their digital catalogue to look at furniture. While these are self-reports from companies who have an incentive to highlight these numbers, the marketing focus is also corroborated by outside market assessments about the AR industry. In a presentation at InsideAR 2013 by ABI research, a leading market research firm, estimated that about 68% of developer investment in the \$670 million dollar AR market is from retail and marketing, and that proportion is expected to continue well into 2018 (see figure 26).



(Figure 26 – Developer Investment in AR, 2013, ABI Research)

There is also a great deal of work being done at these conferences to convince marketers that AR will be beneficial for their industry. One commonly cited statistic is a research report from a marketing research company called Hidden Creative demonstrating the benefits of AR. That report did a side by side comparison of a 2-dimensional display of a product versus AR marketing on consumers. The first measure was overall time spent: 1 minute 23 seconds in the AR condition, compared to only 12 seconds in the 2D condition. The second measure was whether people would consider buying this toy: 45% when shown the 2D model, but 74% in the AR condition. There was also a difference when they were asked to estimate the price they were willing to pay, 5.99 pounds when shown a 2D ad, versus 7.99 pounds when shown an AR ad. At almost every conference I attended, this report was brought up to demonstrate how AR is almost uniquely suited for the purpose of advertisement – attracting attention to and increasing the time spent with the advertisement, increasing the willingness to buy, and increasing the willingness to pay. A speaker at InsideAR 2013 from Villeroy and Bosh explained: “If you have a real life

experience at home, the point of sale becomes the point of experience. [...] This is possible with AR in a way that it wasn't before, and you can have this in your home not just in your showroom." All of these accounts advance a particular vision of AR, one in which the technology is a promising way of achieving some clearly and rigidly defined goals within marketing – gaining attention, increasing affinity, motivating transaction, fulfilling needs, and securing loyalty.

For mobile AR browsers/SDK providers (Metaio, Layar, and Wikitude are examples of companies that offer both, henceforth just 'AR browser companies'), marketing clients are an important part of the AR ecosystem. Each of these AR browser companies have received venture capital (VC) money, which offers certain growth opportunities but subjects them to the pressures and demands of those particular financiers. Layar received 14 million USD from Intel Capital (Intel, 2010), Metaio received funding from anonymous Swiss investors, and Wikitude is funded by Gamma Capital Partners. Historically, VC places pressure on companies to advertise/market the company and grow as quickly as possible (Crain, 2014; Kirsch, Goldfarb, & Gera, 2009). At the same time, they need to gather developers, users, and clients to sustain a business model for after the VC money, to pay back a return on the initial investment. Just like the web browsers of the early 1990s (Bresnahan & Yin, 2007), their goal is to attract as many users as possible, and in order to do that they have to attract developers to create content with their SDK. This is the chicken and egg problem of a two-sided market: producers will make content if there's an audience, but an audience will only gather if there is content – it is hard to get one without the other (Anderson & Gabszewicz, 2006)

Typically these companies have to subsidize one market as a loss leader, while charging the other group in order to generate revenue (Rochet & Tirole, 2003). The business model for

AR browser companies is to provide end users the applications for free, while charging vendors to host/publish the content. Some also require 3rd party developers to pay licensing fees for their software developer kits/AR tools (see Figure 27). The cost of the users is subsidized under this model, making the developers and publishers of content the profit making segment.

The companies in the space are attempting to both meet the needs of developers, as well as attract them to their SDK's with certain technological tracking capabilities, graphics engines, and prices. This particular slide came from a presentation at AWENY 2014, where one of the leaders of the communities was showing people how they might create certain AR experiences using the various SDK's, and the relative strengths and weaknesses of each.

SDK	Purpose	Tracking	Platform	Graphics	Cloud	GPS	License
Qualcomm <u>Vuforia</u>	2D Images, Markers	NFT, Marker, Text	IOS, Android	Unity3D & Alt.	yes	no	Free
Metaio SDK	2D Images, GEO, 3D, Anywhere	NFT, GPS, 3D, SLAM	IOS, Android	Unity3D & Alt.	yes	yes	0-\$10,000
Total Imm. <u>D'fusion</u>	2D Images, faces	NFT, Face	IOS, Android Dev: MS only	Unity3D & Alt.	no	no	0-\$10,000
Wikitude SDK	2D Images, GEO	GPS, NFT	IOS, Android, Blackberry	HTML, Proprietary	yes	yes	0-\$2,300
Layar SDK	2D Images	NFT, QR	IOS, Android	Proprietary	yes	no	\$3,250/ app/yr + \$20/page
13 th Lab Point Cloud	2D Images, 3D, Anywhere	NFT, 3D, SLAM	IOS, Android	Unity3D & Alt.	no	no	0-\$5,000

(Figure 27 – AR Software Developer Kits, Patched Reality 2013)

Metaio, Wikitude, and Layar offer the ability to augment content based on location, as well as certain objects and images using visual search. Other companies on this list are visual search providers only. They all offer different ways of tracking objects, support different operating

systems, use certain graphics engines, and have different licensing fees and models as they compete for certain types of developers.

From Location to Computer Vision

Some of the first uses that the AR browser companies envisioned were location based features that would allow users to browse the city and their surrounding areas. But this particular use has been slow to materialize. Early marketing applications tried using AR to display their business location to users (e.g. find your nearest Tim Horton's), but they found that AR was a rather inefficient way of doing this compared to other location based applications. Early adopters were not particularly inclined to pick up an application, wait for it to load, and find the right content layer with no guarantee that anything was around them (Olsson & Salo, 2011).

Because of this, in recent years, these browser companies have started shifting their focus away from location based uses and more toward AR that relies on computer vision recognition. They kept the marketing focus, but shifted the priorities of what they were working on. This shift was explained at InsideAR 2013, where one marketing manager was explaining their best practices for AR campaigns and explicitly said they would not be entertaining any more location based or 'scavenger hunt' type campaigns. These were some of the initial implementations for AR, but the speaker explained that consumers were not engaging with these AR locational campaigns, and so these were now non-starters. At InsideAR 2013 the conversation had almost completely shifted from creating AR content for locational purposes (navigational, informational, etc.) to making branded content for users. Being reliant on the marketing resources to stay viable for the interim has changed the mission and emphasis for a number of AR companies. Gerald, one of the founders of a browser company, explained that this shift was a

conscious strategy, as his company started to build ties with print magazines and newspapers to augment print advertising. “In the next 5 years, users will expect magazines to be enhanced. Overall, the biggest thing that I'd hope for is [...] the magazines because that means millions of people have a reason to use AR. [...] And then the industry will move. The magazine industry will move, as well as the AR industry (interview).” This represents one example of how the partnerships have started changing the priorities of the AR companies, away from location to augmenting printed content.

The Marketing Problem Space

Marketers are not just passively waiting for AR to form, but are actively dictating the problem spaces that AR companies work on, while also shaping the technology as companies craft their SDKs around their specific expectations. With the print media partnerships, AR has been used to augment certain advertisements and. Some are even wondering whether AR could revitalize print media advertising, or allow for new ways of storytelling in print journalism (Pavlik & Bridges, 2013). The challenge facing companies trying to expand into that domain is how to mold AR technology to satisfy the needs of print advertising. If advertisers want AR companies to augment magazine images, those companies find themselves pulled into the specific problem space of print media. Ariana, an executive at a major AR browsing company, gives several examples of the types of issues that arise:

We said it needed the image to be rich in detail, with good contrast, no repetitive patterns, with no blank areas. [...] Everything was fine in the design phase but then the company printed the sheet in a different color mode (RGB versus CMYK), which changed the contrast of the image. They also cropped the image to fit it in a limited area and improperly placed it over some objects, all of which caused a loss of features for the key target. Another company even had the creative team modify the feature image later on after they gave us the print. Finally, each company has different printing methods, ranging from black and

white newspapers to a wide range of paper/plastic materials, lamination, ink absorption, and glossiness.

Here we see several instances where the marketers actively dictates the projects, and the client relationship with the AR companies determines the problems that AR companies have to work around.

Some of these then become the perceived limitations of AR technology, places that these AR companies need to improve upon specifically for that space. These problems and goals offer one explanation for what the companies are focusing on and how they are trying to address those through SDK upgrades. For one, AR computer vision algorithms have specific image requirements for size, contrast, and distinctive features in order to be recognized by a device. That becomes a problem when the image changes in the printing phase, due to different color palettes for printing that can alter the contrast. Another example comes from the comment about magazine editors and designers cropping/changing the image to alter the size/features of an image. This is common practice in the print media industry, but it affects the ability of the device to recognize it. Additionally, certain conventions of print media also pose problems, such that newspapers are folded along particular creases that could distort an image. Lastly, the materials that are printed on vary from newspaper to laminated magazine pages, all of which can affect glare and the ability for a device to recognize the image. Because actors put AR in partnership with the demands of print media, these standards and workflows become issues that AR developers have to grapple with to create a satisfactory experience.

This partnership also installs a particular type of user (e.g. readers of print media), which further defines the problem space. For example, this particular use case means that they have no control over the users' environment and conditions, such that a person could be anywhere when reading these materials possibly augmented (indoors/outdoors, variable lighting, shadows, etc.).

Other industries have more clearly defined parameters of use, such as indoor applications where lighting could be controlled. The users themselves also do things like fold or rip the materials, as well as rotate the device and the viewing angle: “Once they get into the users’ hands, all bets are off (Ariana).” In Ariana’s presentation to the IARSC, her message was that AR companies needed to fix these technological problems in AR and improve the recognition software so it can work around a variety of these practical changes in color, size, contrast, lighting, and texture. She was also advocating a standardized system of rules and recommendations across the industry for asset features to accommodate these concerns. Other competing browser companies agreed that these were the same problems that they were encountering when trying to serve marketing clients, and that these standards needed to be embedded in the SDK tools their companies create for people to create AR content, which also makes the job of teaching consumers possible. Indeed these are such a prevalent part of these companies’ business models that even competing browser companies are trying to figure out ways to standardize the technology around these problems. These perceived limitations of AR technology are co-constructed and magnified as core problems through its partnership with the print media industry. These browsers are able to recognize images and objects under certain conditions, but it is the print- specific problems of mass production, mass distribution, and geographically dispersed users that dictate the problem space for AR companies.

This demonstrates the co-constitutive relationship between AR technology and advertising – on the one hand AR companies are developing technologies that specifically focus on improving their recognition capabilities because that’s what marketing clients want most. One example of this comes from Metaio, as one of their representatives explained that when they originally partnered with Lego their technology only supported a limited amount of box tops, but

that's because they worked with their R+D department "improving the tracking algorithms and working with Intel for hardware acceleration, LEGO is now able to augmented almost every box in the store." On the other hand, the limitations of computer vision algorithms require the marketing images to have certain characteristics, sufficient contrast, and noticeable borders in order for the experience to work. Some AR companies even offer a service where you can upload an image and it will rate how well they think their algorithms can recognize that image. If a print media company wants to create an image that they can augment, they have to adjust it to the constraints of the AR recognition algorithm, specifically that it should have high degrees of visual detail and contrast, appropriate size and viewing distance, high resolution, etc. This is not just an instance where new media is remediated to take on the characteristics of another medium, but rather an instance where the properties of traditional media are adjusted around the functions of an emerging technology (Bolter & Grusin, 1999). In particular, these kinds of reshaping by the creative teams and print media practices offer just an early example of how the partnership with AR has caused traditional media to shift their content to accommodate AR technologies. As actors in the space are continuing to work on the technological capabilities and explore the possibilities for AR in print media, this mutual reshaping is an instance where AR allows an older media form to take on a form of hypermedia and look more like the internet. This capability is still being imagined, but the earliest inroads that AR made in the space were not to alter the content of print media but the advertising supporting print media. The design and capabilities of emerging technologies do not just affect the content on those media, but also affect the content of existing media industries.

Beyond the advertising features of newspapers/magazines, another tangible instance of how the marketing relationship is shaping AR technology is in the SDKs that AR companies

license out to creators. Qualcomm's Vuforia engine, for example, is a set of tools for content developers to create AR experiences, and is one of the most dominant platforms in the AR market. Qualcomm estimates that the Vuforia platform is the industry leader, as they claim to have 82% of the branded platform share over their competitors. Part of the Vuforia platform includes a computer vision feature that allows devices to recognize real world images, which then allows people to author AR experiences and link them to the images. Ashton, a leader of the Vuforia team, explicitly talks about retail uses as something the design team is extremely conscious of in terms of their decision making and the choices they make with the Vuforia engine. Ashton explains that the Vuforia engine always had marketing roots: "We started Vuforia three years ago, with the goal of bringing a sense of sight to the device, to explore, recognize, and interact, essentially enabling apps to see. People wanted more retail applications, marketers are trying to drive sales, and they can do that with Vuforia recognizing product packaging (interview)." With marketers as their starting base, the development of Vuforia continued to move toward that direction because of the need to demonstrate value and continue to attract customers. Ashton explains again:

It's a nascent industry. For us, ultimately really, we have to be able to show that there's value in doing these things. [...] I think we said, 'Marketers will find this useful.' [...] Initially it's important to lead with a hypothesis, but very quickly you're going to have to start bringing in support from actual customers (interview).

Ashton reveals a telling instance in which this focus had a direct effect on what the technology can do and what they chose to build into the system:

The first version of Vuforia's computer vision algorithm started with 2-D images. [...] Basically you could recognize some simple images, up to 100. But in the later versions we added simple 3-D shapes. Our team chose to *add a feature for cylindrical recognition specifically because it was called for* by the beverages industry, soda companies and coffee companies. This was an important feature for

drink vendors. Marketers were seeing the value in AR and really to connect consumers to physical products (interview).

This revelation is significant for a few reasons. First it demonstrates that the story of AR is not a neutral or a sequential story, where the designers create the technology and marketing agents then deploy it. Instead, the marketing industry was an active client from the beginning and a significant driver of the emerging technology in both their explicit demands but also their presumed needs. Second it provides a tangible example of how the marketing problem space has a direct effect on the developers of technology and their conscious decision-making process for technological innovation. These represent 'sticking points' in AR, similar to what Hughes (1983) described as reverse salients, or places where certain technological, social, or political barriers emerge that prevent a system from being fully operational, and how the technological development that has already occurred retrospectively seeks out those problems in need of a solution.. Relying on marketers to determine what the technology should do is an instance where marketers become a mediator of design, as they attempt to speak for users' needs and determine how they might/should use the technology (Schot & de le Bruheze, 2003). The fact that this campaign could be implemented with limited changes on the production end also made it more attractive than other types of campaigns, because they have all the infrastructure in place to manufacture these products anyway. As AR actors changed the shapes they can recognize, this enables marketers to simply turn that product into the object that potentially activates a branded experience wherever it goes, which broadens the range of places they can reach consumers.

This is a clear instance in which the problem that is called out by clients has pushed the technology in a particular direction. At the same time, SCOT also suggests that the significance of these moves rests not just in how a technology develops but what alternate outcomes are discarded (Bijker, 1995; Pinch & Bijker, 1987). In this case, the companies who produce AR

tools could have chosen to design the computer vision algorithm to recognize other 3D shapes, or continue to address any number of issues still present in 2D objects recognition (e.g. feature extraction, image capture, sensor capture, color conversion, lighting issues), but the marketing needs of an industry pushed them toward cylinders. They also could have chosen not to include that in the standard toolset, but did so at the behest of these clients. Because of this, it affects the experiences that are made possible through AR, and what developers are able to create – it is owing to this decision that soda companies have created advertising campaigns using their cans as markers (see figure 28).



(Figure 28 – Cool Cola Orange, Ball Packaging Europe)

Beyond the type of object the AR algorithms can recognize, marketing companies are also driving the capacity and quantity of objects that the software recognizes. A lead engineer for one of the AR browser companies explained: “One of the problems right now with computer vision on a device is the limited number of images it can hold to recognize. Maybe in some use cases it is ok to have 10, 50 images that are stored that the device can recognize. But right now that is a limitation if your company needs to recognize more than 50 images (ISMAR 2013, Industry Day).” One instance where a company might need more capacity came from an IKEA representative, who explained that their catalogue had 90 images augmented, but “why can we

not have everything to be able to try out? That should be the ambition I think (InsideAR 2013).” This is an instance where the clients are attempting to set their needs as the constructed ambition for the industry, the place they should be headed and how they might measure success. Some of the leading AR software providers are now attempting to address this need for certain market/retail applications. Metaio and Vuforia have implemented cloud capability for marketers. Ashton from Vuforia explains: “We grew from 50 to 100, to now over a million. We can actually do recognition in the cloud, and part of this is actually to tackle retail use cases where you have so many products, and so many images, that it's just impractical to have all those images on the phone because a) it's too large, and b) it's changing too often (interview).” One of the marketing tips given at an InsideAR 2013 presentation was that companies should continually update their content so that people would continue to have incentives to scan AR content. Because their industry use case wanted to foster continuous engagement with the product and repeated AR interactions, cloud storage was expanded to allow marketers to do just that. This was another reason the AR companies wanted to upgrade the storage for their systems, first to enable the requisite number of images but secondly so that marketers could continually update their content to reflect their changing product inventory.

Besides the recognition capabilities and capacity of AR, other features are designed to serve marketing purposes. One of these is the ability of the AR device to constantly scan the real environment for new triggers/markers/objects. For certain marketing applications, sometimes having an application that recognizes one object and scans it is not enough, particularly if multiple elements of the scene can be recognized. During Metaio presentations, they like to boast that their software now has this ability, something they call continuous visual search. The application they chose for the demonstration of this was retail use case, where a collection of

different DVD's are placed on top of one another, with the vision algorithm continuously recognizing the new DVD cover as opposed to a single initial capture (see figure 29). Continuous visual search is another SDK feature that was developed in anticipation of what they presume companies with many products might want from AR, as well as altering how users are able to experience the technology.



(Figure 29 – Metaio Visual Search Demo, metaio.com)

Through these specific use cases and scenarios, we are seeing how the teams creating the software tools are making active design decisions that attempt to answer and often even anticipate marketing needs, what they want to do with the technology, and the environment in which they need to deploy AR. There are also more subtle changes: for instance, all the SDK companies have built in widgets for developers to insert “Purchase Now” buttons that link to the product on the companies’ website. These tools were developed as a result of marketing demand, while also encouraging the further development of marketing uses. Even the language of the AR companies has become seamlessly integrated into the advertising frame, where AR companies promise marketers that they can help motivate users to ‘pull’ content, which is much more highly

valued than 'push' content. The influence of marketing has affected what types of content/images/shapes the algorithms can recognize, the quantity of images, and the ability to update images, and the continuous nature of the scanning.

These represent instances where AR is integrated into a larger technological system, in this instance another industry and distribution model, which creates major constraints on how that technology should be structured and designed (Hughes, 1983). Klein and Kleinman (2002) observe that while applications of SCOT have demonstrated instances where relevant social groups are successful, they often lack a broader explanation for why some groups' meanings had greater relevance than others and how those were resolved. Williams & Edge (1996) similarly highlight the need for examining the social and economic forces that may shape information technologies. What happened in this case highlights the relationship between industry trends, emerging technology markets, and business models, and how the actors developing the technology respond to those demands. This is a continuation of how advertising and new media develop together and so often converge. Marketers' interest in AR is in line with recent trends to integrate products into media, sync products with multiple devices, and integrate products across multiple platforms (Spurgeon, 2008). Where the advertising industry is attempting to build in more interactivity and content creation, AR allows them to do this with integration with their physical product and across many different products.

Thus far, most of the major stakeholders in the space have been multinational corporations who can afford to dabble in experimental marketing strategies. Not only are they attracted by the promise of more consumer reach, engagement, conversion, and loyalty, they also hope AR technology could more broadly associate their brand with a cutting edge and innovative technology. These companies are incredibly powerful, and have started moving not only the

technology itself, but also started affecting the supporting features that accompany AR applications. Some AR browser companies are pitching their technology not just as a tool for engagement, but also as a method of quantifying the effects of marketing. The promise is that AR will allow marketers to not only gain a higher return on investment, but also deliver better metrics for who sees their ad and engages with their ad.

John Wanamaker, a pioneer in advertising, famously said “Half the money I spend on advertising is wasted; the trouble is I don't know which half.” The way an advertising agency used to assess impact is through imperfect estimates like magazine subscription rates, which do not tell them what specific content people are attending to. The demographics of the subscribers offer another clue, but can also be misleading because many could go to businesses and are then circulated among users who cannot be tracked. Where previously they would rely on these imperfect estimates about whether an ad is reaching their targeted demographic, AR offers a potential metric for that. Already some client companies are using AR technology for the purposes of impact evaluation; IKEA reported that they found activity in the AR application where people would place furniture around their home or apartment increased 35% on Sundays. Peter, a representative from one of these AR companies, explained to me that now they have started charging companies for unlocking the measurable data generated by the applications. This is another tool for ongoing behavioral tracking of consumers that is built for internet technologies, similar to cookies or canvas fingerprinting that pulls images from computers to identify users (Acar et al., 2014), or revelations that Facebook saves not only what you post but also anything you type into the box and do not post (Das & Kramer, 2013). AR takes this one step further, by offering companies a way of measuring which users are engaging with a particular physical object.

Other companies like BlippAR offers clients a suite (called AnalyzAR) that allows companies to track behavior, location, and interaction patterns in real-time. These products suggest that many of the early players in AR are not only willing to promote AR as a marketing tool, but are actively soliciting marketing clients and creating secondary tools to analyze the measurables. This is even built into a potential pricing model, where a sales representative from another AR creative company explained to me: “We might even create some campaigns for free, but we’ll own the data. So if the company wants to see the metrics, then they’ll have to pay. This works for us and for them because the only way they pay is I they’re making money.”

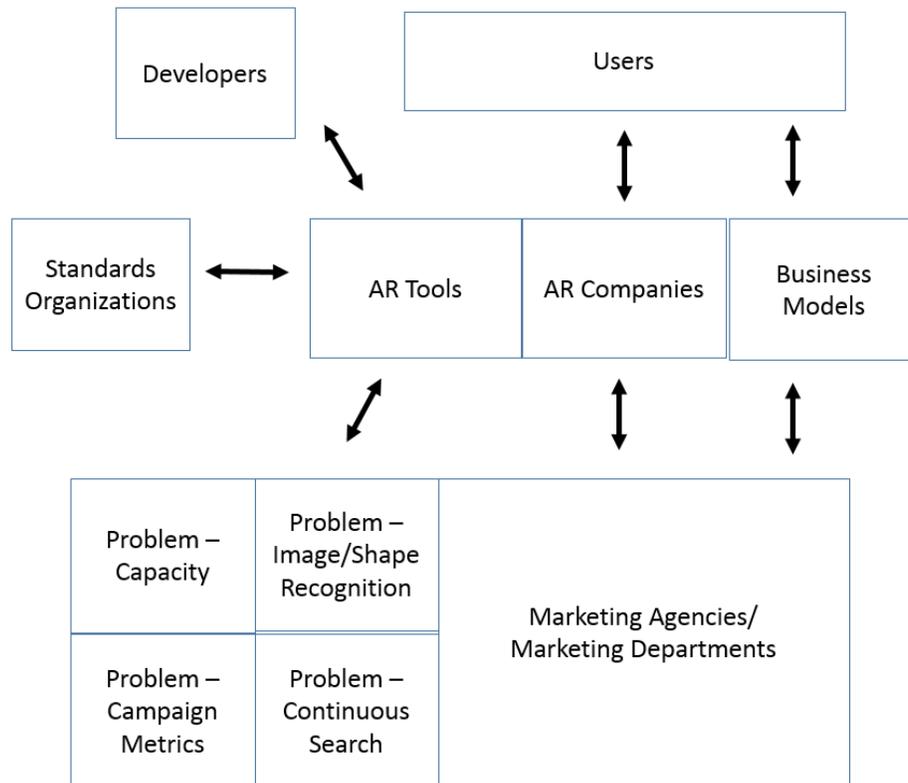
In all of these instances, AR companies are acknowledging that the economic pressures of the emerging technology space shape their practices, specifically that the incentives to get big fast, gain first mover advantage, and generate revenue have shaped how AR companies structure their business models around these marketing resources. A number of representatives acknowledge that their focus has shifted from a user driven model to a marketing customization and licensing model. A few representatives even lamented that they do not wish to engage in this kind of agency work, but that they need to economically. A representative from one company gave a presentation where he flashed the logos of their clients on a slide, but afterwards admitted privately that: “We don’t want to be doing this agency kind of work. We don’t think it is that interesting or driving the industry forward. But we have to because right now it is 90% of our profits.” His comment reveals some of the economic asymmetries and tensions surrounding marketing and AR, and more candidly affirms how large of an influence the marketing sector is for them given how they have structured their licensing business model.

With their business model set up to serve marketing clients, the AR technology and what is possible with the technology continues to move toward that problem space. This is also not

expected to be a one-time investment. Based on industry reactions, it does not appear to be slowing down:

There's a very, very clear tangible benefit to this, and that's why those guys have come back again, and again, [...] I think in retail, in these catalogs spaces, there's definitely been clear evidence that AR adds value to the experience. Our customers are not just wanting an app once, and then they leave. They come back again, and again, so it's very clear to me that they see value in this, whether that's for promoting a product, whether that's for trying to generate sales, whether drive-ins, or traffic, there's clearly a benefit, and value, and that's the reason they're doing it (Ashton, interview).

With more advertising on the way, this relationship is likely to deepen, as AR becomes more prevalent in marketers' thinking about their problems, and AR companies focus more and more on the technology to best serve that problem space. This relationship is depicted below, as the business models of certain AR companies become more embedded with marketing, the tools that developers are able to use also move toward that direction. These companies also are driving standards for AR, which could further solidify these developer tools (see figure 30).



(Figure 30 – Visual Representation of Marketing Influence, Created by Author)

This also has implications for what uses are possible and how users interact with AR, because AR is still rather unknown or loosely defined by most people. The ways that marketing pushes actors, business models, standards, tools, and applications could lay the foundation for how users encounter AR. Because so many marketing applications are centered around their product, brand logo, or other creative image, AR developers must emphasize computer vision and object recognition features, as opposed to other types of location-based or projector AR. This move has implications for the public perception and understanding of the technology, because the definition of AR is still being contested. Within the AR community, a move away from location toward vision could push the definitional boundaries of AR toward that particular orientation, as these applications form a large part of the AR market and are delivering tangible AR experiences. On the consumer side, it could end up shaping the definitional question of AR in a

de facto way, since many people's first encounters and experiences with AR will be with computer vision AR. That might occur while they are shopping and they see a system that recognizes Lego boxes, but with a fluid technology like AR those associations become critical and can harden into definitions over time. The marketing industry is the best positioned group to do this, given that "groups aiming to promote their definition of an artifact must have access to some means of communication, and generally, such access is predicated on the ability to pay (Klein & Kleinman, 2002; p. 45)." This case is not an instance where advertising *about* the technology solidifies a particular social meaning, rather is a case where advertising *through* the technology works to solidify a particular association with and instantiation of the technology.

Beyond how AR gets defined, there are also implications for how AR companies' business models are geared toward marketers in ways that could effectively price out other non-advertising industries from using AR. Ryan, a CEO of an AR creative agency, believes this is a worst case scenario that they might be headed toward:

The worst case scenario for AR [...] is that it's exclusive like a physical billboard. No individual can put a billboard up. It's too expensive. But, brands can inundate you with a message that only they can afford. [...] If it's limited to only the deepest pockets can publish in AR, it will become a medium that's less and less valuable and potentially, without sounding an alarmist myself, enslave us to what information is presented and not necessarily all the information (interview).

Even amongst the marketers, we can already begin to see a trend with the types of companies that can afford to engage with AR right now. It is primarily national and multinational brands that can experiment with an advertising campaign with an uncertain return on investment. For everyday developers, there are tiered pricing models for creating AR applications. Take Metaio, for example: their basic kit is free to deploy with the watermark, while an unlimited SDK removes the watermark for \$3,490-\$5,490. Additional tools, like the aforementioned cloud storage and the continuous visual search, cost a few hundred dollars a month. Supporters of this

model argue that anyone could develop AR, just with the basic toolkit and the watermark. A representative from one of the browser companies explained at AWE NY: “We gotta get ours, right?” Critics argue that these ‘free’ capabilities are limited, and the incentive structure for these AR companies is to maximize the licensing and hosting fees from developers in a way that benefits them while harming the AR industry as a whole. A market report done by ReadWriteWeb research in 2011 found that “more than 1,000 augmented reality projects went live last year, the majority of them created by big, expensive design firms that specialize in AR and charge tens or hundreds of thousands of dollars just to get started (Kirkpatrick, 2011).” Some companies have attempted to win over designers by releasing their software for free, hoping to eventually grow the user market. Which model ultimately survives will play an important role in the trajectory of AR and what is possible through AR, but the scales are currently tipped.

This economic barrier could limit content creation of AR overall, but specifically certain types of applications. Ultimately this will affect the user base that coalesces around AR browsers, as well as their willingness to use them. How this will continue to change is unknown, but it is how early actors have structured and positioned themselves in this space.

Distancing AR from Marketing

Marketing may dominate the exhibition area, but that often stands in stark contrast with the presentations going on in the other room. In the speaker sessions and the panel discussions right next door, entire sessions might go by where no one mentions marketing, unless the session is specifically devoted to the topic. When someone does refer to the marketing contingent, it often occurs as mildly dismissive, or as a precursor to some larger statement, e.g. we need to move beyond the marketing applications to really integrate AR into people’s lives. Many

reactions were often implicit and sometimes explicit attempts to paint marketing as a banal and uninspired application of AR, with sarcastic comments like ‘oh great, another marketing application.’ Instead, most speakers prefer to focus on the latest developments in AR devices, ongoing technical problems and solutions in AR, and exciting use cases like medicine, education, or for first responders. There may be one or two sessions devoted to retail and marketing, but when someone waxes poetic about the possibilities and the future of AR, marketing typically does not come to the forefront. Even at an industry centered conference like AWE2013, the session on retail and marketing was off in one of the secondary presentation rooms. The main theater featured talks about titles about ‘The Augmented World Experience,’ ‘Form Factors of the Augmented World,’ ‘Future AR Technologies.’ ‘Gesture and Interaction Technologies,’ ‘AR Eyewear Face-Off,’ and ‘Cloud, Big Data, and Location Technologies for AR.’ At the AR Summit none of the sessions were specifically geared toward marketing. The speaker sessions focused on ‘The role AR will have in peoples’ business and personal lives,’ ‘What are the technologies pushing AR development forward,’ ‘Enabling the next generation of mobile augmented reality,’ ‘Educational AR,’ and a panel on AR eyewear. At many of the AR conferences, the way that marketing was discussed stood in stark contrast with their presence in the exhibition halls. What these presenters are actively negotiating is not just what the technology is but also what the technology should be. Within these sessions and panels are actors working to align competing interests, distinguish themselves from other stakeholders, and teach newcomers how to promise the technology. The presence of marketing is a point of contention, or at the very least a source of disconnect within the community.

While some AR companies are willingly (or begrudgingly) catering to the marketing/advertising/retail contingent, other stakeholders within the AR community are pushing

back against that relationship. Much of the augmented content that currently is advertising. For some, this is a huge point of concern for AR. Here's Murphy, founder and CEO of a major AR company, articulating what he fears will become of AR: "If all you see is advertising, it's off-putting to some people. I worry that people will literally lump AR in with QR codes. They know how to use it but they don't want to a) because they know it is an advertisement, and b) they have no idea if the QR code might automatically download an application or put their cell phone on some registry list. That would be disastrous for AR (interview)." The comparison to QR codes is a specific discursive strategy, where "connecting a new artifact to an unpopular technology may allow a group to stifle or alter the development of the new artifact (Klein & Kleinman, 2002; p. 46)."

The sheer quantity of advertising was also of concern, specifically that it could potentially crowd out other aspects of the technology. The current model is set up for marketers to be the creators of content, and AR companies would provide the tools needed to create that content, as opposed to alternate models where advertisements might link to or be shown alongside other types of AR content. As an industry that is quite preoccupied with finding the 'killer app,' a prominent topic of discussion at these conferences, this prevalence of marketing underscores the fact that the killer app has yet to be found. Instead, the dominant use is one that might be annoying or superfluous for users, in an environment where a commonly repeated mantra at these conferences is that the 'content is key.' Actors are convinced that the quality of the content is what attracts people to the technology and makes it seem necessary, and in an advertising saturated environment the concern is that "no one would spend money on an AR device if all you could see with it is advertising (Eric, AWE2013)."

Secondly, people in the AR community worry that the marketing applications demonstrate the technological features in the most mundane and uninteresting ways, that do not fully capture the potential of AR. Here's Boris: "Advertising is purely just about attention and that's why the type of investment we give to those is much shorter. [...] The advertising/marketing function is more focused on shiny and simple. It's about the content and not really about the user interaction (interview)."

Not only that, advertising has also been a prominent part of video parodies of AR applications, which play on fears related to uncontrollable and excessive advertising leading to accident and injury. The 'push' element of advertising also shows augmentation as an intrusion which and opens AR up to some of its most dynamic criticisms about safety, overload, distraction, and lack of control over AR content (Liao, 2012). With AR there are people who are worried that the conversation is already being preemptively moved in that direction, particularly because the business models of AR companies that are geared toward attracting those marketers.

Not only could this hurt perception and support of AR, some worry that the advertising association will crowd out other more positive associations with AR technology. These are people who have a vested interest in AR on personal and financial levels, and often have the task of persuading skeptical investors and users that AR will be a major positive force in society. In the speaker halls, most of the discourse is an active attempt to move AR away from advertising and instead link AR to broader goals and societal outcomes. Just one example comes from Gerald, one of the co-founders of an AR browsing company, at ARE2011:

Where's it going to go? I call the phase we are in now the democratization of augmented reality. Of space, the space we are in now, objects, everything. [...] Everybody's gonna make augmented reality. [...] I think augmented reality this way will be an equalizer in the world. Create more of an equilibrium like the web. So information was liberated through the web. Music was liberated through mp3

in any which way. Video was liberated – you can download everything. And I think augmented reality will liberate space.

Those lofty claims, however, are undercut when the primary use for AR is advertising, so instead actors advance visions of liberation, equilibrium, democratization. This discursive work around AR is to link it to particular myths and digital futures that might drive AR forward. At the conferences these are also highly performative, in an effort to build a coalition, where they discursively call upon deeply entrenched moral imperatives to make the technology appear necessary.

Third, opponents argue that many of these advertisements using AR do not improve the product or consumer utility in any way, only possibly the desirability and price. This has raised concerns that the marketing contingent only views AR as a technological gimmick or a short term mechanism for getting attention, one that will wear off as it becomes more commonplace. One of the keynote speakers at AR Summit 2013 warned against this very thing: “The wow factor is a good start, but is only a start. You can’t build recurring business out of it. With the bells and whistles stuff the first time people might say that’s cool, but by the second and third time they’ve seen that already and ask if that’s all AR does.” The fear is that the marketing contingent lacks the commitment to AR and does not care about the future of AR as a technology, rather they are simply using it as the ‘next in line’ technology for their individual promotional needs and will abandon it for the next latest thing. All developers worry about the commitment of their clients, but people mention the advertising industry consistently because they worry about how that industry might push the technology, particularly if their goals are to up the ‘wow’ factor of AR (graphics, color, etc.), as opposed to furthering the development of AR and other features that might actually bring value and utility. They fear the inertia or path dependence that can plague technology, where once the stakeholders begin to move into a

particular problem space, it can become difficult to reverse, take precedence over solving other problems, and gain priority in cases where there are competing needs.

The opponents of advertising have begun to coalesce into an informal coalition around these three concerns, based on how they view the marketers, their intentions, and the effects of advertising on AR technology. At the most extreme end, there are those who view it as an existential threat to the community and the success of AR: “We’re learning in the marketing space [...] how not to use it right now. (Will, interview).” One counterstrategy some are taking for opposing marketing is to combine it with their other advocacy positions such as the form of AR, supporters of HWD may try to link advertising to the mobile form while promising grand futures of AR for headworn devices (see Chapter 3). As most AR is currently accessed through mobile, thus opponents can use the possibility of new devices as a way of both highlighting the limitations of mobile AR and portraying advertising as temporary. This positions the arrival of a new device as a solution to the problems of the status quo, although it remains to be seen why headworn devices would not be just as susceptible to the marketing influence (particularly if their public user base was large enough). Another strategy is definitional, where some actors will define AR narrowly so as to exclude existing marketing applications. This strategy can successfully exert influence in the conferences where similarly minded individuals are in charge of or reflect the membership base (ISMAR primarily). However, it does not address when marketers self-identify their products/campaigns as AR, and the partnerships that are already forming between advertising agencies, corporate marketing divisions and AR companies/laboratories.

A more proactive approach is to advance negative visions of the future of AR and advertising. One of the most prominent depictions is a video called “Augmented (hyper)Reality:

Domestic Robocop,” created by Keiichi Matsuda, an architecture student who creates videos displayed at the AWE conferences. Matsuda explains that his project is an exploration of how:

The latter half of the 20th century saw the built environment merged with media space, [...] taking on new roles related to branding, image and consumerism. Augmented reality may recontextualise the functions of consumerism and architecture, and change in the way in which we operate within it (project description).

The opening scene depicts a dystopia of advertising logos plastering someone’s home entrance (see figure 31). There is also a cacophony of sounds when the advertisements are up, suggesting a nightmare scenario for AR in the future.



(Figure 31 - Augmented (hyper)Reality: Domestic Robocop, Keiichi Matsuda).

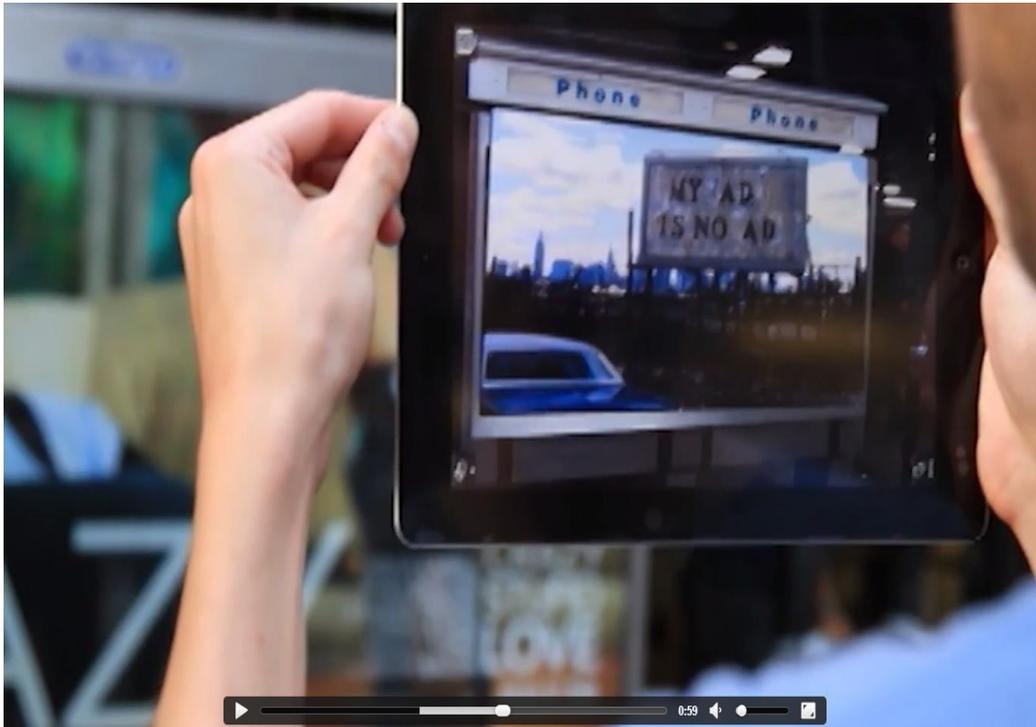
The video is familiar to most people in the AR community, and some use it as a warrant against AR advertising and what might happen if it continues to dominate the AR space. Dennis, a Ph.D. candidate working in an AR laboratory, says the video serves as both cautionary tale and motivating tool: “Imagine Time Square in your backyard. That’s annoying. There’s always concerns about content and control. [These videos] highlights the potential downsides. The dystopian view is there to remind us [scientists] that not everyone is good, and try to not make

this happen.” At InsideAR 2013, another speaker used the video as a call to action, for how the community needs to “stop thinking about AR as packing as much information as we can into the visual scene. We have to let the advertising, the augmented dinosaurs and penguins go, to stop these narratives (Justin).”

The Matsuda video is emblematic of a larger counter discourse against the marketing vision (Pfaffenberger, 1992), where certain actors warn of dystopic futures and subtly denigrate present AR advertising as pointless at best, dangerous as worst. Fear of advertisements are certainly not new, especially in the context of emerging technologies – people in the late 19th century once worried about electricity powering lamps that would beam advertising on the night sky (Marvin, 1988). The lesson there, however, is that those negative futures about advertising did actually shape people’s opinions about the technology and the limits for what they would accept (advertising in the sky was deemed over the line). In the contemporary AR context, these visions are not simply a presentation of a future as a fancier version of the present, rather they are strategic artifacts that attempt to warn people in and outside the AR community of a possible future that we may be headed toward. Taking the relationship to the logical extreme, visual depictions of what that might look like are intended to warn the community against those partnerships.

Another approach is not to contest visions of the future, but to preempt that future by actively deploying AR technology against marketing interests. One of the most important elements of marketing is to control the brand and the messaging behind the brand, and they have enlisted AR as a mechanism for doing so. AR technology can cut both ways, however, and if it can recognize an object to deliver a branded message, activists could attempt to insert their own message. BC ‘Heavy’ Biermann and Jordan Seiler are two AR artists who have been exploring

this intervention, using AR to recognize billboard advertising and subverting its intended meaning through AR (see figure 32). They are explicitly deploying AR to problematize advertising in outdoor spaces, and question the meanings of public places.



(Figure 32 – AR Ad Takeover – Biermann, 2012)

The possibilities are not limited to outdoor advertising, but also using computer vision to recognize any type of product and place counter-information on them. Gerald, a co-founder of an AR browser company, explains:

I like grilled chicken you know nice and light ready you just have to heat it up in the microwave. But you know there's this bio-activist, he found this article by the associated press that in this particular chicken there's anti-biotics. So he just puts that image over it because anybody can augment anything, anywhere. [...] Imagine all the brand managers, the marketing managers, the lawyers you know who try to control this brand. What are they going to do about this if we all do this, see this, and get power through this. It's going to be exciting.

Gerald is describing an instance where different stakeholders utilize the technology toward different ends, and have very different interpretations of that action. On the one hand, people

who create augmentations assert that this is just another means of expressing themselves, and using mobile AR to reengage with and reappropriate the space around them (Liao & Humphreys, 2014). On the other hand, marketers who have considered this possibility will sometimes ask AR companies how they can prevent people from ‘attacking their brand’ or infringing their trademark.

These two different frames of the same activity is another attempt to stabilize meaning and prevent certain applications of the technology. The ability for AR to recognize objects may be helpful for marketing, but only to the extent that the advertisers get to control the messaging. This contestation is still ongoing, as multiple relevant social groups have specific interpretations of the artifact, and what they’re seeing and constructing are quite different objects (Bijker, 1995). In the AR Ad Takeover example, AR is deployed to redress existing power asymmetries in physical advertising, and used to demonstrate that AR can filter that message by using the same recognition tools.

Other worries are that these potential future uses might raise complex legal questions and bring complex issues to the debate. These are instances where people project outward future contestations that might shape the possible uses for the technology:

As branded products become more widely recognized as platforms for receiving digital content, though, those brands will need to be more concerned about who else may be augmenting their content. Trademark law exists in order to maintain the purity of the message conveyed by a mark. The more successful third parties become at repurposing popular marks as triggers for subversive digital commentary, the more concerned the brand owner will be that the distinctive quality of its mark is being diluted. How far the First Amendment’s freedom of speech will allow brand owners to go in regulating such content, however, remains to be seen. [...] You can be sure that advertisers will be clamoring for laws and judicial decisions that protect their ability to control the message conveyed by their brand in every facet of reality (Wassom, 2014).

This usage pits two stakeholder groups against one another with competing rights claims, which demonstrates the next potential area for contestation. The ability of companies to control the augmentations will be an important point of contestation, one where the platforms will have to balance the interests of their stakeholders and property rights against larger public issues of free speech. Much in the same way that the internet allowed for different expressions of speech and did so in ways that challenged legal principles (Lessig, 1999; Solove, 2007), AR is doing so but in a specific brand context (because of its visual recognition capabilities).

Should people's predictions about legal concerns come to pass, that advertisers would clamor for laws to protect their trademark rights, it could be a forum for opponents of marketing to make their case. This strategy is not without costs, however, particularly as Wassom (2014) points out that the legal realm is an unpredictable place to contest competing principles like freedom of speech and fair use versus protected trademark/copyright. There are also other potential interventions marketing companies could engage in that bear watching. Companies could attempt to protect their image from being scanned through technological interventions, similar to the encryption and digital rights management strategies undertaken by the recording and film industries (Gillespie, 2007). On a policy level they could attempt to exert influence with the intermediary platforms of AR, asking them to take down applications and prevent these applications from being hosted. How those issues will be resolved are still yet to be seen, but they can be placed in the larger economic context of the burgeoning AR industry (e.g. if the marketers are paying for AR products, they might have more authority to force a change in policy). For the time being, however, opponents are exploring the possibilities for AR to subvert and deter advertising, hoping perhaps to inspire fear people could turn the technology against their products/goals, and distance marketing from AR.

Lastly, besides using the technology against brands, another approach for opponents of AR marketing could be to proactively push for legal action against marketers. Wassom (2012) has also been keeping an eye on how the Federal Trade Commission is monitoring complaints from consumer advocacy groups that a Doritos AR campaign with virtual concerts was too immersive and potentially misleading. The Complaint and Request for Investigation alleges that:

The [...] experience utilized ‘augmented reality,’ an immersive marketing technique featuring a vivid interactive experience that can be personalized for individual users. Bags of Doritos Late Night chips were printed with a special symbol to serve as a “ticket” for the concert. Flashing that symbol at their webcams would create the appearance of the stage popping out of the bag of chips. [...] Immersive environments use augmented reality techniques to deliberately blur the lines between the real world and the virtual world, making the experience even more compelling, intense, and realistic.

Wassom (2012) writes that there may be a couple causes of action against potential AR campaigns:

One likely candidate is a lawsuit alleging ‘false advertising.’ The federal Lanham Act (which is also the source of federal trademark law) defines false advertising as ‘any false designation of origin, false or misleading description of fact, or false or misleading representation of fact, which [...] in commercial advertising or promotion, misrepresents the nature, characteristics, qualities, or geographic origin of his or her or another person’s goods, services, or commercial activities.’

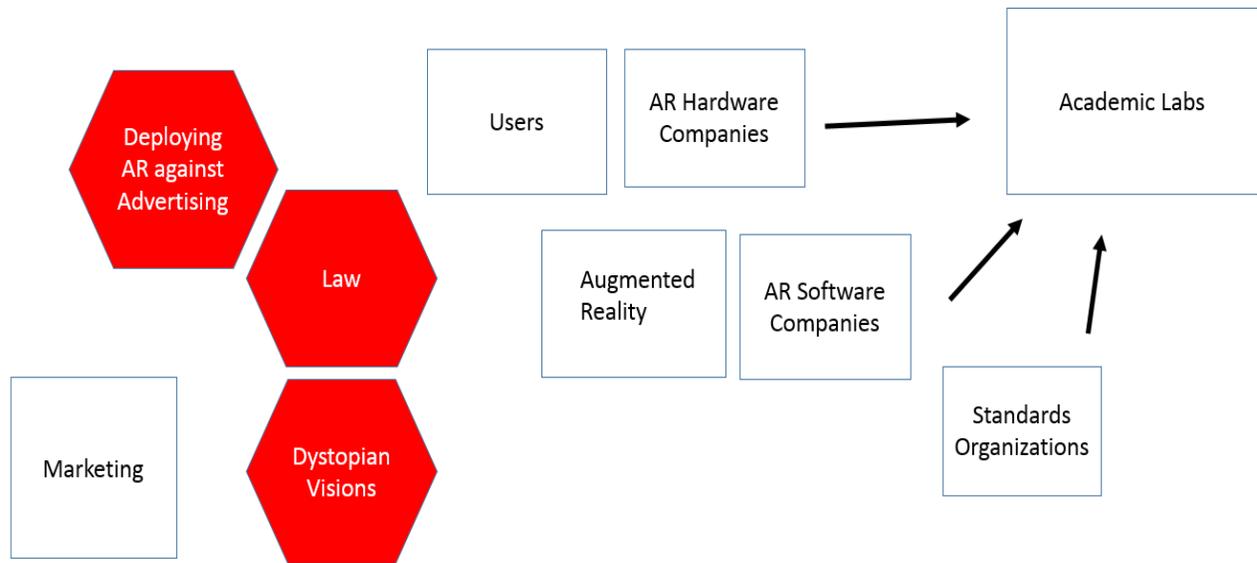
Because AR is often used to digitally enhance a product’s appearance, it would be easy for companies to violate strict false advertising laws. Wassom explains that this could happen either mistakenly, as ads placed on AR browsing platforms “try to convey or recognize more data than they’re able to [...]. Under the wrong set of circumstances, that might end up conveying information that is false and has a material impact on a consumer.” Conversely it could happen if companies use AR to ‘overpolish’ their product in ways that may be misleading.

For stakeholders concerned about the marketing relationship in AR, the legal route could be an effective strategy for minimizing if not eliminating that presence. Bran explains that it is

important to keep a close eye on the early legal rulings here, as a broad interpretation of ‘false advertising’ regarding AR would be a negative deterrent for marketers looking at AR:

I actually don’t know how far they [FTC] would go in critiquing immersive technologies but I’m slightly apprehensive of the kinds of generic statements they might say. [...] What worries me is how a generic statement that might come out of a ruling might be applied to the technology in general. Once you made a precedent that says AR can potentially be deceptive, it could taint the whole field [for marketing], and people just want to keep a safe distance from it (interview).

Thus far actors have been taking a wait and see approach with the legal outcomes and how early decisions might inform actors’ willingness to proceed with AR, but if the legal and regulatory actors intervene it could change the landscape for marketing and AR dramatically. The political/legal option is not without its downsides, however, as the FTC is under no legal obligation to act and can delay action for years pending investigation. Also, there is no guarantee that they would issue a broad ruling for AR advertising that would serve as a deterrent against marketers, rather they could limit the findings to the specific case/industry or create prongs or benchmarks to determine what constitutes ‘deceptive.’ With this particular Doritos case the complaint alleges that the advertising is targeted to a vulnerable population of teenagers, which could become the issue of contention rather than anything specific to AR. Lastly, the FTC could rule in favor of the companies, which might eliminate any deterrent effect and instead attract marketers. Further on down the line, if it came down to a battle of political lobbying, the marketing companies and their clients would have a decided advantage in resources, although there been some early steps for AR to form a trade industry group. Given that there are AR companies with deep ties to the marketing industry, pursuing such a strategy would undoubtedly be hotly contested even within the AR community. Instead of allowing AR to be closely linked with marketing, opponents are trying to distance AR from those uses through a variety of strategies, in hopes of steering AR back toward their original visions (see Figure 33)



(Figure 33 - Visual Representation of Marketing Goals, Created by Author)

Back and Forth with Marketing

Marketers have been largely undeterred by this opposition. They argue that they are a positive influence for AR because they are the ones making the investments and pushing the boundaries of the technology. Online advertising has fundamentally changed the business models of many media industries, while also providing individual corporations with a more efficient means of reaching consumers: “online advertising is disrupting all aspects of the global advertising industry [...] from how creative work is done, to how advertising campaigns are run, and to how advertising is bought and sold (Evans, 2009; p. 38).” AR may be the next frontier of digital advertising, however, as the promise of AR for marketing is that it can turn physical products themselves into a new advertising space. There’s an element of ‘not being left behind’ with AR. BlippAR, a visual search browser company that creates marketing solutions, plays up this concern in their attempts to attract marketing clients: “those brands that consider it ‘futuristic’ or ‘not for them right now’ will be lost in the wake of those that embrace it, as

quickly as six months from now (BlippAR blog, 2012).” It is this motivation that is driving them to risk capital on an untested medium. A representative from ABI market research observes: “Developer investment in AR related projects 650-700 Million US. It is forecasted to grow to 2.5 billion. Most of that has been retail and marketing, not too surprisingly, that has been driving the market (Justin, InsideAR 2013).”

There is, however, some legitimacy to the concern that some marketers are not necessarily invested in AR as a technology, per se. At one of the marketing presentations at InsideAR 2013, Ball Manufacturing was quite realistic about the technological limitations and failures they had when running AR campaigns. They also reported statistics from some of their early campaign assessments, which concluded that the AR ‘buzz’ is good but the campaigns were largely unsuccessful in translating that buzz to sales. Others thought about AR merely as a means of gaining ‘video virality,’ not as an effective standalone marketing device. For the most part, they did not romanticize their commitment to AR as a technology. Even with the findings from the early studies of AR’s effectiveness, skeptics attribute that only to novelty factor. It is these types of attitudes that spur concerns that the marketers will ‘smash and grab,’ damaging the long term potential for those whose futures are more intricately tied to AR.

Direct marketing has simply been the most obvious use of the technology, but some companies are not constraining themselves to a narrow view of AR as just a tool for direct marketing. A few companies have used their experience with AR marketing as a launching point for broadening the uses and applications of AR in their products. Examples of these include Band-Aid, which has paired their line of children’s Band-Aids with an augmented Muppet experience that is triggered by the product, where the Muppets will sing a song while children are getting bandaged (see figure 34).



(Figure 34 – Band-Aid for Kids, <http://www.band-aid.com/magic-vision>)

Ashton talks about this with great excitement, that companies are now exploring other avenues for AR: “Kermit the frog, and some other characters come up, and they play in bands, and kids are have a smile on their face. But suddenly this experience is not about promoting Band-Aids, [...] it becomes a digital feature of the physical product (interview).” The hope here is that AR can extend beyond the outside packaging of the product, and eventually become an integral part of the product that aids in the purpose of the product (e.g. the experience and the process of healing). The visual/auditory qualities of AR are especially well suited to adding meaning to products, which is a major reason they are so active in the AR space.

Another example of this comes from Mattel, which partnered with an AR company that had an exhibit at ISMAR2013. Karen, one of the creators of the AR application, explained that Mattel and other toy companies have been losing revenue to gaming applications on smartphones and tablets. A recent Nielsen (2014) survey confirmed this, showing that mobile/tablet gaming had more than doubled from 9% to 19% from 2011 to 2013, in a sample of users age 13 and older. The hope was that they could pair their physical products with AR to recapture some of that market share. Their toy was a new racetrack that comes with an AR application. The

racetrack set even comes with a place to mount an iPad, on which an augmented dinosaur appears that the toy car would ‘hit’ each time it physically came around the track, eventually defeating the dinosaur. If the car did not go around the track enough times or went by too slowly, the dinosaur wins. By integrating the toy with a smartphone or tablet device, it broadens the possibilities of the physical toy through digital augmentations. The success of these types of augmentations have yet to be seen and they are a small segment of AR in retail, but they demonstrate a broader purpose for AR than just advertising which might address some critic’s concerns.

One company at AWE2013 describes their partnership with a L’Oreal campaign, which augmented the viewer’s own hands to see what different colors of nail polish might look like on their hands. This was described as a hugely successful campaign, with interactions in the hundreds of thousands (another instance in which the metrics from AR are important). Another retail application that has practical benefits is the IKEA furniture application, where potential customers can see a piece of AR furniture in their own house (and adjust the color) to see whether it fits in their space and how it looks aesthetically. These considerations are especially important for large items such as furniture which are labor intensive and costly to move and install (see figure 35).



(Figure 35 – IKEA Catalog App 2014, www.metaio.com)

Some opponents of marketing still gripe, however, that even these ‘useful’ applications are still solidly in the category of marketing. Helping people shop and decorate are only useful in those localized contexts, but it still might portray AR as something frivolous and non-essential.

Furthermore, pitching AR as a data analytics tool for marketers utilizes AR technology as a tracking device for people’s location, behaviors, and what they are looking at. In a data environment where technology has already allowed companies to gather enormous amounts of data and ‘dossiers’ on individual (Turow, 2011), that usage of AR could further cause public concern over monitoring that gives companies data about behavior in people’s homes (e.g. people redecorate their homes more on Sundays).

No Advertising Model?

While many of the current players are supporting and embracing advertising, Google is demonstrating a way that these concerns could be managed at the device and platform level. If AR platforms and devices simply do not allow advertising content, that would limit the

distribution possibilities for advertising. Google's Project Glass, which recently announced that applications that included commercial advertisements would be banned. While the specifics of how they plan to implement this policy are unclear (e.g. what counts as advertising, whether all advertising will be banned or some accepted), it seems antithetical to the rest of Google's business model that relies heavily on advertising. Why did they make such a seemingly rash and counterintuitive decision? First, banning advertising serves a discursive purpose, by rebutting some of the parodies of Google Glass that showed unwanted content popping up and that advertising will overload their users. Second, preventing marketing applications at the start may make room for other forms of content to emerge and thrive. Lastly, it may strategically aim to grow the user base, as opposed to immediately maximizing profit. After the user base is created (one side of the market), Google may then have more flexibility for permitting advertisements later. What Google is struggling with is what all the other companies are dealing with, specifically that there is a need to fund the technology, develop the technology, and legitimate the innovation. The companies who are catering to the marketing contingent have the funding, but that also takes the development in a particular direction and makes it more difficult to legitimate the innovation with stakeholders who are averse to advertising. Google is attempting to bracket that by enforcing a developer policy, but that also limits an important source of revenue. How Google enforces its policy may offer one model for limiting AR advertising, but it may not be replicable for other companies who do not have the same resources to forego advertising revenue at the outset. It also requires an immense degree of control over the system, over the devices and how applications are loaded, who has access to the SDK, and the resources to absorb this untapped revenue.

Conclusion

These debates are a continuation of the debate over print and outdoor public advertising that has been happening for more than a century. AR does, however, introduce new dynamics to the issue of public advertising that are unique to the affordances of the technology. First, where historically the battles for advertising placement occurred over a finite physical space, AR promises a virtual, non-exclusive space where one ad does not preclude the possibility of another ad in the same spot (Liao, 2012). Second, AR content can be delivered by recognizing a users' geospatial coordinates or through image recognition, which offers more ubiquitous potential for advertising and branding to go beyond publicly regulated physical space. Lastly, it has the potential to merge tracking information about purchasing habits with users' location and what they are physically looking at, offering unprecedented targeting ability to an individual. Along with industries that are trying to expand their market, it is also interesting to note that some of the industries interested in AR are those who have undergone significant transition and market losses in the digital age, and are attempting to use AR to return the emphasis to a physical product (e.g. print media, toys, etc.). These companies view AR not as a luxury, but as a necessity, the next frontier to help reinvigorate their industries.

With emerging network technologies, early actors in the space play a huge role in shaping the network, as they make essential decisions with regards to the architecture, regulation, and conventions that determine the trajectory of the technology (Abbate, 1999). These technological systems are socially shaped, and the culture of the producers shape the technology in powerful ways (Castells, 2001). With AR, many of the current early producers and publishers are not the hackers and virtual communitarians that shaped the internet (Abbate, 1999; Castells,

2001), rather it is the commercial and enterprise stakeholders that are driving its economic models and its content.

This has implications for not only the uses of the technology, but the content and design possibilities of AR. It is also distinct from other media industries like radio in the U.S., which originally resisted an advertising model only to eventually centralize around that as a revenue model (Douglas, 1987; Wu, 2010). In its early stages, the AR industry has been tightly structured around marketing resources, promising the technology as such, selling services that fit the marketing vision, and fitting the technology to the form.

While advertising may seem like an inevitable or essential part of media, many scholars have demonstrated that it was in fact specific actors who created the systems that allowed for advertising and incentivized it as a business model (Williams, 1990; Wu, 2010). Flooding the content space with advertising also takes a different path than other historical examples – the purported purpose of advertising in modern newspapers was that it allows journalists to deliver news that informed readers, at a quality that might not otherwise have been possible. Early radio advocates in the 1930's, even some of the major broadcasters, were reluctant about direct advertising because they feared that people would not accept it (Douglas, 1987). With radio in the United States, Herbert Hoover famously proclaimed that “It is inconceivable, that we should allow so great a possibility for service, for news, for entertainment, for education, and for vital commercial purposes to be drowned out by advertising.” As a way of selling without directly advertising, some companies subtly snuck it in through sponsorship of shows and ‘educational’ programming like Gillette delivering a lecture on the history of beards (Wu, 2010). Eventually companies began to sponsor the shows themselves (Douglas, 1987). Stakeholder groups eventually brought legal challenges through the courts to challenge Hoover’s authority to

regulate radio (Wu, 2010). This shift in advertising fundamentally changed the business model and the nature of radio by giving the industry:

An irrepressible incentive not just to broadcast more but to control and centralize the medium. [...] To see why, compare the older model: When revenues came from the sale of radio sets, it was desirable to have as many people broadcasting as possible—nonprofits, churches, and other noncommercial entities. The more broadcasters, the more inducement for the consumer to buy a radio, and the more income for the industry. Once advertisements were introduced, radio became a zero-sum game for the attention of listeners. Each station wanted the largest possible audience listening to its programming and its advertisements. In this way advertising made rivals of one time friends, commercial and nonprofit radio (Wu, 2010; p. 76-77).

This decision gave rise to a particular economic model that shaped not only the content of radio, but the way that the industry is run. Radio is now almost exclusively funded by advertising, while also motivating individuals to change the laws for broadcasting and content creation. Radio demonstrates just one instance where the presence of advertising was not inevitable, but the work of both broadcasters and advertisers. The changing economic model of advertising also had the effect of pushing a particular trajectory of radio, at the expense of alternative models (Douglas, 1987).

Television was also fundamentally shaped by its relationship to advertising, but again the configuration was highly contingent on the actors. This is evident in comparisons of the television industry across continents, where in the United States, advertising “became the central feature around which radio and television were organized, as well as the main source by which they were financed. In Britain it was as late as the 1950s in television, and the early 1970s in radio, that advertising became a feature of broadcasting (Williams, 1990; p. 61).” Advertising was not merely an interruption in the broadcasts, it fundamentally changed the flow of programs, played a role in the themes and content of the shows, and altered the rhythm and the style of the programming (Williams, 1990). This chapter demonstrates how marketing has already bypassed

the question facing these other media forms, namely whether direct advertising should be allowed. What actors are trying to reconcile is the degree to which advertising should play a role, whether the emphasis should be to use the advertising revenue to create and bring other non-advertising content into existence, or continue down a path where advertising is the content. More recently, online advertising has evolved into a multibillion dollar industry, but continues to support different types of content development and changing properties of the web (Evans, 2009; McStay, 2010). The question of how advertising functions in the AR community is the major concern being played out at these conferences, where opponents to advertising fear that tying the technology to advertising from the outset, in the absence of other non-advertising content, could sour people on the technology.

A segment of the community has formed in opposition to these advertising goals, however, and are contesting it by using the technology against advertising, anticipating legal challenges, and advancing negative visions of AR advertising. The result of this contestation will have important implications for the trajectory the technology takes. AR is just the latest place that the advertising industry is exploring, for its ability to reach people where they are and to encourage people to engage with physical products. For better or worse, that relationship is playing a huge role in funding and shaping the business models that are prevalent in AR. This in turn shapes the technology and its trajectory by emphasizing their problems and constraints on the industry. While this fills a resource gap for them, it also affects not only the technology itself but the ways that they are promising the technology.

Their prominent role also has the potential to solidify a particular definition of AR, as well as public associations with AR. The pricing models of many companies are built around these applications, and they may drive other types of applications out of the marketplace. This

could solidify the association that AR is only useful for marketing reasons, or that all the content on AR is advertising. To prevent these associations from limiting and/or damaging the technology, certain stakeholders are starting to rebuke the marketers, by trying to define their applications as not AR, supporting alternate forms, and advancing negative visions of an advertising-heavy AR future. There are also users proactively deploying the technology to subvert branding and appealing to more powerful legal institutions. Unquestionably, the issue will continue to be a major part of the ongoing discussion for the AR community and will motivate future moves within the industry. The contestation over the marketing and retail applications of AR as well as the various perspectives and viewpoints surrounding it reveal how certain stakeholders align around and take action regarding AR, ultimately shifting our understanding of the technology itself.

CHAPTER 6

POSITIONING BEFORE THE STANDARDS WAR

“You remember when back in the day you could write your webpage for Netscape, but if you tried to open it with Internet Explorer it looked like garbage? That’s where AR is now.”

(Riley, Associate Professor of Computer Science, interview)

“There is this vision that you only have standards once everything's set and solved. What we have found, to make good standards that's way too late. What you need is the standards to be formed while the innovation is happening. You need the standards to inform the innovation.”

(Percy, Chief Engineer for an SSO, interview)

The importance of standards in the development and adoption of technology has been demonstrated many times over. Standards are an essential part of all technology, but are especially critical for ‘network technologies’ where the benefits for users are directly related to the number of people connected to the technology (Shapiro & Varian, 1999). This is seen most clearly with communication technologies, where the value of a particular device goes down dramatically if I am unable to communicate with certain people simply because they are using a different type of device. At the same time, there is great economic reward to be had for the winner: if enough people are using one particular standard, there is often a ‘tipping’ point where nearly everyone switches to that standard. Because of this ‘winner-take-all’ goal, it is in these network markets where we have seen the fiercest ‘standards wars’ to capture a complete market share. Some recent examples include VHS defeating Betamax (Cusumano, Mylonadis, & Rosenbloom, 1992), Microsoft Word defeating Wordperfect (Liebowitz & Margolis, 1999), and Internet Explorer defeating Netscape (Cusumano & Yoffie, 1998).

These types of standards wars are so common and consequential that some scholars have devoted significant attention to discerning patterns from them, identifying certain factors that might lead to success, identifying relative strengths and weaknesses, and when to engage in certain tactics when fighting a standards war (Bresnahan & Yin, 2007; Cusumano & Yoffie,

1998; Shapiro & Varian, 1999). Some suggest that preempting other standards from arising through first mover advantage is critical (Shapiro & Varian, 1999), but it is not always definitive and can sometimes be reversed through other factors (Bresnahan & Yin, 2007). Others illustrate how certain early commitments to a standard can be turned into weaknesses due to inflexibility and inertia, and how battles over standards are a sort of technical judo (Cusumano & Yoffie, 1998). What these historical examples make clear is that technological superiority does not guarantee a standards victory; rather, success more often depends on a coalition of actors making strategic decisions and enlisting other material advantages in production, intellectual property rights, powerful user groups, etc. (Cusumano, Mylonadis, & Rosenbloom, 1991; Liebowitz & Margolis, 1999; Shapiro & Varian, 1999). The importance of standards, however, cannot be underestimated. On a technological level, because the crux of many networked systems is interoperability, perhaps no industry is as concerned about standards as the IT industry (Cargill, 1997; Greenstein & Stango, 2007). On an economic level, the value of standards for industries as a whole has been demonstrated many times (Blind, 2004).

From a communication perspective, however, international standards have been an understudied area of communication policy and global media, even though they play a critical role in how content can be delivered, how people use technology, as well as how it gets regulated (Sterne, 2012). Some of that difficulty reflects the messy business of standards, as they are often created out of an emerging, interconnected web of actors, much of it done behind closed doors and conducted under the cloak of secrecy (Cargill, 1997; Sterne, 2012). It is also an international affair, as we will see, which can make it prohibitively costly for scholars to investigate the spaces in which standards are discussed and determined. A few case studies have looked at the practices of particular standards setting organizations (SSOs), to see how the process and structure of

organizations affect the development of standards (Epstein, 2010; Greenstein & Rysman, 2007) as well as the effectiveness of certain SSOs at creating standards in a timely manner (Simcoe, 2007). The strategies that have allowed one technology to become a standard have been retroactively studied (Bresnahan & Yin, 2007; Cusamano & Yoffie, 1998; Shapiro & Varian, 1999), but this also demonstrates a methodological difficulty for researchers studying standards as it is difficult to predict beforehand which standard is worth following, and which one will ultimately be adopted.

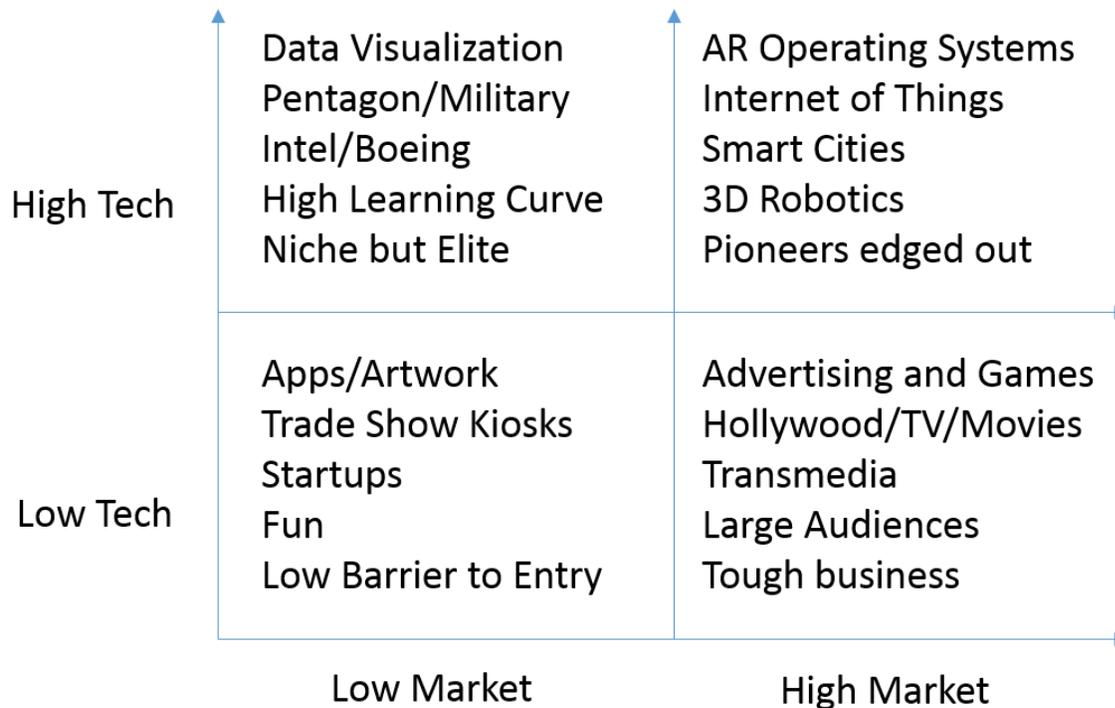
Despite a great deal of literature devoted to how standards battles are fought (and how/when/should companies should engage in them), many of these instances are cases in which fully formed proprietary formats and technologies are battling to become ‘The Standard.’ Much less is known about how a range of standards even become options, or rather how certain stakeholders build coalitions around different standards in anticipation of eventually engaging in these standards wars. Because the standards discussion at this point is not over physical products but over enabling standards that will link multiple devices, many stakeholders are involved. In this particular instance AR is still emerging, so at this stage the standards battle is not over the hardware/software per se, or between two standards with well-established commitments; it is stakeholders guessing at what applications might ultimately be used, the language that will ultimately describe AR, and how people will develop and use the technology, while also being strategic in setting terms that will be amenable to them. The quotes in the epigraph highlight the challenges involved, but not necessarily the best way to overcome them (or even why it is necessary to do so). The standards discussion at this stage is not as well formed as debating specific technological features and benefits (e.g. VHS versus Betamax), rather early actors are trying to anticipate whether a user should have to download a separate browser application to

access AR or do it through an existing browser, what language to use to annotate that content, what capabilities that language should include, what existing standards to build on top of, and the format that content should be transferred in.

One solution being advanced by a group of supporters for a common markup language for describing AR locations, scenes, and objects, called Augmented Reality Markup Language (ARML). In supporting ARML, a number of community leaders argued for the necessity of this approach, drew on specific historical narratives to rally supporters, and aligned themselves with existing technological standards and standards setting organizations. Eventually it was the basis for an implementation by several major mobile AR browsers in the space, who started out as competitors but eventually began to coalesce around ARML to expand their network effects and team up against competing standards and approaches. It took not only a dedicated core of standards advocates, but also an intermediary meeting across standards setting organizations specifically devoted to AR to facilitate discussion and build trust amongst industry competitors. Despite all the work it took to move ARML forward (or perhaps because of the decisions they made), their proposed solution was then opposed by other groups with differing visions for AR. Several actors broke off to work on alternate standards for web-based AR, which would structure AR as an extension of existing web browsers. This chapter explores the early contestation over these standards, the pressures and factors motivating certain actions, and the outcome of these agreements. My aim is to extend our empirical understanding of how standards creation works in emerging technology spaces, how organizations create the standards in anticipation of larger market battles, and the broader implications the resulting standards (and what is left off of it) might have for AR technology as a whole.

Different Standards to Support Different Futures

At AWE2013, science fiction author Bruce Sterling gave one of the keynote addresses about the future of AR.⁴ In it he proffered four different possible futures for AR, depending on what quadrant it fell on two axis – low/high technology on one axis, and low/high market on another (see figure 36).



(Figure 36 – Sterling Keynote, Created by Author)

The status quo is the lower left hand quadrant, where AR is low tech and low market, but the message was that any of these futures could be possible depending on the actions of the

⁴ Science fiction is intentionally omitted in this analysis, and even in this instance where it is being advanced by a science fiction writer, the focus is still on a realistic vision for AR. However, Sterling has been an invited keynote to the AWE and ARE conferences for several years, and along with other science fiction authors (e.g. Vernon Vinge, Daniel Suarez), have played an important role as the keynote speakers and judges for conference competitions. In technology there is a sort of legitimacy to making predictions about the future, and Sterling and Vinge were one of the first to make projections about AR. Their presence as the visionary and expert, along with other science fiction visionaries, speaks to the overlap between fiction and prediction, the desire for any sort of certainty in projecting the future, and the places which desirable (and undesirable) long term futures may come from for certain actors.

community. The quadrants themselves are interesting in that they are intended to be provocative and persuasive, while also highlighting some of the fears people might have for each type of technological future. What these quadrants really call for, however, is a level of standardization currently missing from the AR industry. Percy, an executive at a major standards setting organization (SSO), explained how the quadrants represent very different standards, which will be driven by the priorities of the members:

We're a consortium so the members set the priorities and wherever the quadrant is where the members drive to that's where we, as a consortium, would go. Would the standards be different if we ended up in different quadrants? Yeah. [...] One community that we serve is the military with its precise needs for its precision location and accuracy and semantics and its history of being heavy on the engineering side. [...] If you look at those quadrants for AR and you read the requirements for each of those quadrants, the technical requirements would be different. A standard that is well suited for one quadrant would be thrown out in the other quadrant if it's too complex, too heavyweight. Not suitable for mass market (interview).

The futures implied in each quadrants are incredibly different, and carry with them different assumptions of what standards are needed for possible implementations, goals, and requirements. If moving up a quadrant to high technology/low market is the goal, standards that serve industry actors are still necessary, but they might be need to be more exacting, but with much less emphasis on network effects (see chapter 7). Any move to the right toward the high market sector would also require some kind of standardization across devices, platforms, and content, ones that makes the technology easy to access and create for the most number of people. Many of the actors that are working on the markup language standards are attempting to move toward the high market spectrum.

The choice of standards is hugely important, because with networked technologies a standard is only as valuable as the number of users who adopt it. This is sometimes referred to as Metcalfe's law, which states that the value of a network grows exponentially as the number of

users grow. In these cases, companies that control a significant piece of these networks have the potential to grow exponentially as well (Cusumano & Yoffie, 1998). These standards require a degree of consensus among the participants around a vision for the technology, as well as sturdy coalitions of early actors to help determine the direction that these standards will follow.

For a convergent technology like AR, existing capabilities already rely on a variety of different technologies with thousands of standards. Particularly with AR on mobile devices, preexisting standards for wireless, 3G and 4G data, camera, and other standards are already in place. Beyond the hardware specifications, however, are AR-specific standards having to do with transmitting location, creating and sharing 3D representations, database management, video compression, and file formats that AR might rely upon. Thousands of specifications play a role, ranging from well established standards to ones that are still in development. While some existing standards are sufficient, others are not because they were not created with AR specifically in mind. For this new category of standards, the path they take “depends on the details of the historical sequence in which individual choices occurred, that is, on the path the process of adoption took (David & Greenstein, 1990; p. 6).” Rather than focusing on the battles that occur after different competing standards emerge (Greenstein & Rysman, 2007), this analysis attempts to follow the strategic positioning before the war, namely how an emergent AR standard came to be, the arguments that people made in support of different candidate standards, and how different groups diverged based on different approaches to the problem.⁵

What this account focuses on instead is not the standards themselves as predetermined specifications, but instead on how the technological decisions are made. I examine the places

⁵ This chapter will describe certain basic differences between the proposals, but leave the precise technical differences within particular solutions and technical capabilities to other more qualified to report on those (Butchart, 2011; Kim, Kim, & Perey, 2011; Lechner & Tripp, 2010; Visser, 2011).

where these standards are constructed, the decisions that drove their development, and the coalitions that had to form in order to even make them a potential option. Along the way certain factors play a role in shaping the standard, some strategic and some not, but all of it is the outcome of a social and political process. Some of these decisions may ultimately have important implications down the line, both in determining which standard might be adopted and what people are able to do with AR. This chapter examines some of their stated and unstated motivations behind standardizing, how actors made it authoritative and persuasive to people and organizations, and some of the contestation and implications of the standard.

The Early Market Conditions

Location has been an important component for AR systems, as they must know where the user is and what points of interest are around them. One of the first mobile AR systems to demonstrate the importance of location was created at Columbia University, where a large computer backpack carried GPS and compass information about the user (Feiner et al., 1997). It took almost another decade for those capabilities to scale to a point where they become widely available on a consumer handheld device: in 2008 Wikitude published the first AR application available to end consumers. Layar, Junaio, and ARGON were launched as AR browsers shortly after in 2009. Other AR applications came to market, some aiming to display whatever content was there in a given location, while others sought to find content specific to a location (e.g. videos created in the area). The way that many of these browsers first conceived of AR was in terms of location, i.e. that people would want to pull out their phone browser and see information around them about historical landmarks, apartments to rent, places to eat, etc. This was an

attempt to build on the previous infrastructure laid by geo-locational mobile apps which attempted to do the same, just without AR elements.

In a winner-take-all environment that is characterized by VC funding (Crain, 2014; Kirsch, Goldfarb, & Gera, 2009), there is little incentive to converge on standards. Each AR browser established its own way of storing AR content and serving that content to their proprietary browser (Visser, 2010). Early on, then, there was no interoperability between AR browsers – content could not be shared between browsers, there were no search discovery features, and content developers had to either choose a particular browser platform or format their content and metadata in multiple interfaces and formats (Butchart, 2011). These were the economic pressures that created the present conditions described by Riley in the epigraph, where multiple browsers could not view or open the same content.

Like previous browser battles between Netscape and Internet Explorer (Bresnahan & Yin, 2007; Cusamano & Yoffie, 1998), there are huge incentives for becoming the standard browser – they can set the specific technical features that would define AR experience design and the ways in which users interacted with AR applications. The more users they have, the more developers and investment they attract to their browser. But even after several years, no AR company was making significant progress toward capturing the user market such that they felt the need to engage in head to head competition. Some estimates placed the number of active AR users at 1% of the smartphone market, which is a rather limited market (Tamarjan, 2012). A few stakeholders from these AR browser companies began to blame the lack of interoperability as the reason they were not gaining widespread adoption. With multiple browsers competing for an early market share and smartphone adoption increasing exponentially, the perceived barrier to mobile AR was no longer a hardware question, but a data standard barrier (Lechner & Tripp, 2010). Some

stakeholders began to call for AR browsers to unify around one standard, rather than engage in a prolonged battle over their specific way of doing things.

KML, ARML, and KARML

Because AR browsers offer a specific type of location based service, many had made the decision to build on top of an existing standard for describing location. Keyhole Markup Language (KML) was originally created by Google to allow users to overlay content on maps and services like Google Earth, and quickly became one of the most commonly used standards in describing location (Weiss-Malik, 2008). Google has since turned control of KML over to the Open Geospatial Consortium (OGC), in an effort to promote the spread of KML as an open standard. KML first allows developers to identify a place based on longitude and latitude, as well as other coordinates like tilt, heading, and altitude. Then, it allows developers to place a set of features (place marks, images, polygons, 3D models, textual descriptions, etc.) on top of that location in a web browser or application like Google Maps.

Martin, the CTO of Wikitude, was driving the push towards ARML, explained that KML was insufficient for AR. The basic framework of KML can be useful for describing where AR locations are, but those advocating for a new AR standard argued that “the information available via KML is simply not sufficient to tap the full potential of AR applications (Lechner & Tripp, 2010; p 1).” KML lacks the ability to list the address and phone number of POI in a machine readable manner, which would not allow users to “call some hotels, view some nice pictures, be guided to the hotel address, even view their webpage etc. (Lechner & Tripp, 2010; p 1).” This presumed ‘need’ to dial business phone numbers within an AR experience is itself an assumption of what a future user might want to do using AR, but was one that was accepted as a common

enough feature where AR might need to move beyond KML. To deal with content from a variety of sources, ARML 1.0 also added the field for the content provider, so that each placemark is linked to a particular source. This would allow AR to go beyond simply identifying the coordinates of a particular location, but to allow multiple sources of content to label and add to it however they see fit. An AR application would be able to show provider specific information, separate content layers from different providers, and allow creators to place content from multiple providers onto one display. ARML 1.0 made the decision to build on KML, however, because it is such a dominant standard. As we will see, this has implications down the line for the process that ARML undergoes.

Another possible option came out of Georgia Tech's Augmented Environments Laboratory, called KARML, which they used for their browser called ARGON. Like ARML 1.0, KARML is also an extension of KML, but it focused on different issues with KML that make it less than optimal for AR. According to the specification, "the most important limitation of KML is the lack of control the user has over the display and placement of labels, icons and the HTML content in balloons⁶" The advocates of KARML also note that "one drawback to using KML is the lack of a notion of relative positioning: all points and even the vertices of geometry elements in KML are defined in terms of longitude, latitude, and altitude. (Hill et al., 2010: p. 234)" KARML attempts to fix this by adding the ability to define location relative to another placemark, the user, and place-markers relative to trackable features e.g. markers or objects (Visser, 2010).

The teams behind these proposed markup languages made important determinations about which existing standards to draw upon, as well as what needed to be added in order to

⁶ <https://research.cc.gatech.edu/polaris/content/karml-reference>

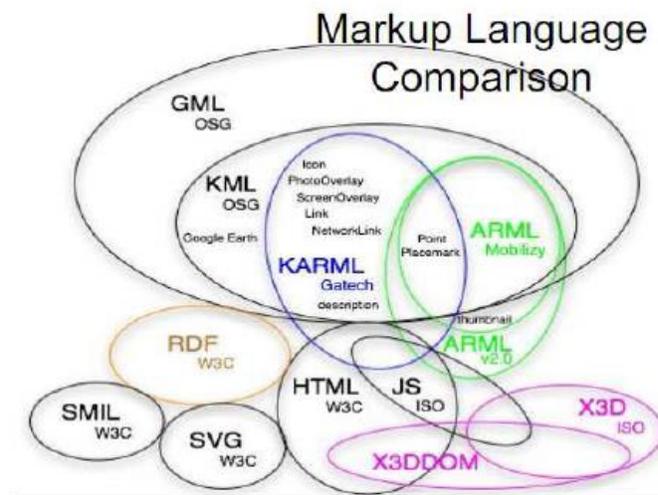
support AR. The authors of KARML explicitly acknowledge that KML was not the ‘best’ location standard that they found, but that they chose KML because it was so deeply embedded into the existing standards environment:

We found many that are more comprehensive and geospatially accurate than KML. However, given our focus on developing tools for general consumption, we felt that the significant penetration of KML into everyday applications such as Google Maps (GM), Yahoo Maps and applications throughout various domains made it a difficult choice to avoid (Hill et al., 2010; p. 236-237).

This speaks to a particular motivation the actors have for AR, namely that the goal is ‘general adoption.’ In the choice of KML, they demonstrate their priority is to move toward the high market spectrum as opposed to up on the technology spectrum, which privileges ubiquity over certain technological capabilities. In thinking about ubiquity and location, the OGC report on the KML standard introduces it as a mass market standard that already serves millions of users, most of whom are non-experts with respect to the geospatial domain. In this example we see the complexity, plurality, and path dependence of the standards environment, where a successful lock-in of a de facto standard can influence actors to build new standards on top of it (David & Greenstein, 1990). Here we see an instance where the choice of what to build upon comes with its own assumptions and consequences for how the technology is likely to be used. The near-ubiquitous assumption that we will use AR to find nearby businesses or hotels justified the need for a new feature to be built into the standard.

The technological relationship between these two proposed markup languages and other standards is portrayed visually, in a diagram presented at the 1st International AR Standards Meeting in Seoul, South Korea (See figure 37). You can see that both KARML and ARML are both extensions of KML, but slightly different in the portions of KML they adopt and the areas they build on. Building off existing standards is common in computer programming; Lampland

and Star (2009) argue that these multiple layers of standardization tend to reinforce each other in a network of nested relationships. This problem is amplified by information technologies, as “standards for computers are perhaps best considered as a complex web of interlocking standards where each depends on the existence of others (Busch, 2011; p. 65).” In this particular instance, two emerging candidate standards for describing AR assets drew upon the same existing location standard specifically because they were aiming for ubiquity and interoperability with existing map based applications.



(Figure 37 – Markup Language Comparison, Lemordant, 2010)

The International AR Standards Community (IARSC)

The 1st IARSC Meeting was held in October 2010, initiated by a collection of people who were interested in pushing AR standards. This was roughly a year after the first AR browsers launched, which demonstrates that at the time of release there was little concern for standards. The IARSC is an interesting organization, as it is not a formal standards organizations or a consortia that develops standards. The IARSC does not have formal processes or voting procedures, nor are any standards housed within the community. Rather, it was convened

because there was recognition that AR is emerging in “standards bodies such as ISO, W3D Consortium, Open Geospatial Consortium and Open Mobile Alliance. [...] A global view of all AR standards activities needs to be maintained, providing those who are implementing draft standards in their projects and products a complete catalog of all relevant specifications and potential liaisons (Perey, 2010).” It has become an important place where different standards organizations could stay updated as to the activities of others. Importantly, it was not a formal SSO that led the development, rather an independent group that attempted to coordinate activities amongst SSO’s as well as actors from a particular industry. These types of non-binding spaces are becoming more prominent in standards and internet governance, and have been found to be important places where discursive and political agendas are put into place (Epstein, 2010).

Meetings were usually two day events, organized and held in international locations (Barcelona, Geneva, Seoul, New York City, etc.) every few months to keep people updated on recent developments and to allow different task forces within the organization to report on their progress. At these meetings, representatives from the different SSOs such as Khronos group, IEEE, OGC, W3C, and others would give presentations about their standards work that might be related to AR. Industry leaders were also in attendance, to explain what they might need from these organizations and talk about some of the challenges that the industry faces. Competing organizations would have an opportunity to meet face to face, and establish a relationship with people across organizations. The chief technical officers of a number of major browser companies would attend these meetings, to discuss their needs as an industry and to coalesce around and find places where standards could be helpful. It was here that important discussions about standards took place, that would eventually play a major role in the standards set for AR;

this group also later served as an important advocacy group for the value of the standards they had developed.

While much of the standards literature has focused on the SSOs themselves and their processes (Greenstein & Rysman, 2007; Simcoe, 2007), there has been much less work on the coordination efforts between these organizations and the significance of these intermediaries. Many technology industries are plagued by an overabundance of SSOs, each creating their own standards that are not interoperable with one another, setting up the conditions for future standards battles that can fragment the marketplace (Cargill & Bolin, 2007). What often occurs is that individual SSOs engage in standards development that might be in the short term benefit of the SSO, or in the interests of their members, but that ultimately push past the ‘carrying capacity’ for standards in a given industry and are detrimental to the industry as a whole (Cargill & Bolin, 2007). This fragmentation was a real possibility for AR, given that AR relied on multiple elements of existing hardware and software standards that were already the domain of many different SSOs. The intermediary space created by the IARSC helped to bring all of these SSOs together to coordinate activities and prevent fragmentation. Particularly because the collective focus was on one industry as a whole, this meeting organized these activities across groups and helped to reduce redundancies in development.

The model of having intermediary groups specifically oriented for industries to coordinate between SSOs was successful in this case, but it was also highly contingent and fortuitous that it was able to do so. Maintaining, organizing, and putting together these meetings was spearheaded by one person, Tina, and it was a huge endeavor. Often it was owing to Tina’s personal connection and initiative that many of the attendees came. In this instance it was one leader that was the driving force behind the meetings, and who also did the secondary work of

cataloguing the meetings online and keeping volunteer activities on track. The documentation of these meetings allowed people to stay updated on developments in relevant standards, and the meetings themselves served as deadlines and motivating factors for people to work on their respective task forces. The group also provided a voice for standardization that would get others in the community thinking about standards and working on standards if they wanted to, as representatives would give presentations at major conferences (ARE2012, AWE2013, ISMAR2012, etc.). Without the IARSC, certain partnerships and necessary alliances crucial to the standards process might not have formed, in particular with competing companies building trust with one another.

In fact, the group almost disbanded under these pressures. At the end of the 9th IARSC Meeting in 2013, in a move that surprised the group, Tina announced that she would not be able to continue running the group. She would not be scheduling the next meeting, and the whole endeavor would shut down unless another organization could help support the group. Tina explained that the work that was now necessary to coordinate the groups' activities went beyond the volunteer efforts and limited financial resources, and that a more institutionalized system would need to take over. In a loose organization of people who are volunteering their time to a particular aspect of standards, the relationships that found and support it can be fragile, and can disappear. As we will see, it was the recurring presence of the meetings that laid the groundwork for many of the developments in AR standards, and it was the revival of the meeting by the OGC that led to a significant agreement between AR browsers.

Placing ARML into a Consortium

Around 2011, an attempt was made to turn ARML into a formal standard (ARML 2.0) that would go through an official ratification status and approval by an SSO. This is one tangible instance where the IARSC meetings were crucial, as ARML formed as a result of conversation with representatives from the Open Geospatial Consortium (OGC) at one of these meetings. George, the chief technical officer from the OGC, explains that the meeting was the key connection: “One of the things that came up then was the AR community standards meetings, run by (Tina). [...] That was how I got associated with AR (interview).” There had been a presentation of ARML and what its proponents hoped to accomplish with the standard, which piqued the interest of the OGC. The IARSC served as a forum where that conversation could occur, and helped move ARML to a SSO.

OGC is a standards consortium, making it distinct from formal standards development organizations (SDOs). It is one of the largest consortia in the world, with over 479 companies as members. Their focus is on geospatial mapping, and their slogan is ‘Making Location Count.’ Cargill (1997) explains: “The difference between consortia and formal committees is a matter of openness; SDOs must admit anyone who appears at their door while consortia have the right to request money as a prerequisite for participation (p. 29).” As a consortium OGC is a member-based system, but also adopts formalized structures and processes similar to those of SDOs, to ensure a balance of interests amongst parties and a democratic governance structure. Certain standards are separated into domain working groups and specific standard working groups. At the meetings each working group is given time to discuss its progress, in between larger plenary meetings where broader motions and standards are voted on. Getting a standard passed through the OGC requires a number of steps, such as the addition of required modifications and periods for public comment amongst the members. Consortia can expedite the process of developing

standards, but their dependence on their membership creates a delicate balance where too much independence could lead to an unused or an overly-constrained specification. (Cargill & Bolin, 2007). The OGC is also just one of many SSOs that have a hand in AR standards-setting, each with their own domains of focus and expertise – some specialize in hardware, mobile, web, etc. The OGC’s interest in ARML stemmed in part from the emphasis on location in AR, specifically the types of applications that needed to know where in the world a user was, in order to access augmented information

Beyond the practical goals of ARML, there was also a specific technological motivation for developing the standard with OGC. As mentioned previously, ARML reused KML concepts to describe the location of particular places. Initially the decision was made strategically to recognize the mass market vision of AR, but in this instance it also made it both easier and more attractive to integrate into the OGC process, because KML is an OGC standard. Richard, an executive of the OGC explains: “the OGC is primarily concerned with location, and we believe that feature will be integral for AR. When Martin came to us with the ARML standard, we saw that it reused concepts defined in OGC's KML standard, so it made sense to create a standards working group within the OGC and have it go through the process of becoming a candidate standard (interview).” Percy, another executive at the OGC, also cited the use of KML as a warrant for OGC’s involvement: “KML became an OGC standard, and some AR folks picked up KML to extend it for AR. That got our interest. They were studying it, to essentially use position based, that is Geo-based AR, and having KML be the document format for Augmented Reality artifacts. [...] Immediately they saw they wanted to grow what their vision was. They needed to base on something like KML (interview).” This technological choice became a political and discursive advantage for bringing ARML to the OGC, because the standard they produced would

likely be more compatible with other members of the consortium and also further proliferate an existing standard of the consortium. Having it develop at the OGC also promised to promote its adoption by member companies and gain credibility. This represents an instance where the layering of standards is not only important technically but also for its institutional alignment.

Concerns about ARML 2.0 and OGC

The initial announcement that ARML was going into the OGC process to become ARML 2.0 raised some concern in the community, about the choice of SSO and the implications that move would have for other SSO participants. The World Wide Web Consortium (W3C) had also begun exploring the possibilities for AR standards, with a working group called the ‘AR Community Group.’ Their approach was a bit different than ARML, namely that AR should be web-based and draw on existing web standards, rather than requiring separate standards for standalone AR browsers. KARML supporters were also uncertain what the move meant for their model. Here are a few of those concerns from Riley:

I feel that this effort is premature. [...] I don't think any of us (including researchers like me) really know what needs to be in these standards and tools, since AR is still not being used by very many people for very many things, and certainly not in the architectural scenario these standards will impact. [...] But when you start talking about "events" I get nervous. I'd much rather see an informal effort by those of us (you at Wikitude, my team, perhaps others) who are actually building on top of KML and building javascript libraries for AR, to work together to be compatible where we agree, and go our own way when we don't, and then see things evolve basic on real people actually doing things with the various browser and so on. (September 4, 2011; W3C Listserv)

Timing is of great concern here: because standards are intended to facilitate interoperability, they also have a strong potential to create path dependence (Busch, 2011). Using the language of standards setting theory (Bresnahan & Yin, 2007), Riley is making a demand side economics argument, that consumers have yet to articulate what the systems actually need. An early

standard could lock-in a certain set of uses ill-suited for what people want to use the technology for, that would persist due to its first mover advantage. This advocacy represents a ‘future technology’ approach to standardizing, where they attempt to anticipate what technologies and features people will ultimately need, negotiate through group design, and (hopefully) produce a set of best practices (Cargill & Bolin, 2007).

A move into an SSO consortium also risks shutting out other approaches to AR, and some within other organizations worried about how compatible a new standard would be with existing web standards:

As for your point about ARchitect/ARML 2.0 being a web based standard...I think a more accurate classification would be as a standard that ‘uses the web as an overlay’. Using a shim that overlays a UIWebView or similar over a video background is a step in the right direction at the presentation level. But this is a long way from where a truly web based AR standard should be (September 18, 2011; W3C Listserv).

This was a part of a larger conversation as to whether ARML 2.0 would be a native web standard, which would eventually be a point of contention between the two standards approaches for how data should be displayed and accessed in AR: whether it should be integrated with existing mobile web browsers or as a markup language that simply ports the web into an AR specific browser.

Others worried about the openness of the OGC process, whether Wikitude was making a move for its own commercial interests, and whether they would be coordinating their activities with other standards groups. Some expressed concern about the origin of the standard as a Wikitude-centric one: “If you have a larger list of efforts, it would have been more useful to include them, rather than making it appear quite so ‘Wikitude-centric’. Folks will be far more interested in contributing if it appears to be more inclusive; as it stands, the document feels a bit too focused on your company, which won't serve you well (September 4, 2011; W3C Listserv).”

There was also concern that Wikitude was only entering the OGC to leverage a powerful SSO to promote their own interests: “While I realize their bias is toward their own commercial interests, it would seem to undermine the position of OGC as a standards organization to have a small group of people leverage them as a platform to promote their commercial product” (September 4, 2011; W3C Listserv). These are not idle concerns. Consortia have ‘pay to play’ requirements, with membership fees ranging between \$500 to \$55,000 depending on the institutional affiliation and level of participation desired. This is also compounded by secondary fees as “the travel and meeting requirements [...] constitute a membership limitation, as very few private citizens have the ability to travel to all of the meetings of an international technical committee where the technology is decided (Cargill & Bolin, 2007; p. 324).” Second, consortia do not have representative participation, they can simply be a group of like-minded companies that focus on the problems they deem important. The fear was that through the OGC process and the de facto participation restriction on the direction of the standard, a new standard could be deployed by member companies as a proprietary strategy for leverage in dominating a market (Busch, 2011; Sterne, 2012).

The Tim Berners-Lee Vision

For ARML 2.0 to gain participation and support from other people working on AR standards, they needed to address these concerns. Martin responded emphatically that now was an important time to standardize, even with imperfect knowledge of what AR would become:

While I do agree that AR is not used by masses of people yet, I strongly disagree that AR standards are still not required. In my opinion, a standard the AR community agrees on will help the industry grow significantly, if (as in *IF*) the standard takes into account that it will require extension in the future. Still, we all know that AR applications are out for quite some time now (with a lot more to come every week), and I guess all of us will agree that they all have significant

overlaps in their functionalities. As far as I'm concerned, this already justifies working on a standard for AR (September 15, 2011; W3C Listserv).

This is one attempt to answer the philosophical chicken and egg question of standards, whether there needs to be a group of users to know what to build into a standard or a standard needs to be in place before users will emerge (Cargill, 1997). Mark offers instead the 'current practice school' of standardizing, arguing that abstracting a specification from existing products is the preferred model. This of course does suit him, as the company he represents is one of the early leaders in the AR browser space; "the current practice school rewards the innovator by allowing a time-to-market and market-share advantage, while embracing stability in the market and rapid deployment of technology (Cargill & Bolin, 2007; p. 325)."

Martin then bolsters this by drawing on a historical narrative, one which he believes will be the analogue to this move:

Figure how HTML was created - it started out with a couple of tags, and I'm pretty sure Tim [Berners-Lee] did not know precisely how the Web will be shaped in the future. Yet, it was extensible, and turned out to be successful. I think that - while the "Web Story" is a little bit different from the "AR Story" - it still makes a good "reference story" (September 4, 2011; W3C Listserv).

This is the story that is often invoked to garner support for ARML; the comparison between AR and the internet was made in the first position paper advocating the standard:

We see an analogy between AR today and the Internet 20 years ago. Before Tim Berners-Lee proposed the first version of the HTML standard for the Internet, numerous ways existed to define the visibility and structure of content, but none of them had a broad acceptance through the various Internet browsers available. Tim Berners-Lee then started with a small set of pre-defined tags what would later become the first HTML standard. Twenty years later, HTML (of course in a very extended version) is still used as the main language for building webpages. Taking into consideration the situation 20 years ago on the Internet and the major step Tim Berners-Lee took, we can see that AR is moving in the same direction now (Lechner & Tripp, 2010; p. 2).

Advocates of the ARML standard often use the Web as the analogy to push for a standard markup language, arguing that HTML was an important prerequisite for its growth. The choice of this particular historical narrative is certainly intentional, and is intended to draw upon several key themes and rally a specific group of people. First, it is one of the canonical stories about the success of standards, strongly making the case that a standard must be in place for the technology to grow. Equating the standard with the success of the technology implies that AR's success will similarly depend on setting a markup language standard of its own. The argument being advanced was that in order to obtain a future where AR is the next Web, the creation and adoption of a markup language is necessary. Objecting to this would delay the benefits for the industry, while also subtly suggesting to competitors who want to hold out that they would be left behind the dominant standard.

One of the factors that contributed to the success of HTML was its extensibility: “Berners-Lee planned to create a hypertext system that would link files located on computers around the world, forming a “world wide web” of information. To the idea of hypertext he added the use of multimedia; his system included not only text based information but also images, and later versions would add audio and visual. The Web's use of hypertext and multimedia drastically changed the look and feel of using the internet (Abbatte, 1999; p. 214).” ARML advocates are capitalizing on this to support their pragmatic approach, that building a simple notation system that is extensible and can be built to accommodate future uses. The Berners-Lee example serves as a direct rebuttal to critics who question the timing of the ARML 2.0 standard, specifically that it can take a narrow form in the interim but build outward like HTML. Drawing on this vision, the timing goes one way – standards are the prerequisite to building a user market,

not the other way around. The worry about unanticipated uses is rendered unimportant, given that Tim Berners-Lee also did not know what users would ultimately do with his annotations.

The Tim Berners-Lee story also becomes a way to alleviate concerns of certain stakeholders, given that the standard was not one that benefited any one sector exclusively, but was instead beneficial for the internet as a whole. Drawing the analogy was a way to suggest that AR could take a similar mutually beneficial path. As a matter of trust, entering the OGC was one way of demonstrating that the result would not be a standard that only benefited one company in the marketplace. Here's Martin:

The second thing I want to discuss is the statement in various emails that the OGC ARML 2.0 SWG [Standards Working Group] was formed because Wikitude has commercial interests in getting ARML standardized. Well - first of all, before a company can join OGC, it has to sign a contract that everything developed in the OGC is available free of charge for everyone. ARML will not be an exception. And to be honest, we never thought about "licensing" ARML - it simply wouldn't work, and it wouldn't be used by anyone if one has to pay for a standard - seriously, who would?! Second: after all, when a new standard was developed, it's not mission-critical for a company, including Wikitude, to define a new standard in the industry it is working in. Call it AR_XYZ-Standard - the company will not receive a huge benefit if it created a standard. The standard is free of charge, that's it. There is fairly limited business potential in a new standard (September 17, 2011; W3C Listserv).

Modelling ARML after Berners-Lee approach with HTML, which was free of licensing, Mark explains that they will go down the same path because the goal is ubiquity. Invoking the name of an internet visionary also incentivizes people to work on the standard. In a field where the original founders' contributions retrospectively mythologized (Streeter, 2011), equating work on this particular standard with that level of impact is a powerful motivator. One of the fears that Sterling highlighted in his keynote is that in the high-tech and high-market future for AR, some of the early pioneers of the technology get pushed out and forgotten. Drawing on the Tim-

Berners Lee story is one way of suggesting that this would not happen, but also to equate their participation on the standard with that of an internet visionary.

With assurances that participation would be open and that the process would not inherently privilege one party (beyond the monetary/travel fees of joining the OGC), the developers of ARML and KARML agreed to work on ARML 2.0 under the charter of the OGC. Boris explains: “Initially [some] were quite against ARML 1.0 as it was a direct competitor to KARML...but in the end it seemed like they accepted that they could work together under the OGC on creating a new ARML 2.0 that was based on both ARML 1.0 and KARML (interview).” Riley explains that the analogy to the Web was a motivation – “it was like the early days with the interoperable browsers, [...] if we had content that we wanted to move between environments, that’s where markup is appropriate (interview).” The technological decisions and the paths they took are deployed as metaphors, which “not only help us think about the future; they are a resource deployed by a variety of actors to shape the future. [...] It is actors who choose to repeat old metaphors and introduce new ones (Wyatt, 2004; p. 257-258).” These particular deployment of past visionaries and metaphors represent important discursive resources to align particular actors and deflect criticism.

ARML 2.0 and the OGC

Once ARML 2.0 became a formal standards working group (SWG) in the OGC, it was subject to certain rules and procedures of the OGC. These include time periods for voting on specifications, and other periods where the specification must be open for comment to the rest of the membership. In preparation of the first draft, the group began working toward making the standard more flexible and less Wikitude-centric, to respond to some of the earlier criticism of

ARML 1.0. Simon explains: “At the first two meetings nobody says anything or makes any decisions, but the overall management of the group quickly broke down the problem set into some chunks that were very workable to try to make it more standards friendly and less obviously a single-header solution (interview).” One move was to integrate some of the KARML standard, particularly the specificity they developed for changing the augmentations and listing their location in relation to the user. Riley explains what he was trying to add into the standard: “In ARML 2.0 the things we pushed into it were notions of association content with more location rich data, not just the longitude latitude, but full position orientation, being able to relate one object to another. [...] So the ability to more intimately associate things with the physical world. So part of the reason I got involved in ARML 2.0 was to push it in a more useful direction (interview).” This illustrates how important who participates is for building support for the standard: early participants were able to mold the standard to accommodate their preferred visions for how people might use the technology. The KARML contingent was pushing the importance of finer level associations between objects for more sophisticated uses of AR into the specification.

Revising ARML 2.0 also meant incorporating more OGC standards, under the guidance of certain OGC veterans:

They did a good job of adopting a fairly flexible and open commercial specification to a very OGC-like standard that incorporated some of the other OGC componentry, such as GML, the Geography Markup language, and others, to make sure that it fit with the rest of the standards. [...] There was some very good guidance from some seasoned OGC members, to make sure the extensibility was not taken to the extreme where it became a single data platform again (Simon, interview).

Percy was one of these OGC veterans, and this was his perspective:

I see it from an OGC point of view. [...] In the sense of AR being a medium, it becomes a new medium for all of the geographic geospatial information that has

been built up into governments and other archives, what people would call "Spatial Data Infrastructure," which is a term for the collective knowledge that we have (interview).

One of the arguments ARML 2.0 supporters made to convince people to go with the OGC was that it would ensure that the outcome was not proprietary or tipped toward a single vendor, but it also resulted in the standard moving toward what the OGC's preferred practices with regards to location. Percy explains: "I remember two points that frame this discussion well about AR. One of them is the position topic and registration is part of that. The other is the semantics. The first one is a solved problem as far as the standards go. We [the OGC] have a real long history of communicating precise geolocation, which we applied to the ARML 2.0 standard (interview)." Evan, an OGC member who participated in the ARML 2.0 process, added: "Obviously OGC has some very longstanding ideas about how geometry and topology are represented according to reference systems. Those things made their way into ARML 2.0 (interview)." What happened to ARML 2.0 highlights how entering a consortium might offer the standard more legitimacy and openness in the participation process, but it also pushes it into alignment with that SSOs preexisting standards. That shaping occurs first because of the expertise of who was available to participate, but also because the institutional structure incentivizes it. Making these changes could also help the candidacy of ARML 2.0 because the OGC membership who must ultimately approve the standard may be more amenable to one that reflects their own established understanding of how to treat location data.

One worry was that putting ARML 2.0 through the OGC might push the standard toward a larger and more exacting location standard (their area of expertise), with complex location semantic descriptions, and away from something usable. Evan, an OGC member that was involved in the ARML 2.0 process reflects on what happened and lamented:

OGC has a lot of people from different groups representing different organizations. [...] It's not totally positive. OGC works best, in my opinion, for things that big national mapping agencies want to do. Maybe not so well for things that are out of the OGC mainstream. [...] Basically, if the goal is something small, specific answers to a specific problem [...] I don't think we're there with AR, [...] and I don't think that fit very well with the OGC processes (interview).

Others felt that going with the OGC meant that location based uses of AR would be emphasized at the expense of other functions: “The OGC wanted to have a major stake in the AR standards land grab so they were happy to provide a home for a standard based on their existing KML. And the G in OGC made it logical they'd be happy with an AR standard based heavily on location (Boris, interview).” Standards organizations and consortium have had turf battles over standards areas; in many ways their standards are their ‘products’ that they have to sell to different companies and clients. The more successful an organization’s standard is, the more prestige they gain as a generator of standards, and the more they have a motivation to spread that standard to other domains. Consortia generally have clear commercial goals in mind (Busch, 2011; Cargill & Bolin, 2007), which also push their interests toward a particular vision of monetizing the technology. These comments coming from outside organizations may be precursors to future battles, as other SSOs pursue their own AR standards.

As the interests of the OGC push ARML 2.0 toward enabling these location uses, it has larger implications for the technology writ large. As mentioned in chapter 4, the definition of AR is still being contested, and standardization that bases AR on location tagging might drive particular uses and applications of AR at the expense of others. The OGC process is another moment in which the visions and uses the authors wanted for ARML 2.0 are made plain. A requirement for OGC candidate standards is that the group articulates future use cases for the particular standard, both to demonstrate a need for standardization and to illuminate what kinds

of AR applications they envision. That document, which I was only able to access as an OGC member, justifies the need for ARML with five hypothetical use cases – 1) a tourism use case for any person, 2) a computer vision use case for any person, 3) a ‘site visit’ use case for a technician, 4) an underground locator use case for a technician, and 5) an emergency responder use case for a police officer. With the exception of the second one, all of these involve scenarios where a particular device would need to know the precise location of the user. The first geospatial use case is the most detailed, documenting how a user of any age, gender, or cultural background might be at a tourist location and wish to select an augmented tour, complete with text, audio, and 3D representations from a variety of authors. The presumed user in this setting would be any person interested in exploring a particular locational site. The emphasis here is the ability to aggregate content from a variety of authors, which points to a mass market usage of AR where the goal is toward the ‘high market’ quadrants. The second computer vision use case is also for public purposes, as it documents a user trying to find an instruction manual for a coffee machine. Use cases three through five are specific to how outdoor workers could utilize location-specific AR.

These document is an important artifact for understanding the standards creation process. First, it demonstrates the way that these standards must be articulated around typical applications, procedurally as well as practically. These use cases justify the need for the standard as well as framing the problem that they are trying to solve, and in the process of developing the standard these use cases serve as points of reference for what they should be doing (as well as what is out of scope). The tourism use case, for example, envisions not only a system that determines location and displays content, but allows users to place their own augmentations around them, send photo compositions with augmentations to social networking sites, and

rate/comment on the space around them. The purported uses that prove compelling then become the goals that the markup language is designed to serve. Second, these broad use cases (written by the authors of the standard) incorporate many assumptions about the types of devices that might be used to access AR. Evan, the author of many of these use cases, explained that this was strategic and that people were trying to insert their own preferred use cases:

They're related to understanding why you want to do AR in the first place. That's partly use cases that lead to people getting some benefit. [...] There's certainly a value in information about your surroundings as you're traveling around to know what's there. [...] It's perfect to me, anyway; everybody loves their own use cases (interview).

In the end, all five use cases were specific to mobile device usage, as opposed to HWD or projector AR. Here is an excerpt from the first use case:

The user is of any age, gender or cultural background. The user is new to a particular location [...]. In advance, using the latest technologies, the site's management has created interactive experiences for mobile AR platforms using historical information in one or more of these formats: text, sounds (speech), images, 3D models and animations. [...] Given the position and orientation, the system welcomes the user to the site. [...] By pointing at the building (or where a building was), a reconstruction of the original site (or the building) appears. [...] The user wants to be in a photo with the statue. This is done using the user-facing camera. The composition results with the statue, the user and the fictional rendering of the artist leaning on the other side of the statue. The user sends this photo montage to friends via a social network. This portion of the use case can be extended to include capturing a video of the user, location and digital element (ARML 2.0 Use Cases, 2012).

This is a pretty standard use case for mobile AR, but notable in that it attempts to be as generic as possible in terms of its user. In support of the mobile form, their advantage is ubiquity of device, therefore it is to their advantage to posit an everyday example that 'a user of any age, gender, or cultural background' could participate in. It also projects certain uses that anyone could engage in, like pointing their mobile device at a particular object. Even certain deployments in these future visions necessitate a mobile device, e.g. a user-facing camera, which

a HWD would likely not have. This also presumes certain future uses that people would want to engage in, such as using mobile AR to create and share fictional AR renderings with them in the scene.

Here we see how in the procedural justification over standards, these futures work to not only solidify particular uses and users, but also to secure conditions that could ultimately intervene in other disputed areas of AR (like the device question of mobile versus HWD battle in chapter 3). Both of the public use cases listed in support of ARML 2.0 were written with explicitly commercial purposes in mind, which standards consortia generally take seriously (Busch, 2011). These ‘presumed use cases’ also emphasize the locational element of AR and uses in which people and devices need to annotate location, which is the primary focus of the OGC. Two of the use cases even specify that the users would be “equipped with a connected mobile device equipped with one or more positioning technologies, accelerometers, compass, and camera (ARML 2.0 Use Cases).” These offer insights for the vision that ARML 2.0 is trying to enact and the uses that they hope to serve, mediated through the OGC process.

ARML 2.0 and the Common AR Interchange Format (CARIF)

The ARML 2.0 standards working group within OGC has been making progress toward becoming a candidate standard, and have released the specification for public comments by other OGC members. As I write, ARML 2.0 is now in the implementation phase of the approval process, where the OGC requires that there be real world applications of the standard. Boris explains: “For ARML to become an official OGC standard they require two independent implementations and up until recently there has been no real commitment from anyone [...] to really do a full implementation (interview).” The perceived intractability of the browser

companies to implement ARML 2.0 was a barrier, as they would be the primary benefactors of the standard. The cost of converting all of their data and content into ARML 2.0 was one concern, while some wondered if it would even be possible:

You could kind of make a Frankenstein monster thing out of the pieces from those companies plus others [...] Super expensive, that's the problem. [...] And some of the things might not be convertible, [...] say in one you might have local references with respect to some landmarks, rather than using latitude and longitude. You might not have the data that you need in order to make the conversion between the local, relative references and latitude and longitude. You have one dataset that's all reduced to WGS [World Geodetic System] 84 latitude and longitude, and you have another that's lots of stuff in a local coordinate system with respect to some landmarks (Evan, Interview)

Because browser companies are based all over the world and have content that spans all over the globe, differences in how they already describe location might make the conversion prohibitive in terms of cost. Evan explains: “You could do a survey, which would get you enough control information which would allow you to make the conversion, but it's not a matter of just running some algorithm on some numbers, and turning them into some different numbers (interview).” Thus far, none of the companies have fully implemented ARML 2.0, and there has not been much movement on that front. This is a byproduct of the economics of early stages of emerging technology development, driven by a winner-take-all environment, which actively disincentivized standards cooperation and created silos that became difficult to reverse through standards. Each of the different browsers all used their own approaches that make it more difficult to standardize, or agree on what the standard should be. Sharing that data and their particular approaches could also represent a competitive loss for each company.

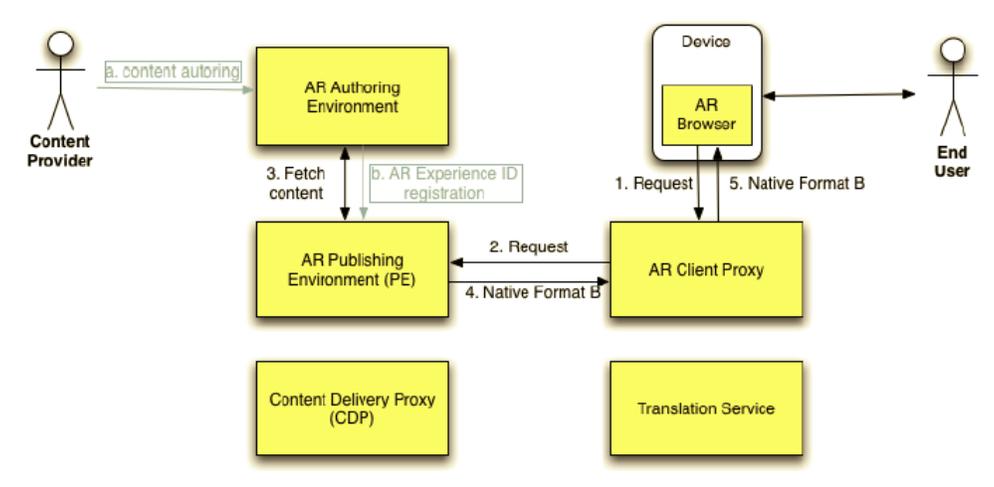
While the first few years were spent trying to capture market share, there were increasing economic incentives to standardize across the AR browser companies. Because of this, standardizing could ultimately increase the number of users they could claim their browsers

could reach, which would in turn increase their attractiveness to developers and marketers (see chapter 5). Second, these companies were faced with increasing competition from other AR companies. For instance, some companies like BlippAR and Vuforia were working on systems that did not do locational AR, instead solely focusing on vision based AR. Other browsers like CityLens by Nokia had partnered with Microsoft, and came pre-installed on Windows phones like the Lumia 920. That particular strategy comes straight out of the playbook of the browser wars between Netscape and Internet Explorer, where distribution on an existing operating system and platform was viewed as the primary advantage for a browser that did not have first mover status. Internet Explorer came preinstalled on the Windows Operating System, which resulted in a wave of new users that were able to slow Netscape's market advantage (Bresnahan & Yin, 2007). This strategy is not without its downsides, however, as it ended being the focus of a multimillion dollar lawsuit. AR browsers are attempting to catch up on the distribution by building their own partnerships – Layar was preinstalled on the Samsung Galaxy S. In addition to competition from other companies, the group in the W3C was pushing a web standards approach, which had the largest current distribution on existing browsers.

Faced with these other models, coalescing around ARML 2.0 was one way for AR specific browser companies to unify their user base/existing content to solidify their first mover advantage. In the summer of 2013 and into the fall of 2013, there started to be some movement on the implementation front for ARML 2.0. However, it was also in the early summer of 2013 that the IARSC had disbanded, leaving a void in the space for coordinating these activities. As part of an effort to revitalize the IARSC, representatives from Wikitude and SEAC02 (a mobile AR browser based in Italy) began discussions with members from Layar and Metaio. Short of a full native rewrite and implementation of ARML 2.0, however, which there was not much

enthusiasm for, they explored another option, which was to use ARML 2.0 as an interchange format. This would allow companies to still have their own native format but have an external translation service convert their content to ARML 2.0 so it could be read in other browsers. These discussions would eventually revive the IARSC from its hiatus, as the OGC stepped in to house the meetings, which made sense given their investment in AR. There was some symmetry to this move, because it was owing to these meetings that the standard was placed in the OGC, and now the ARML partnership would absorb the IARSC. These moves demonstrate how contingent these standards are, as a result of complex institutional requirements that enable these standards to develop.

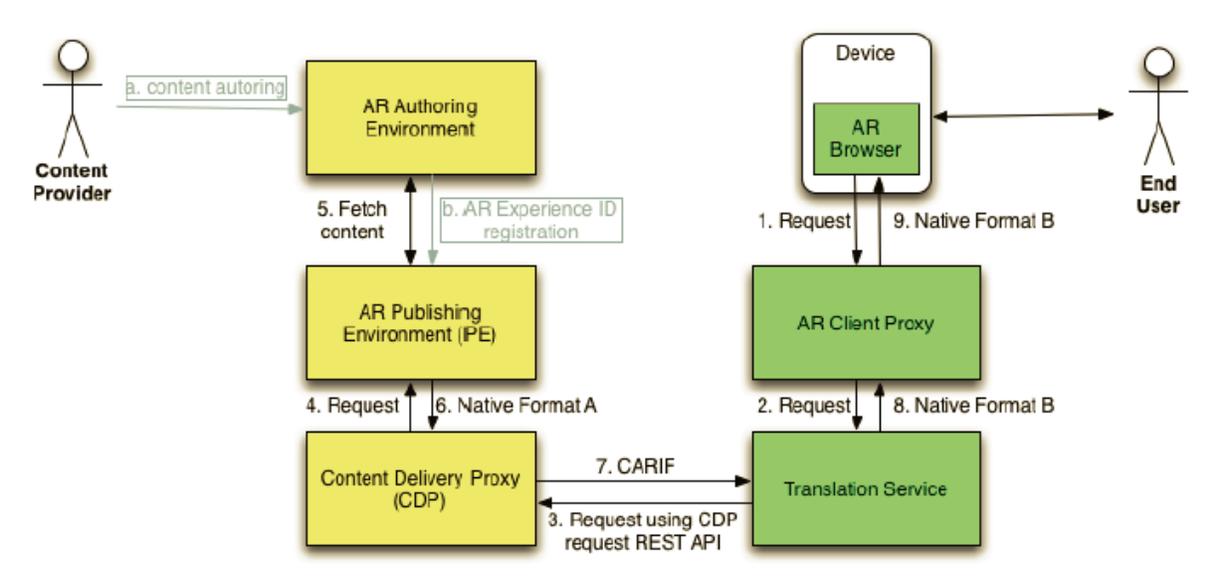
Figure 38 comes from the draft document, outlining how the current browser system works where browsers can only accommodate requests from their native environment, and content providers need to write for that format to publish in that system.



(Figure 38 – Draft Document of the AR Browser Interoperability Engineering Task Force, www.perey.com)

The proposed implementation introduces two additional components that would allow requests to be handled across browsers. In the event that a user attempts to access content through

browser A that is hosted by browser B, the client will reroute that request to a translation service that then fetches the content, converts it from native format B to a system called the Common AR Interchange Format (CARIF), which can be read by browser A and delivered to the user (see figure 39). This particular solution was crafted to make things interoperable, but also to protect the interests of each browser company. Tina explains that in this system, companies do not have to share their content or data with their competitors. This was a major point of contention in the conversion discussions, because most stakeholders wanted to preserve their own AR content and compete on the content/developer level while still being interoperable on the user side. This core interest to protect their content represents an economic barrier that needs to be accounted for, even as they sit down to cooperate. The solution also means companies do not need to engage in the costs of converting all existing data into ARML 2.0, instead they just need to make sure that the basic features of their content (id, location, title, image, description) and action type (open website, play video, play audio, send email, call a phone number, send a text message, route to destination) can be converted from one browser syntax to another.



(Figure 39 – Proposed Interchange Format, AR Browser Interoperability Engineering Task Force, www.perey.com)

The announcement of CARIF via an OGC press release heralded this development and movement into the implementation phase. The three largest browsers in the world (Wikitude, Layar, and Junaio) agreed to this implementation, which they argued would attract a much larger audience to AR and give users a larger catalog of AR experiences to discover and explore. This again speaks to the mass market goals they have for AR browsing, in particular moving more towards ubiquity and transmedia. The ability to convert images, videos, and audio files across platforms allows content publishers to “be able to offer AR experiences with their content to larger potential audiences (e.g., all users of AR browsers that support interoperability) with equal or lower effort (costs) of preparing/producing AR browser-based experiences with their digital assets” while content developers can “invest in innovation [...] in preparation of highly engaging and interactive experiences (OGC Press Release, 2014).” All of this, they promise, is intended to secure investments from content publishers and attract developers. The revenue stream gained from pooling and attracting more users is particularly important to them, as many browsers are not sustaining as many users as they had initially hoped (Olsson & Salo, 2011). Kirk, an executive for Layar, explains that cooperating on standards was critical to building that base and realizing a critical mass of AR users:

AR is a new medium that will change the way we look at the physical world, linking it to the digital world. In order for this medium to become ubiquitous and easy to use, it is necessary to create standards so that content publishers can rely on their creations being viewed by the largest possible group of end users, regardless of the application they use; like the web that can be browsed with multiple browsers thanks to the W3C standards (OGC Press Release, 2014).

Again the Tim Berners-Lee narrative is referenced in the public announcement of the agreement, through the web standards. Representatives from three competitor browsers ultimately united

under this same web-browser analogy that had been pushed from the beginning as a prerequisite for a larger market, but protected their own interests with a standards solution that does not require a complete rewrite of their native systems. This early decision to move toward CARIF as a basic feature system but not a full conversion of native systems sets up a situation where ongoing competition will necessitate additional cooperation over standards in the future if they want AR to become fully interoperable.

First, CARIF does not standardize the creation of AR content, as people are able to write in ARML 2.0 for CARIF, but can also choose to write using different native languages. This is spun as an advantage in the press release, and likely will be an area where these AR browsers continue to compete, providing features and enhancements specific to their own browsers:

“Developers of AR experiences will be able to *choose the AR experience authoring environment they prefer or is best suited to a project* without sacrificing the ‘basic’ experience they can offer their clients' target audiences (OGC Press Release, 2014, emphasis mine).” Continuing to compete over the developer space in their specific authoring environments is an instance where these browsers are not fully interoperable, and are subject to the limitations of what is covered in the ARML 2.0 conversion. The historical analogue to this move is Netscape, who early on hoped to convince online merchants and developers to use their particular server software so that Netscape customers would have the most optimized experience: “Although this cheated, a bit, on the promise of purely ‘open,’ interchangeable standards, users of other browsers might find that the road [...] was a little rough and the stores at the end of the road disappointingly dull (Cusumano & Yoffie, 1998; p. 28).” The difference, however, is that while HTML was the standard across browsers, ARML 2.0 is not how these companies annotate their own content. Under the CARIF agreement, using any one browser would still allow for a more seamless

experience with the content on their servers, while simultaneously excluding things that currently are not covered under the interchange format. The extent to which CARIF is able to keep up with these new types of features also remains to be seen.

The initial choices of what features are interchangeable are also revealing of these stakeholder priorities and predictions. The original vision for why ARML was even necessary, calling a hotel, were enacted through the CARIF agreement because those uses might be the type of user actions that could incentivize corporate developers. The interoperability demonstration held by Layar, Metaio, and Wikitude tested the ‘hotels’ link from other browsers, as well as ‘eating and drinking,’ ‘shopping,’ and ‘tourism.’ These were the kinds of future applications that had justified the need for ARML 2.0 in the first place, and they are the ones that commercial vendors wanted to be available across competing browsers. The features they included in the conversion (text, image, video, audio, call, SMS, URI) also reflect the presumed commercial uses of AR. Other types of AR content, however, such as 3D models or renderings for more complex uses like historical content or artistic expression are still siloed. The need for users to access art across browsers, for example, is seen as a less important consideration when expanding the network size than other types of (commercial) content. At the 10th IARSC Meeting, Martin explained that about 90% of their existing database would be interoperable under the CARIF agreement. ARML 2.0 and CARIF might be a framework to build upon there, but clearly there are differing motivations and priorities for protecting certain data and interests that did not make it into the first draft of CARIF. There are also admittedly technological barriers to 3D file sharing and conversion, but then that omission becomes a byproduct of the decision to do interchange conversion rather than a native standard. This still allows them to compete over

the developer space and their SDK packages, but also moves a step further and promises that developers can ‘write once, run anywhere.’

The other notable omission was the computer vision requests, and the algorithms by which each browser attempts to recognize natural objects. This had originally been part of the goal for ARML 2.0, which we know from one of the use cases written in support of the standard: a common user using an AR a device to recognize a broken coffee machine and downloading the manual. Standardizing computer vision had also appeared prominently in earlier drafts of ARML 2.0 (see figure 40).

Computer Vision-based AR		
Feature	The security features of a 10-dollar-note	
Anchor		A US 10 Dollar-note (along with the location of the security features on the note).
VisualAsset		Some buttons that can be pressed to get more information on a particular security feature
Result		As soon as the 10 Dollar note is detected in the scene, the VisualAssets are projected onto the note in the correct positions.

(Figure 40 – ARML 2.0 Specification Version 1.0.2, Open Geospatial Consortium)

In the CARIF negotiations, however, visual recognition fell out. An informant who was present at those talks explained that it was a major point of contention within the talks. One company who had started off as part of the task force dropped out because the issue:

SEAC02, the browser company from Italy, brought up standardizing computer vision requests. The other companies declined, instead focusing their efforts to standardize on location based requests. They feel like computer vision, the

accuracy, the speed of recognition, and the number of images are critical places where they can still compete for market share on a browser level. SEAC02 was one of the companies who was on the organizing committee for CARIF, but once it became about location only they decided they did not want to be a part of it (personal correspondence).

One reason they might have excluded it from the agreement is that it involved interests beyond just the browser companies, as some AR providers rely on 3rd party companies for their computer vision algorithms. Morton, an executive from a major computer vision company, expressed his opposition to standards at this moment:

Our standard bodies tend to be willing to standardize things for the sake of standardizing, and this is no point. I think in the end we had like three different institutions trying to do standards around Augmented Reality. [...] I think that is too early, because trying to put everything that is possible right now with a markup language would be like trying to have written the HTML formats at the beginning (interview).

Again we hear the charge of it ‘being too early,’ specifically in reference to HTML and disputing the Tim Berners-Lee vision, that AR needs users to identify what they want to do with it. Here’s Morton again: “From my perspective, we don't need to impose anything. [...] If they want to work on it and they are real at some point, I would probably accommodate that and see if there is value for me and for my customers. Until then, I'm not that worried about what is going on.” Morton’s company, it should be noted, is also currently funded by significant venture capital money, which might also contributes to his reluctance.

The CARIF agreement is being heralded as a major step forward in the interoperability of AR amongst the three largest browsers in the world. They launched an interoperability demonstration at the Mobile World Congress 2014 in Barcelona, and it was successful at showing the same content across those three browsers. The agreement was brought together in large part by the people involved, through numerous voluntary meetings and participation in a de facto group like the IARSC. Many of the players believed that Tina was major factor, having

created, organized, and headed up nine standards meetings all around the world, where there was common dialogue and face to face meetings that built a working relationship amongst the different companies and standards organizations. In the end, it was the relationship between the CTOs that was established through the International AR Standards meetings that played a major role in facilitating the agreement. It was not just creating the committee within the OGC, but the years of meetings beforehand – Martin explained at the 10th IARSC Meeting that it was about building trust and confidence in each other that enabled them to move toward a common vision. Sitting in the same room at those meetings also allowed them to convince representatives from other companies that it was not a standard to push the interests of any one company, rather it would benefit everyone:

I remember a discussion I had with [Kirk] a couple of months ago, where we said "whoever creates the standard, it doesn't really matter". In fact, a standard is like a set of key/value pairs, packaged in different structures. It should be no problem for anyone in the AR industry to adapt to a new standard, different from the one it used previously. (September 17, 2011; W3C Listserv).

Martin was also a huge champion of ARML 1.0 and 2.0, but also of standards more generally, traveling to AR standards meetings all over the world in Geneva, Barcelona, New York City, Austin, Basel, Amsterdam, Washington D.C. and Atlanta. That level of dedication and work to put together something he felt was an industry need is what laid the foundation for the agreement. The OGC is also volunteer driven, which makes the standards process often halted, or split between people's actual work. The CTO's of these companies credited Tina with incentivizing them toward pragmatic action: "Actually already more than 4 years ago, we started meeting on a regular basis as part of the AR Standards Community. [Tina] did a great job at pushing [us] to actively pursue that goal (Groten, 2014)." Actors explained that their ability to

work together closely on this interchange format was largely due to the familiarity and trust fostered at these meetings.

That said, there are implications to how the decision was made and the things that were left out. Interoperability in terms of information systems can take many forms: interoperability on a systems level, syntax level, and semantic level. The CARIF exchange approach operationalizes interoperability at a syntax level of different data formats, but it is only partial data interoperability, through a mediated conversion, as perceived by users, and only for certain requests. Having the location standardized and readable across platforms without the computer vision standardized on the user side means that the initial effort of CARIF excludes a major portion of AR. Scanning an object with AR and sending that request through the cloud would not be converted to ARML 2.0, because ARML 2.0 currently does not include language for how to accommodate those requests. This is a significant omission, particularly as so much of the content is moving toward computer vision, with the growing partnerships between browser companies and marketing companies creating AR for packaging and print media (see chapter 5). It seems that these constraints might have been the very reason that computer vision was excluded from the CARIF agreements, because the business models of the browser companies are specifically shifting toward computer vision and partnering with computer vision companies who are still in the early stages of VC funding. What we see here is that the venture capital pressures on multiple sectors of the AR industry also affects the standardization process. Even as the browser companies recognize the need to move past the VC competition stage to fend off larger competitors, adjacent technological areas and partner companies are still in competition and therefore not as amenable to cooperation. So even as ARML 2.0 was created specifically to accommodate computer vision anchors, that element was left out of the conversions in the first

round of the CARIF agreement. SEAC02, one of the browsers that wanted to standardize computer vision, left because it was not on the table. The speed and accuracy of one's computer vision algorithm becomes a new area of contestation. Perhaps an agreement on location was possible simply because that the mobile AR browser companies have not had much success attracting a critical mass of location-based users, and therefore that is a less essential place where they are competing.

The standards groups has effectively pushed any decision to include computer vision into the future – this is explicitly referenced in the draft document for CARIF. Under the tag for `anchor/geometry/gml:point`, which is a way of placing an augmentation relative to another object or anchor more specific than location coordinates, the authors write that it could ultimately support “other anchor-types (like computer vision-based anchors) in the future (AR Browser Interoperability Architecture, 2013).” This leaves open the possibility that it could eventually be part of the standardized conversion, but does not account for the institutional pressures that kept it out in the first place, or the participants who left the discussion because it was not included. It also leaves in place the problem of browser silos, where a particular method of retrieving AR still requires a native application. Tim Berners-Lee's vision was the discursive move that brought these competitors together, but their CARIF solution is still admittedly a ‘basic translation’ that leaves some native AR features out.

Nevertheless, the decision to come together and cooperate on CARIF represents a significant step toward interoperability of AR browsers, and does not foreclose the possibility of taking additional steps. While some might point out the limitations of the CARIF agreement, supporters could simply point out that this is the best enactment of Berners-Lee's vision. The ‘extensibility’ of the standard was something that was important to Martin from the beginning: “I

already laid out [...] that we will keep a close eye on the extensibility of ARML 2.0 to not close doors in the future - I think that's another very important point (September 17, 2011; W3C Listserv).” Whether the stakeholders can agree on the extensions, however, is another story altogether.

Alternate Approaches, Accessing the Tim Berners-Lee Vision

ARML supporters initially drew on Tim Berners-Lee’s vision as a warrant for standardizing AR browsers, but as their standard developed, one of its critics, ironically, was based in the World Wide Web consortium (W3C), an organization founded by Berners-Lee. The Augmented Web Community Group, a working group in W3C, headed by Boris, argues that instead of creating a whole new markup language for AR specific browsers, the industry should develop standards that would allow AR to be created for and viewed on traditional web browsers. Boris argues that serialization, the process of translating data structures into an exchange format to then be reconstructed in another computer environment, is the wrong approach and works on the wrong problems. It does not fully cover what is already possible with AR, and will require constant updating to keep up with the more complex ways that developers are writing for AR. Even if there was a full implementation of ARML 2.0, Boris argues that it ignores much of the work that has been done with existing web browser standards: “While it (ARML 2.0) can be parsed by a web browser, it is not a web based standard (interview).” For an organization that has been working on web standards since its inception, this strikes them as a major omission for building on the original Berners-Lee creation, the World Wide Web.

Instead the approach they are advocating is an “Augmented Web.” In his demonstration of his company’s BuildAR tool, which allows developers to create AR on web browsers, Boris argues: “When you look at the interactivity and 3D capabilities, it's already equal to and in some

cases beyond today's proprietary mobile AR browsers. The ability for it to run in mainstream web browsers (even mobile ones) means it will rapidly be available to a much larger audience than proprietary Mobile AR ever was.” In Boris’s presentation at ISMAR 2013, he points to growing market trends as a reason for pursuing this strategy: “Chrome and Firefox have over 60% of the browser market, and interoperability is key here. That means we can have people writing for 1 billion devices, across different devices and platforms.” These mobile web browsers also come preinstalled on many smartphones, which highlights another barrier to standalone mobile AR browsers, where new users must download a separate application and search for content. This view diverges sharply from how AR browser companies are thinking about users – rather than building tools in an AR environment that allow you to annotate places, users accessing AR through Chrome and Firefox would have all the functionality of the internet, with AR as just another feature of the web, as opposed to a completely separate entity.

The tools for designing these web based applications, BuildAR, also positions itself as the open solution and boasts that it is: “based on open web standards, built on an open source awe.js library, can support a wide range of 3D models, special effects, and spatialised audio.” The video demonstrations of BuildAR highlight that web browsers are beginning to support location-based AR as well as marker-based AR, approximating the experiences that the ‘proprietary Mobile AR’ companies offer. In doing so they can also juxtapose the openness of existing standards against the proprietary format of AR browsers. ARML 2.0 as a standard may be open, but they still have their own systems and are so far only willing to convert data into ARML 2.0. The W3C standards group is calling upon the philosophical, ideological, and practical authority of open source to advance their project. In standards battles for network technologies this is often an area of contestation, as companies will put forward proprietary

standards to secure parts of the market for themselves. This move toward openness is not only meant to attract certain developers to their side, but also to leverage the commitments of existing AR browser companies against them.

ARGON, another browser application begun as a project at Georgia Tech, has rejected the CARIF agreement. Their approach also attempts to build on existing markup languages and web standards, with a goal of making AR accessible to the widest possible group of developers and users through a web-centric platform model. Developers can create AR experiences for ARGON using web technologies, which broadens the number of possible developers. Their criticism of ARML 2.0 is that it standardizes only a limited set of functions that one might want to do with AR (e.g. location descriptions only). Riley was a part of many of the early discussions for ARML 2.0, but even after the process is still skeptical:

I'm still ambivalent about the whole thing. [...] If we had content that we wanted to move between environments that's where markup is appropriate. However in general the things that you want to do with AR don't lend themselves well to that. Do I want to go do the Layar thing or the Wikitude where I have some 2D presentation of content that's tied to a location, that might have a URL associated with it to go off and get more? [...] As soon as you want to do an app, have interactivity, anything interesting those markup languages never seem to work because they never let you do enough to do anything interesting (interview).

Beyond location, a developer might need to annotate other objects on the scene, and in order to make a useful augmentation they would need other information, such as where objects are in relation to one another, what their physical relationship is, how to display on the screen, whether it is occluded by other objects, etc. ARML 2.0 currently only allows content to be attached to a geospatial location, which is fairly limited in application. These would be some of the extensions to the markup language that would need to be negotiated if ARML 2.0 were to become the dominant standard, but for now ARGON is not planning to make conversions for ARML 2.0

ARGON represents another way of standardizing, with the same mass market goals as ARML 2.0. It is less defined in terms of its specification, but advocates argue that it does not inherently privilege location-based queries; with their standard, AR could include 3D models embedded in specific websites or computer vision software, integrated into a web browser. Critics argue that this is still many years away from being possible, and the barrier is outside the control of the community, dependent on which these browsers and mobile platforms will open up their APIs to allow for AR development.

Despite their different standardization approaches, both sides are attempting to invoke the Tim Berners-Lee story as a blueprint for persuading and organizing participation. They deploy the web analogy in different ways and highlight competing parts of the story, however. One side is attempting to create standard tools and markup languages, and using the canonical story as an analogy to the current situation with a new technology. They are using the Berners-Lee story to prescribe a similar action that the community should take, which is to standardize based on existing AR browser needs as a prerequisite to creating a market of AR users. The other interpretation is that the story as a whole is not the analogy, rather it is the explanation for how we got to this point where web browsers reach a mass market of users. For them the analogy is not a warrant to standardize a different set of browsers, rather that the simplest way to build off of that vision would be to include AR as a feature layered on top of existing web standards. Each group tries to access the Berners-Lee story their own way, but the W3C appeal to origin is made easier by their organizational affiliation.

Conclusion

This chapter highlights a series of economic factors, organizational alignments, and individual decisions that brought people together into certain coalitions to work on AR standards. As multiple organizations have made early inroads in the same space, many now believe that linking the infrastructure together is necessary to launch the technology into the mainstream. The question is in part technological, but much of the work is done by actors who persuasively advocate for their approach as well as certain market and organizational pressures. Because of different perspectives and commitments, other groups left the process and began working on different approaches to standardizing AR. What the CARIF agreement and W3C Augmented Web Community Group proposals illustrate is that although two different organizations and groups might share the same market and audience goals, the same persuasive historical trajectory they wish to emulate, the same presumed users and uses, and the same discursive commitment to ‘openness,’ they can nevertheless have dramatically different proposals for accomplishing that goal. This is only one element of the standards story, but studying the emergence of these formats highlights “smaller registers like software, operating standards, and codes, as well as larger registers like infrastructures, international corporate consortia, and whole technical systems (Sterne, 2012; p. 11).

This chapter begins with a standard, created as a result of interlocking priorities between people, organizations, economic factors, and technological visions. ARML went from a company specific specification to a candidate standard (ARML 2.0) in a large SSO, to the basis for an interchange format between the largest AR browsers. It built on existing location standards (KML) with particular uses in mind, which allowed it to align with a particular SSO but also pulled it more toward that organization’s focus on location.

These moves foreshadow a possible standards battle between organizations, which is why this chapter is more about the positioning that goes on before the contest. Whichever approach ultimately gains adoption (and none may ultimately succeed, or be something else entirely) remains to be seen, as well as how each standard continues to evolve. For each approach, however, there have been tradeoffs, as a result of specific decisions that the actors made throughout the process. A series of decisions about which existing standards to build on top of, which organizations to align with, which people can and choose to work on it, and which perspectives to include in the standard itself led to this point – all of which has implications for what a person can do when they pull out their device, how they are able to access content, from where they are able to access that content, and how that content appears to them. These negotiations also guide developer motivations for working on certain platforms and what types of AR experiences they can create. For someone using an AR browser, it might seem like another mundane, taken for granted capability that developed naturally; however the story of these AR standards discussions show that there was and continues to be nothing inevitable about these technologies, and how they move from isolated applications to networks of possibilities.

CHAPTER 7

INDUSTRIAL AR

In reading the last few chapters, it might be easy to get the impression that the AR industry is primarily focused on the consumer market. Indeed, this has been largely true for the past few years, which is why much of my account has focused on the contestation over the design of consumer AR devices, debates over which technologies count as exemplars of AR, the influence of marketers interested in reaching everyday users, and the development of markup standards to build AR content for everyday uses. The burst of consumer AR enthusiasm has extended to the media, who have been talking about AR primarily in terms of the new consumer devices, particularly Google Glass. Technology blogs care most about consumer possibilities their readers might enjoy, company promotional videos show how everyday users could find their way around the city with their device, and even programs like the *Daily Show* and *Saturday Night Live* parody Google Glass wearers. From the first conference that I attended, designers have been incessantly talking about the ‘killer app for AR,’ the one that will make consumers need a particular device and push AR over the top. Now, several years later, as consumers have been slow to appear, some have started to get impatient and wonder if the killer app is coming at all, and whether they should continue to push for that market. In recent years, a growing set of actors have started advocating a turn away from consumer markets and toward industrial applications.

These industrial advocates offer several warrants for this move – not only the small consumer market and their diminishing expectations that one will materialize, but also a fear that the consumer focus unduly heightens negative concerns about AR, and a faith in the industrial market and are the resources those sectors might provide. To them, consumer AR is viewed as a

solution looking for a problem. One lament is that many of the existing applications have the same fundamental problem: “they all assume that people want this firehose of AR information and visualization all the time (Riley, interview).” Certainly the major conference themes like “always on, always augmented” are positing that kind of ubiquity as the goal for AR. Industrial advocates argue that the expectation that anyone wants to engage with AR in this way, however, is simply a questionable assumption on the part of consumer AR advocates. It may ultimately be that most users simply do not wish to engage with the world and other people using AR in this way.

Another problem, industrial AR advocates argue, is that these uses do not take into account existing technologies that can already accomplish the purported tasks without AR, in ways that are arguably more efficient. If people want to navigate, they can pull out Google Maps or use services like Yelp. If they want more information about something, rather than scan it with AR they enter a word or two into a search engine. As Dennis explains, a dominant assumption about public consumers is that: “the public is finicky. They need it to work and to work quickly, and it needs to be better than other options. Otherwise there’s nothing compelling them to use it and they’ll never come back to it (interview).” This portrayal of consumers in AR is part of a growing belief that consumers are changing in the new media environment, and that advertisers and designers now have to think about and fret about consumers’ proclivities to be more active in their demands and quick to reject technologies (Turow, 2011). At the same time, the consumer focus has also given rise to concerns about the likelihood of legal and ethical controversies around facial recognition, privacy (Roesner, Kohno, & Molnar, 2014), and subversive recording (Denning, Dehlawi, & Kohno, 2014), which these actors are hoping to avoid.

Subsequently, the contestation is over whether these problems are even the ones that the AR community should focus on. Industrial advocates argue that in order to affect a change in consumer behavior that normalizes AR use, there needs to be an certainty that when users take out their phone to augment something, the content they're seeking would be there. This requires that the CoP work toward building the infrastructure for consumer AR, with an aligned developer base and consensual standards (see Chapter 6). This is no small task, and may be mutually exclusive in terms of the human and financial resources it will require; they are still far from reaching the level of ubiquity some believe is necessary. In places where there is already a high concentration of AR content, filtering it becomes the major problem: "Say I'm in Times Square, there could be literally thousands of strands of content attached to that location. The challenge then becomes how to filter that information in a way that's useful (Riley, interview)." In other areas that have fewer users or points of interest, the opposite problem can occur where there's little or no content at all, which would slow the adoption of AR. In all of this the payoff is uncertain, as it remains to be seen whether users even want to use AR in this way, with the most recent survey suggesting that most people do not engage with AR that frequently, for a variety of technological, practical, and social reasons (Olsson & Salo, 2011).

For those supporting industrial AR, these are critical reasons to pull resources away from the consumer problem space and target AR much more narrowly toward industrial and commercial applications. The term AR even originated in industry, from Boeing engineers (Caudell & Mizell, 1992). Many industrial supporters are academics working in research laboratories, who have long had connections with the industry space. At ISMAR 2013 they even hosted an industry day preconference, which was devoted to discussing the things industry might need from the AR academic space, and vice versa. A number of AR companies have started

customizing solutions for businesses, including the hardware, software, and database work necessary in setting up a client-specific AR system. In recent years, representatives from these industries have been attending these meetings, to further explore the possibilities of AR.

This chapter examines the growing coalition of AR advocates and developers who are not interested in everyday consumers at all, but are instead focused on industrial applications of AR or ‘business to business’ (B2B) partnerships. Because consumer use of technology garners much of the public and academic attention, industrial uses are understudied. But these industrial partnerships are growing, and represent another important stakeholder group, one that is not interested in AR as a two-sided market, rather as a specific tool to streamline their operations. This turn emphasizes different capabilities of AR, designed to address very specific problems in particular settings. And in doing so, these partnerships have started shifting the priorities of the entire community toward those industries and shaping the design of the AR technologies away from consumer gadgetry.

While this industrial focus may avoid certain concerns wrestled with above, it brings with it a different set of actors, its own technical barriers, and problems specific to the industrial space. All of this demonstrates how malleable AR technology is, that so much of what the technology looks like is ultimately shaped by the people surrounding the technology, the social groups they enlist, and the specific problems those groups emphasize. The contrast between consumer and industrial AR shows how this is not just as a local contestation over a specific device or a specific application, rather it is a fundamental disagreement in the community over the future uses to build for and the industries to prioritize, the infrastructure necessary to support those aims, and whether these trade off with one other and can be pursued simultaneously or sequentially.

AR for Maintenance, Repair, Assembly, and Warehouse Applications

A common industrial use case for AR is one that leverages the visual displays to show multiple steps of a work task to an employee. AR lends itself well to particular sequential tasks, like maintenance and repair. One video presentation at AWENY 2014 depicted a scenario where a repair person used AR to fix an air conditioner unit, and they were able to use an AR application to recognize the device and see the steps they needed to fix it. “Instructions will be easier to understand if they are available, not as manuals with text and pictures, but as computer-generated information superimposed upon the relevant equipment, and showing the relevant tasks that need to be done step-by-step and how to do them (Yuan, Ong, & Nee, 2008; p. 2720).” This vision of AR giving instructions is also quite common in the academic space, with one presentation at ISMAR 2013 showing how AR could help show astronauts how to open a hatch door. There have been some early pilot programs and experimental pieces that actually attempt to test these applications. Henderson and Feiner (2009) tested the effectiveness of an AR system for a user repairing and maintaining an armored personnel carrier turret. Regenbrecht and colleagues (2003) explained how they were testing systems for servicing and maintaining larger systems, like changing a filter on a space station, maintaining an engine, and diagnosing a tram system. Much of the research in industrial AR is aimed at finding the most efficient displays to do this, the optimal user interfaces to use, and how to precisely track the real world objects (Yuan, Ong, & Nee, 2008).

The empirical question that researchers often ask around AR in these environments is whether these systems will increase operator efficiency and effectiveness on a single task. This is how these systems are evaluated in experiments, by operationalizing efficiency and effectiveness

in terms of a persons' movement (e.g. range of motion, head rotation, view direction), task completion time, time to locate certain objects, and finally ease of use and intuitiveness (Henderson & Feiner, 2009; 2011). These studies tested AR against other systems for delivering instructions (e.g. compared to a screen or a traditional printed manual), and found that the AR condition was more effective than other displays in terms of reducing head movement and completion time (see figure 41). Another study found that an AR system displaying a light over a particular sequence of buttons could increase task efficiency and effectiveness, reaction time, and movement (Marner, Irlitti, & Thomas, 2013).



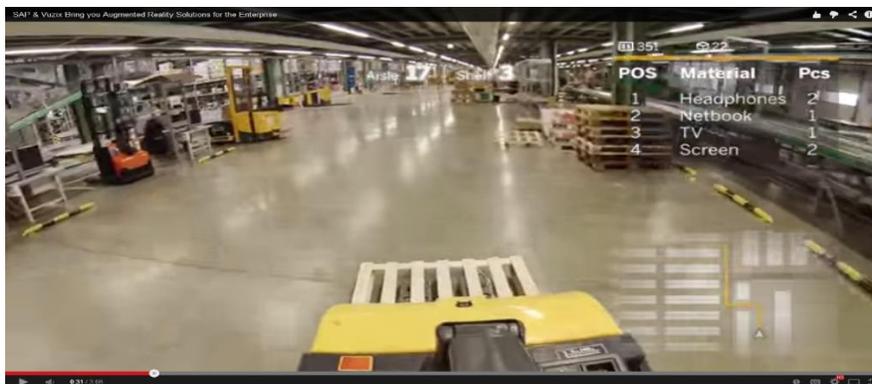
(Figure 41 – Maintenance Application, Henderson & Feiner, 2009)

These studies demonstrate how industrial visions shape the research agenda of certain members, and how their work subsequently supports the argument for industrial uses. Of course, these types of use cases lend themselves quite well to AR systems, because the object they are maintaining provides a consistent and stable reference object for the AR technologies to visually recognize. While all of these examples are based on possible practical applications, the more difficult use cases like repairing a jet engine are often featured, to make a persuasive case for AR

(and because the research was funded by the military), that AR can be useful even in tasks that have an extremely high degree of difficulty.

A similar variation to the maintenance tasks are the warehousing use cases. In particular there is significant interest in utilizing AR with warehousing, which is an important part of the supply chain that requires a significant amount of physical labor. Errors in this area can cause problems down the supply chain, all the way to the consumer. Warehousing has also become more complicated, as there are now so many items that need to be shipped, stocked, and moved around, all in a way that properly utilizes the space and the people in it.

A video produced by S&P and Vuzix was then shown, depicting a (hypothetical) forklift operator wearing a pair of AR glasses while navigating the warehouse. The AR glasses would display the objects that needed to be picked up, and then display information directing the operator spatially to where those objects needed to go (see Figure 42).



(Figure 42 – Warehouse Application, Vuzix and SAP Video)

These AR applications demonstrate the potential benefits of AR for organizing spatial information, while also doing so in a space that is already highly technologically organized and well defined. Where the maintenance and repair tasks were focused on a specific object, this application demonstrates how AR can annotate large spaces and show spatial relationships between objects in space, while offering the same benefits of technological efficiency and

effectiveness. The warehouse video imagines that the AR operator could ‘confirm’ that they have completed a certain task, which would update company records. This suggests one way AR could be integrated with their existing systems infrastructure, in a way that they hope will eliminate redundant work.

The video is an attempt to demonstrate how AR might reduce human error in picking and placing products and reduce accidents, but the unstated implication is that it could help managers take inventory and see what is happening in these warehouses, all of which has the effect of reducing costs. The idea that human error is a critical problem in these environments is common, particularly as the speaker described something called the ‘tetris problem’ in warehousing, where one might have to stack large objects on top of one another to maximize the space in a particular shipment. This is easy to do if all the objects are the same size and weight, but if they are not, these companies often worry about how warehouse employees will work through the tetris problem. Stacking certain items of varying sizes and weights could damage items, or be stacked in a way that makes inefficient use of the space. In this scenario, AR is being prescribed as a technology that can visually show spatial relationships, transmit remote expert knowledge, and delegate certain tasks the device as opposed to the local worker.

Beyond repairing or moving objects, AR can also help employees assemble products. In many industries, a correct and efficient assembly process is essential for a successful system of production (Yuan, Ong, & Nee, 2008; Regenbrecht et al., 2003). Industrial AR advocates argue that AR could solve two important problems in the assembly process – learning of the assembly task and context shifting while doing the task. These are often related, as in the learning process people might look away from the task to look at manuals or other objects in their environment. The concern is that: “These changes in attention will consume valuable time, [...] may lead to

reduced productivity, increased assembly times and errors, repetitive motions and strain injuries (Yuan, Ong, & Nee, 2008; p. 2718).” Because AR makes it possible to show things directly in the field of view, it could theoretically eliminate the need for context switching during assembly. Experimental studies testing these claims have found that AR can help increase productivity and reduce errors (Renhart & Patron, 2003), with 3D overlays of instructions reducing errors by 82% (Tang, Owen, Biocca, & Mou, 2003).

The same technology that might help a worker navigate the warehouse can also help a surveyor navigate a landscape. At ISMAR 2012, I met an industry representative from one of the leading earth moving companies, who explained that they were interested in equipping their bulldozer drivers with AR eyewear, to help them ‘see through’ the ground to where pipes and obstacles might be, to inform where they should and should not dig. At the 10th IARSC meeting, one industry representative from the Electric Power Research Institute presented about their vision for implementing AR to see where underground transmission lines are, as well as help survey cities after natural disasters to see where buildings should have been (see figure 43).

Not only are these industrial visions fundamentally different applications for AR, they have much different requirements for the technology. Where marketing, advertising, and entertainment applications directed toward consumers emphasize the recognition of 2-D printed images and the registration necessary to display content over them, these maintenance use cases require AR to recognize complex 3-D objects like jet engines or air conditioners, and from a variety of user perspectives. This adds another technical dimension to industrial AR, as the changes in perspective requires dealing with issues like depth, shadows, luminance, and changes in size. One temporary solution demonstrated at AWENY 2014 was to recognize a 2-D image of the machine, which would then overlay a 3-D AR model built beforehand using computer

assisted design (CAD) modeling. This would be something that could ‘stick,’ but it was limited in the angles one could recognize the object from. The first versions also were only able to display one step, and it could not recognize when that step was complete or display subsequent steps. The larger goal is to have a system that can recognize a complex 3-D object and display 3-D objects that account for changes in viewing angle and the movement of the device. The maintenance and assembly uses also require the ability to recognize when a certain step was complete, either through gesture recognition or continuous recognition of the object, as well as to change AR displays in sequential order from one step to another (see figure 41). The outdoor industrial uses impose a whole different set of requirements altogether, particularly the need to identify location coordinates, know where the user is looking, and display visualizations over those with a high level of accuracy (see figure 43).



(Figure 43 – AR Power Lines, Electrical Power Research Institute, www.perey.com)

The nature of many of the industrial tasks (repair, maintenance, assembly, etc.) also call for some level of input or gesture recognition, as well as the ability to sequence between multiple AR experiences as opposed to singular experiences. The context is also different, as the consumer uses could anywhere, some of these industrial visions have more clearly defined areas where certain conditions can be held static.

These contextual demands demonstrate how consumer and industrial AR pull the technology in different directions, given their distinct priorities. These different visions and demands on the technology demonstrate how relevant stakeholder groups can reorient the goals of a particular technology and push it toward those problem spaces (Bijker, 1995; Pinch & Bijker, 1987). In this instance, it completely reorients the goals from widespread distribution and use to a much narrower set of clients, but in doing so changes the technological requirements, priorities, and contexts they need to accommodate.

The Move to Industry in the AR CoP

While industrial uses for AR have been considered for a long time, the academic focus has primarily used industrial contexts as justification to explore technological problems. In my tours of the academic lab sites, the students working in the lab explained that they were building prototypes for industrial AR, with an eye toward experimentally testing certain functions. Some of these were full on industrial equipment that was actually used in manufacturing plants. The move toward actual implementation of these systems, however, was largely absent. At ISMAR 2013 one of the members of the steering committee lamented: “computer science academics and researchers often express disdain for simple, known technologies, instead they need to be on the edge. They push the edges, but do not refine existing technologies.” The training and incentive

structure for students, then, is to research new technologies and algorithms, rather than improving or implementing existing technology. This gap is well noted in the academic community, as Georgel (2011) did a survey of industrial AR applications across the work cycle of industry (design, commissioning, decommissioning, manufacturing, maintenance/inspection), and found that while many applications were partnered with industry in prototype phases, only two applications were out in the field. At ISMAR 2013 they held an industry day designed to facilitate an exchange between academia and industry, an instance where these intermediary groups could come together to discuss common goals and partnerships. But the structural incentives behind ISMAR and academia strongly incentivize testing in these contexts, not implementation. Here's Bruce: "Perhaps those use cases are too boring for academics. You won't get an ISMAR paper for building the technology, but we need it for products (ISMAR industry day 2013)." This comment speaks to what is and is not valued in the ISMAR forum, and how broader incentives can direct certain work toward particular ends. At ISMAR, the industrial/occupational contexts are used as the parameters of the inquiry, but the real contribution is to engineering in terms of tracking, calibration, registration, display, and input.

While there has been some degree of overlap between academic laboratories and specific industries (typically in the design stage), the recent move toward deeper partnerships and implementations is being initiated by AR companies who are hedging on the consumer market and looking for new resources. This shift represents a change in how actors are promising the technology, the different actors they have to persuade, and the different contexts they have to account for in design and planning. Rather than persuading people that AR might do something cool or fun, their arguments are strictly about operational efficiency in a specific context. The information exchange that happens at these conferences also facilitates these moves, as industry

actors are explicitly drawing from the authority and work of these academic studies to persuade corporate actors to adopt AR. Here's Walter, a CEO of an eyewear company, explaining:

At the end of the day businesses are in business to make money. It's gotta be a value proposition. There has to be something that really makes something more efficient, more productive whether its currency, money, all those typical value propositions. [...] There was this study [...] that showed upward of 30% efficiency of time on task when you have information in front of your eyes (interview).

Walter goes on to explain how he utilizes this study to quantify the potential impacts of AR when meeting with clients:

We go to them and say: we can save their worker 20 minutes, or 30 minutes a day. [...] we know that person works two-thousand hours in a year, or two-hundred fifty days a year, and we know what we pay that employee. We know that we're saving them 30 minutes a day, and we also know that there are five-thousand, ten-thousand, twenty-thousand workers that do that exact same job function at that company (interview).

These experimental pieces demonstrating AR's efficiency becomes an important way that actors persuade companies and make these appeals to the industrial stakeholders that might implement these systems. The other way that these conference spaces become important is through the exchange of pilot programs and success stories about them. Here's Steve, who works for a ship construction company: "we need to see at least a few implementations that are successful so we can point to them, and to demonstrate that we need to do this to keep up. We need more people to work in industrial AR, and then after a few successful implementations, I can go to my boss and make a recommendation." Knowing how these other industries are engaging with AR becomes an important resource to draw on and provides the impetus for these AR champions to make the case in their respective organizations/industries. Other companies like Metaio have begun to answer these needs, and devoting resources to solving the industrial display problems –

at InsideAR 2013 they were showing how they could have a moving 3D AR object slide into a piece of machinery.

Industry partnerships also provide important motivations to spur major changes in the AR space. At AWENY 2014, Riley was explaining that one of the most successful AR applications he ever worked on was in partnership with a poultry processing plant. At these plants they had to process meat by cooking them to a certain temperature to prevent spoilage before they shipped. The AR system was designed to show colors of the chicken that would help indicate when it was precooked to the proper shipping temperature. This was a matter of millions of dollars, because if it was undercooked there was the risk of health issues, recalls, and possible litigation. If it was overcooked even a little bit it would shrink the chicken (which they sold by net weight) as well as make the chicken more dry, less tender, and generally of lower quality. This case demonstrated the ‘AR sweet spot,’ leveraging the unique technological capabilities of AR and applying to the very specific needs within an industry. The broader point that the professor was making with this story, however, was that when AR finds the particular niche application that solves a problem for an industry, companies willing to spend millions of dollars on a particular system because the net savings for them are so great. The industrial interest and willingness to invest in these complex systems is another warrant for AR to move toward that space. Walter explains that the shift is happening slowly, but there are tangible signs that companies are beginning to write AR into their budgets like they would other technologies: “They considered AR part of the mix. In other words, just as they have made considerations on software and buying laptops (interview).”

Existing industrial client relationships are starting to provide a source of revenue for AR companies, but they change what they require of AR and the types of work these AR companies

have to do. For example, one of Walter's clients is a major worldwide retail chain, who is thinking about having their customer service employees wear AR systems while in the store. Currently, within each store there are certain sales floor reps designated to particular areas of the store (e.g. electronics, auto, etc.), assigned based on the expertise of each employee. One of the factors motivating their interest in AR is to save on training employees into these floor specializations:

They're trying to lower the cost of training. You have a low-skilled person that can be productive on day one, because essentially, we're providing that information. [...] They've got 140,000 units of product, how is one person, a sales associate on the floor able to be an expert on all those items? [...] We can feed that, and pipe that, through the use of just looking at a bar code right into that sales associate's field of view. If a customer comes up and asks them about a particular product, they can list off benefits and values very easily (interview).

Several people from the company were going down to the retailer headquarters to help set up the system, where their task was to take the existing product database systems and convert that content to an AR platform. Here, the problem is not attracting content developers, but having the expertise to turn existing infrastructure and systems into ones that can accommodate and support AR devices. Their task is out in the field, serving as consultants for this company, trying to figure out how to integrate a new system into their day to day business practices.

Walter and Riley are part of the coalition trying to persuade people in the AR community to move away from consumer AR, arguing that they stand to benefit greatly from supporting and participating in these industrial partnerships. Part of their argument is that pushing these complex solutions to industry makes the experts of these system highly valued commodities, and would in effect establish their group as the expert class of AR practitioners. Walter was explaining that their company has more clients than they can take on, simply because each type of solution needs to be so customized to each industry. The successful partnerships with these types of

production, manufacturing, and service industries offer another avenue for bringing more investment into the community (without waiting for adoption, certain devices to come out, or doing AR marketing campaigns), which could pull more members of the CoP in that direction. At AWENY 2014, I spoke to the CEO of an AR company making this move, and he observed: “a few years back all anyone was talking about was consumer AR. Now where is everyone? That’s why we’ve been focused on enterprise applications.”

This shift redefines the problem space for AR and brackets off certain design and technology concerns about AR devices. First, it does not have to attract voluntary users, promising them that a certain device and application is worth using. The promises are instead directed at a higher level of management. Second, they do not need to leverage or attract any third party developers, because the content they need to use is already owned by the industries. Lastly, it brackets the aesthetics question for HWD, because industrial users are not voluntarily using these devices and do not need to ‘look cool.’

This last one is an important point, because many of the earliest HWD devices have significant physical and aesthetic limitations that many argued would deter consumers from volunteering to wear them. Industrial clients are less concerned about these issues of aesthetics and comfort if there is a larger organizational objective that AR can help with. At InsideAR 2013, these motivations are made explicit, that AR should target industrial uses in order to minimize this persistently tricky problem for AR development: “employers can just make their employees wear this stuff. I’m sorry if that sounds draconian, but that’s the way it is (Kim, InsideAR 2013).” It is at these conferences, amidst other members of the community, where AR advocates can freely disclose their intentions for making certain moves. This also means that the uses envisioned are specifically targeted toward solving problems with lower level employees

who are engaged in physical labor (e.g. warehouses stockers, repair workers, factory and assembly line workers), who could be subject to wearing these devices by management. Redefining the problem space to industrial settings changes how advocates promise AR and the work they have to do to make these systems viable, but it also affects how hardware companies are designing devices that service these needs.

Industrial AR Hardware

Founded in 1997, Vuzix started as a company that designed AR eyewear for personal entertainment solutions. In recent years, unlike other companies, they have specifically targeted enterprise and industrial AR uses (see chapter 3). With the M100, their latest product, they explicitly state: “the M100 is designed primarily for enterprise, commercial, and medical applications (see figure 44).”



(Figure 44 – Vuzix M100 Advertisement, www.vuzix.com)

Through this device, we can begin to see how certain assumptions about the users and the industrial tasks they will need to perform get embedded into the design of the M100.

Unlike mobile AR devices, Vuzix’s enterprise applications require a headworn form, because users need their hands to perform the work. And unlike Google Glas, Vuzix placed the screen directly in the field of view of the user, as it is intended to be central to the users’

particular task and cannot be ‘out of the way.’ This can be understood as a move to limit context shifting, which is an important part of maintenance, repair, and assembly uses. Only by placing the display directly in the line of sight can an employee continuously focus on the task without looking away. And, anticipating some of the field and warehouse contexts where this device could be used, they stress that wireless internet is not necessary to run these applications, preferring to store the relevant content on the device itself. Here is an instance where a problem is turned into an advantage in a different context – data transfer is a problem for many HWD that do not have access to mobile networks; here they highlight industrial applications where data transfer would not be necessary.

The viewing screen, unlike Glass’s exposed screen, comes in an enclosed box. One of the executives at Vuzix explained in a presentation that this was a conscious decision, given what they presumed about the user (worker) environment: “If you see here in this mining example, where people are working on a hillside, they can use the M100 to show where you need to dig. This is one advantage of having the enclosed eyebox [...] if you’re out in the ambient lighting, in an outdoor environment, the augmentations can get washed out. With a darkened eyebox this doesn’t happen (AWE NY 2014).” This is specific to the earth moving industry, and an instance where the presumed context necessitates a particular design. Vuzix also explicitly points to the ‘ruggedness’ of their device, as well as durability being a tested and required for future devices: “We had a client who was trying out a small wearable watch. And what they told us was that at least once a week it would get dropped and run over by something. Our devices have to be able to stand up to that environment, which we demonstrate by shipping to the Navy. They have to be thrown down on the ground and stomped on to show that this technology can benefit an enterprise (AWENY 2014).” The enclosed eyebox is part of the ruggedness claim, where having

the lens exposed makes it much more fragile. And Vuzix also offers the option of an ‘industrial packout,’ additional equipment specifically intended for industrial uses. This includes an auxiliary battery pack that is intended to solve the ‘device battery problem’ that continues to plague the headworn space, and are necessary for the specific types of uses they are envisioning, where workers may be out in the field without the ability to charge their device for hours.

The purpose of the device becomes clear through the design as well as the places where they anticipate their use. The Vuzix M100 demonstrates the social construction of these artifacts, as certain design decisions were made specifically to address an industrial stakeholder group and solve problems that may arise in those settings. These powerful groups have coalesced around an alternative set of AR visions, such that these larger promises (e.g. limiting context shifting and improving efficiency) and presumed requirements (e.g. outdoor lighting, durability, and battery) drive the design of the artifact. These are instances where certain priorities in the material design are emphasized over others, one being that Vuzix is more concerned about functionality than trying to make decisions based on style or aesthetics. The other decisions about how one can engage with it and the things it accounts for draw a ‘script’ that attempts to anticipate the surrounding context in which the technology enters (Akrich, 1992). The negotiation of that script, however, is a complex one that remains to be seen.

Honda, AR, and Practical Issues with Industrial Implementation

While certain problems are less pressing with industrial HWD (aesthetics, battery, multiple components), other potential issues are elevated, ones that affect how these devices get implemented in these contexts, how people attribute meaning to these devices, and the possible rejection of the technologies by employees in these settings. Thought industrial AR advocates do

argue that workers can be ordered to wear these devices, in fact any AR implementation enters into a complex relationship between management and employees, between workplace norms and the experienced meaning of these devices. The naïve versions of these persuasive visions portray the technology functioning perfectly and fitting perfectly into the work environment. But many studies have shown that even the best technological design and planning cannot foresee how users ultimately adopt the technology (Oudshoorn & Pinch, 2003; Pinch & Kline, 2004), particularly in workplace settings (Suchman, 2006; Zuboff, 1988). Some of the design components, no matter how well intentioned, may not be appropriate in certain settings and power configurations. Beyond that, certain social meanings and perceptions get embedded into these devices that can affect the regularization of these technologies (Pfaffenberger, 1992).

At the 9th IARSC meeting, one of the headline events was a showcase of 20 years of AR eyewear, with over 30 pieces that had been donated from various companies and collectors. It was held at an art gallery in the Chelsea district of Manhattan, and it was a night of celebrating and reminiscing about these novelty pieces. Two exhibitors were leaders in the field, who had been working on eyewear for decades. They went over each one piece by piece, telling stories about it, covering its history, and the technological capabilities of each. However, when they got to one particular unit they told a much different story, about a failure not necessarily because of technological reasons but in its implementation with an industrial client.

In 2003, Honda committed to buying 3,800 units of the Microvision Nomad, which were intended for use by their service technicians. The system came with a wearable display that would provide auto technicians with access to repair information and online vehicle history (see figure 45).



(Figure 45 – Microvision Nomad, www.cnet.com)

It was the classic maintenance and repair case, and it seemed like the idea had a reasonable chance of success. It had the commitment of the industrial client, who made a huge investment in the device. It served a clear purpose, for service technicians to see the vehicle history and access the repair manual so they could provide “faster, more accurate service and repairs with critical information superimposed in their vision as they work (Microvision Press Release, 2004).” They essentially had 100% market adoption, with distribution to all of the dealerships and an exclusive agreement with a major manufacturer. There was also organizational commitment to the device, considering Honda had made a huge investment in the technology and shipped a unit to every one of their dealerships.

Despite these advantages, the Honda implementation AR was a failure, with many reasons why it did not work the way it was intended, ranging from technological to social. One issue was the design of the headset itself, which protruded from the head, making it difficult for mechanics to get around and under the car. The display itself was also an issue, as the edges of the screen limited their peripheral vision. Additionally, the system required users to wear a video

control module on their waist, a place where mechanics sometimes carry their tools with a tool belt. Putting on the device was also a multi-step process.

Just as important as the technological issues, however, were the social issues associated with wearing these devices. Fred explained at the eyewear exhibit: “It is the matter of pride amongst mechanics that they know instinctively or tacitly know what they’re doing without consulting a manual. Wearing one of these devices would be akin to saying that you would need this technological aid to tell you what to do, admitting to all of your colleagues that you lack skill and expertise.” As a result very few Honda technicians used the system, and Fred explained that “you could probably go to all of the Honda dealerships today and ask them about it and they’d still have the Nomad in a backroom somewhere.”

The design of technological systems is embedded with certain assumptions about the user, the task the system attempts to solve, and the context in which it will be used. Designers of these technologies have to make certain hypotheses about the world into which the object is inserted, and the technology is constituted by the back and forth between the designer’s projected user and real user (Akrich, 1992). In industrial implementations of AR technology, workers are expected to conform to the social visions of the technologies creators. However, what the Honda example illustrates is adoption cannot always be imposed, and the implementation of these technologies may complicate many political, organizational, and social commitments in these local setting. The execution and design of the system failed in the local context because the users overtly and covertly challenged the meaning of the artifact. The design had legitimate limitations, but these were not necessarily prohibitive. The larger issue, however, was that for these service technicians and mechanics to adopt the artifact would be an admission that there was something wrong with the old way of doing things, which would have been a challenge to

their performance of expertise and authority derived from their experience. The system also challenged their notions of tacit knowledge and how one might go about fixing a particular problem using their own experience and instincts.

This demonstrates how technologies are enacted within the structures that they enter, and in this instance the technology ran counter to the norms of the group and their recurrent social practices, as well as an active threat to their sense of self in the work context (Orlikowski, 2000). Because of this, no one was willing to step up into the role that the designers envisaged. These motivations resulted in what Orlikowski (2000) labels as inertia, where users have very little interest in using the technology that threatens the status quo. Future implementations that continue to rely on experimental task and movement efficiency numbers may miss the power dynamics in the space. This groupwide disavowal, or the setting of social norms that associated the use of technology with a particular status, is an example of how users can define their own roles that confronted the technology, and rendered the technological object a chimera (Akrich, 1992). On the one hand, people are pushing these visions, making these persuasive, trying to solve specific technological problems that the space calls for, and designing artifacts that account for the presumed user and industrial settings. But they are also negotiating over the meaning of the device that ultimately occurs in the users' reactions to the technology and the ways the technology specifies certain relationships in the users' environment.

Conclusion

The industrial market has been a part of the AR promise since the beginning, but it got pushed aside in recent years as consumer AR began to seem possible. As that development has slowed, there has been a turn back to industrial, as some stakeholders attempt to shift the focus of recent

hardware and software development in AR back on industrial problems and industrial implementations. Across numerous settings and industries, they are making very specific promises for how AR might benefit them organizationally and financially. These conferences are a place where knowledge about how to promise those is taught and examples that might drive the industrial promise are exchanged, as well as a place to persuade more members of the CoP to move toward industrial AR. Consumer and industrial AR are not strictly mutually exclusive, but they do differ in priorities across all areas of development, including technological capabilities, physical design, standards, and content. These actors are attempting to move up the technology spectrum to serve a different industry with more rigid technological demands given the content and context. On the technological side, industrial AR pushes for more complex recognition capabilities, 3-D models, highly accurate registration with the physical world, certain types of interaction capabilities, and sequential displays of AR. Consumer AR advocates focus on solving problems within the confines of existing technological devices, such as creating better recognition of 2-D images and simple 3-D shapes, expanding the storage capacity, and creating AR visualization/animation options that are more colorful and playful. The physical design of the Vuzix device demonstrates the various ways they are attempting to cater to these presumed industrial uses, with a centered screen and enclosed lens, in contrast to the stylish sunglass prototypes seen in Chapter 3. Industrial AR also rely on a different set of standards (e.g. CAD models) and may necessitate more rigorous location standards. Lastly, the work that people have to do to serve both industries are quite different – where the consumer space is focused on attracting developers and creating content to populate the space, the industrial work requires converting databases and systems to allow client-specific content to be accessed through AR.

Different actors have different strategic reasons for making this turn, from the economic investment that industry could bring to genuine concerns/pessimism about the consumer space, but the way they are promising these have important implications for the technology and the social meaning surrounding it. Because the visions they are building for presume a lower level employee who can be ordered to wear these devices, they can bracket the aesthetic and comfort problems with HWD. This also means they have to convince the upper level managers of these organizations because they can make people wear them, which means the benefits must be geared toward this group. This is an instance where AR actors are attempting to redefine the problem space by appealing to power, an instance in which the strength of the relevant social group has direct implications for the technology itself. Taking the SCOT approach to see how alternate configurations might have formed, this strategy could be contrasted with an alternate model. The industrial visions could be held constant, but instead AR advocates could have attempted to make the system so effective and so comfortable that industrial workers actually wanted to wear these devices, and marketed it directly to workers accordingly. That they are not doing this speaks to their anxiety about the AR design space, their need for resources in the current moment as opposed to far off in the future, the people who have control over the budget decisions in the workplace, and the highly specific work problems they are trying to address through AR, some of which requires access to data as a prerequisite to making the system work. The process that leads to organizational adoption of AR is an important one to map because “no one feels the prescription of new information and communication technologies (ICTs) more than those who are forced to work with them every day and whose livelihoods depend on them (Slack & Wise, 2006; p. 153).” As the Honda example demonstrates, imposing AR in the industrial

space as an exercise of power may bracket certain design issues, but elevates other problems specific to the organization.

How this relationship unfolds in the AR CoP is still yet to be seen. It may be that more actors move toward the industrial space, or actors continue on with the consumer space. It seems likely that there will continue to be development on both, with the question being whether one becomes dominant first and leads to the other, one becomes dominant and crowds out the other, both continue to develop concurrently, or neither gains traction. Already there has been some ebb and flow within the AR CoP between these two visions, as actors are competing to attract the most material and human resources to work toward one or the other. Each group has their own defined set of problems, which may ultimately converge or lead to solutions that are intractable and require further standardization (i.e. location requirements, 3-D object augmentation, etc.). As the relationship between the industrial and consumer space deepens and continues to develop, this chapter points to places where the industrial space has already started shaping the material features of the technology, its potential capabilities, and potential meanings of the artifact.

CHAPTER 8

CONCLUSION

Like any emerging technologies, AR now exists in a state of flux, a mixture of ideas and hardware, projection and speculation, continuity and discontinuity. The people working on these technologies are an eclectic group of professionals and amateurs, old-timers and newcomers, makers and sellers, business-minded executives and casual enthusiasts. Together they have been working to bring this thing to life, in cooperation and competition with one another, individually and collectively, meeting a few times every year to update, celebrate, and prognosticate. When I attended my first meeting, I was in awe of the promise and the possibilities to come. After the first major conference I attended, I heard it from so many people that I was convinced that in the upcoming year, AR would become a major technology in our everyday lives, with new devices and new applications. That this dissertation is about the power of futures is circular, then, because it is owing to those first persuasive futures that drew me to study AR.

It was only in going back to the space over and over again did I realize that these predictions were overstated and even mistaken, occasionally deliberately so, but that their accuracy was relatively unimportant. At the next conference there was simply a revival of the same promise, sometimes even delivered by the same people. I began to notice the similarities in these conferences and meetings, almost like watching the same play over and over again. The subtle changes, then, were not in the presentations of the futures or the goals they wanted everyone to work toward, rather they were in people's explanations for why that promise had not come to pass, what new thing was going to bring it about, and how actors maneuvered around those developments and disappointments.

The other noticeable changes came not in the leaders, who were generally present consistently throughout, but rather in the rotating group of supporting players who were being socialized into the community, learning to promise the technology in the same way, and taking on new roles and tasks in the community. One way of understanding these conferences, then, comes through in the leaders' presentation of the technology, showing others how to promise the technology, directing them toward certain tasks that might realize those goals, and aligning the group together broadly in support of AR against naysayers and skeptics. While this certainly happens and is the goal like any other trade industry, the group was not unified or cohesive about certain critical issues. What I observed was another layer of complexity, with longstanding and newly forming tensions that appeared among different factions.

Each of these chapters tells a different part of the story: the dynamics involved, the stakeholders involved, the resources they draw on, their motivations for taking certain actions, and the tactics they engage in to do so. Chapter 3 highlights the contestation over the physical design of the technology, not just as small adjustments over features, but one in which a longstanding vision (HWD) runs up against a tantalizingly new and tangible alternative (mobile). This represents a moment where the community is faced with a problem, either to diverge from the presumed historical trajectory and embrace this form, which comes with its own set of actors and promises, or to stay the course and dismiss this as a form that will fail to deliver on the more glamorous promise of AR. The ongoing clash over the technological, market, and normative futures of each form reveals not only the stakes of the struggle and the concerns each coalition has about the other, but also helps to explain how these social issues end up embedded into the design of the technology. Mobile advocates are trying to make their devices more robust, with more cameras and more sensors, in an effort to address critiques of technological limitation and

limited application. Headworn designers are focusing on the wearability and aesthetics of their devices, in an effort to prepare and create a market for these devices. Some are integrating their device with the mobile form to coopt technological/market advantages, others building in social cues to the device to allay the fear of non-users, and still others banning certain applications through policy to limit criticisms of certain undesirable futures.

As stakeholders attempt to stabilize the technology, these mappings help to understand the groups that are forming on both sides, the issues that each is raising, and how those shape the contours of the technology itself. The device question in one of the most tangible places where we see these efforts, where actors are engaging in specific discursive work to make their form seem viable and secure certain powerful allies. While there is internal competition between the coalition about specific technological capability and designs, they are more broadly trying to fend off alternate forms and the advantages they have. This demonstrates the complexity of SCOT and how actors are thinking about a wider range of stakeholders the innovation process. Where some of the canonical case studies of SCOT found actors attempting to create design solutions for a specific industry problem (e.g. inflatable tires), in this case we can see that laden in those decisions are not only attempts to distinguish between close competitors, but also to engage with broader opponents. How stakeholders move within the community between these two forms will have important implications for what the technology ultimately looks like and how it can be used.

Chapter 4 is in some ways an extension of this dispute, but beyond any physical design issue, the contestation is over the very definition of the technology itself – and therefore about the community itself, who belongs, and who has authority to speak for the technology. Each articulation of “what AR is” is an attempt to include and exclude certain technologies outright, as

well as the actors associated with them as inside or outside of the community. The definitions themselves emerge not from any one source, rather they are rooted in historical commitments from other disciplines and attempt to distinguish AR from other technologies. The drive to articulate this protected class of things as “AR” was to secure certain types of institutional resources and legitimacy, while also securing the authority to define the technology. Faced with newcomers who are trying to include emerging devices and applications into the category, again the community is confronted with the question of whether to accommodate these technologies (and the stakeholders who accompany them) or defend against them.

At the expense of welcoming new disciplines and perspectives, certain segments within the community are holding the definition paramount. Some do so to preserve their own interests, others have legitimate concerns about alternate definitions, and still others are using the definition to exclude forms and uses they do not like. The definitional question is another way actors are attempting to stabilize relationships, as a step for negotiating authority. Where much of the CoP literature has focused on the positive elements of these groups, how they pass on knowledge, and how organizations can support CoP, what happened in this instance demonstrates how definitions can be used to segment off the population, exclude newcomers, or condition the terms of their participation. Either through a priori enforcement or a series of classification schemas, the ways that these definitions get solidified will ultimately affect the people who are in the community, the problems that are worth solving, and who has authority over the technology.

Chapter 5 begins to explore the tensions between early industry actors, their business models and sources of funding, and the relationships that are forming between them. A strong marketing and retail presence has taken hold in AR, in part because of the need for investment in

AR and in part because marketing actors seem the most excited about the possible reach of consumer AR. This has implications for the technology on a few fronts – the problem space of another industry becomes the one that AR actors must work on, which has directed the types of content that AR software can recognize and has expanded the capabilities of new devices in terms of storage, updating speed, and continuous visual search capacity. In addition to how these partnerships shape AR technology, they have also started to affect the content characteristics of other media as well, so as to be compatible with the emerging AR technologies. The prevalence of these examples has been worrisome to certain actors, who have countered with negative visions of advertising, used AR to subvert advertising, and even considering legal challenges to advertising. Examining how the AR community deals with the marketing contingent will go a long way toward understanding why certain technological features are developing, as well as people's willingness to engage with the technology. What is occurring in the AR industry is also a microcosm for the changing economic pressures of the emerging technology space, and how they are particularly influenced and beholden to these early clients and stakeholders. The fact that it is advertising allows us to examine how a recurring pressures is being exerted on a modern media form, particularly in relation to past media that faced similar struggles.

Chapter 6 examines the early standards discussion forming around AR, some of the factors contributing to the creation of these standards, the interlocking relationship between standards, and the institutional alignment of standards organizations. In this case different groups rallied around a particular historical narrative of technology, the Web, to argue for a standard markup language for AR. The process it subsequently underwent was the result of the social negotiation surrounding the standard, who chose to participate in the process, and the justifications for why the standard was necessary. The implementation of the standard (and what

was left off of it) further revealed tensions among the stakeholder groups, as well as limitations in each's approach. An opposing group promoting a different standard approach broke off, even while still claiming the same historical narrative. These two options formed from a particular standards community, one that brought the SSOs and industry actors under the same umbrella organization to form these partnerships. The early standards discussion paves the way for understanding a potential standards battle, as well as how these standards are saddled with commitments in the process of becoming an option.

Chapter 7 closes with a recent turn away from consumer AR, back toward industrial uses of AR as the place to secure investment. This contingent is aiming to deploy AR technologies into work environments, which changes the promises they make about the technology and changes the requirements for the technologies. The industrial influences can be seen in the way they reconfigure the problem space, the ways that industries are thinking about implementing AR, and in the material design of industrial AR devices. One of the big advantages of industrial AR is that employers can simply require their employees to wear these devices, though they cannot always dictate the social meanings that get embedded into the technology and the subtle ways that people negotiate their use. A high profile failure of industrial AR demonstrates how industrial technologies can be a site of negotiation for many other issues within an organization.

These chapters come together to explain how a group is coming together to stabilize relationships between different components of an artifact and between different actors of the technological activity. Contestation over the models and theories of technological innovation and diffusion that arise from these historical accounts may argue that certain perspectives and factors were overlooked, overemphasized, or overly predictive. The various actors, stages, frames, systems, and boundaries get debated as analytical categories and for their explanatory and

predictive potential. Where historians have an advantage compared to real-time mappings, in that they know the outcome of a particular socio-technical arrangement, real-time mappings can assert without conjecture the actual concerns that were being played out at any given moment.

While there can never be a perfect mapping of who the relevant social groups are, in this case there were clear actor groups and presumed actor groups who were defining the problem space for the technology, with varying degrees of overlap in both participation and technological frames for approaching certain designs. Instead of a unified group pushing a singular technology forward, actors were attempting to carve out clear boundaries through definitions and classifications of the technology, to preserve authority over the technology. The technology was not developing in isolation, rather there were broad economic forces and specific industry clients that are helping to define the problem spaces and pushing the actors working on the technologies in particular directions. Even early in the process, several groups were trying to unify a coalition and reverse certain commitments through standards, trying to replicate the historical trajectory (and struggles) of previous technologies. Lastly, a set of stakeholder groups are pushing the technology toward industrial different groups, with very specific prescriptions for the technology and how they might change the dynamics of certain workplaces. In all these ways, these groups are demonstrating the interpretive flexibility of material technologies and as well as our social understanding of them.

These chapters document how this group of people worked to bring an emerging technology into existence, at a time where very little was settled. The story is not a clean narrative of a brilliant discovery and steady diffusion, rather it is about how a CoP repeatedly and strategically met regularly to work on and train people how to promise the technology, attract investors, develop innovative technologies, create applications, and write standards.

Despite their efforts to unify they were not necessarily working toward a unified vision, rather different internal groups are hoping to affect the direction that these technologies take, which we see through a series of confrontations, negotiations, and attempts to consolidate power.

The ongoing development of AR will direct the research to specific sites, but the questions that I raised in the introduction remain the ones that need to be answered. Whether through AR is deployed through mobile devices or HWD, for consumer or industrial applications, as AR browsers or web browsers, for marketing purposes or for content alongside marketing, these will play a important role in shaping how people use the technology and how we understand the technology. The question of what content is possible and allowed is an ongoing part of the negotiation, as actors are negotiating this through design, policy, and standards. If people adopt AR or organizations implement AR, those sites of user deployment, negotiation, and contestation will be another place to see the social meanings of the technology being forged. The context of those implementations will also be important to study for their social implications, as the Nomad case demonstrates that the way these are implemented is crucial to understanding the social meaning that are assigned to these artifacts.

Hopefully this dissertation can serve as a starting point for future researchers to understand how AR actors were thinking about some of these issues, taking actions to make them possible, and contesting the technology across many different domains in ways that both shape its material form as well as set up some of the social, economic, and legal battles yet to come. Rather than having to piece together these developments or ignore the factors that led to these commitments through retrospective clues and secondary sources, this account offers a mapping of a particular moment of technological negotiation as it happened. AR can be another place to understand the continual process of stabilization, the shifting relationships between

production, consumption, materiality, and content, the development of new media industries, and ultimately the ways that all these developments are interrelated and shape and constrain how people encounter and use the technology.

APPENDIX

BACKGROUND AND WORK PRACTICE

I'd like to start with a little bit about your background. What led you to AR, and how did you first get involved in working with augmented reality?

Tell me about what projects you are currently working on. Could you walk me through what the goal of the project is?

[prompt for what is promising about this project, why they are working on it, what differentiates this from other AR projects]

MAPPING THE FIELD

AR is a growing field with a lot of people and interests involved – at this moment, who would you say are the key players and stakeholders who are important to the development of AR?

[follow up: different priorities and goals of each actor category, mapping who they view to be the center and periphery of the industry]

What group or groups would you say you belong to, represent, or most strongly associate with?

What places, companies, and/or academic institutions would you say are really driving AR development?

Who would you say are the leading voices and people who are influential to AR?

Would you say there is a generally accepted definition of AR or is there still some disagreement about what is and isn't AR?

PROMISING AREAS OF AR

What fields would you say are really embracing AR technology?

**-Follow up - why do you think those fields in particular are enthusiastic about AR?
How is AR being successfully presented to these fields?**

What fields do you think are really promising areas for AR development?

- Follow up - why do you think those areas are promising?

In what fields do you think AR could have a significant impact?

-Follow up – are there any applications or technologies you can think of that would be beneficial for a particular field? How would you explain to someone in that field why AR would help them?

Can you give me an example of a project or application that really illustrates the promise of AR technology?

What type of claims do you hear people making when they talk about the potential of AR?

Follow up: Who is making these claims? What are they trying to accomplish?

What's the coolest AR application you've seen?

PESSIMISM ABOUT AR

Are there any claims or promises that people make with AR that you think are over the top?

[Follow Up: what specifically do you think is exaggerated? What do you think causes them to make these claims?]

Are there areas or technologies within AR that you think are being over-hyped by those who cover the industry?

What concerns do people raise about AR?

[Follow Up: what do you think about those concerns? How legitimate do you think those concerns are? Are there concerns which you think are overblown? How much do people that you know/work with think about those concerns?]

Are there areas or technologies within AR that you think won't find traction?

Are there fields or places where you think AR technology shouldn't be used? What would be an example of something that concerns you?

Military? Crime? Privacy issues? Porn?

What legal issues do you foresee AR raising?

BARRIERS

Do you think there Is there a technological barrier that is holding AR back? And if so, what do you think that is?

Do you think there are social barriers that is preventing people from adopting/using AR?

In what area does AR need to improve the most to reach its full potential?

Put another way – if you could magically make one thing possible with AR, what would it be?

HISTORICAL TRAJECTORY OF AR

How did we arrive at this particular moment in the development of AR?

What would you say were the important milestones in the development of AR?

When did you first think AR was possible?

When did you see your first demonstration of AR?

THE ROLE OF CONFERENCES (FOR CONFERENCE SPEAKERS)

How many AR conferences have you attended?

What do you do when you are at the conferences?

Prompt: If speakers, who. If exhibitions, what are they looking for.

Do you find the conferences to be productive/helpful?

Why? Why not?

Who typically attends these conferences?

What would change, if anything, if they stopped having conferences?

CONSTRUCTING THE USER

Who are the current users of AR?

What kinds of applications are they using a lot now or getting excited about?

What do you think attracts people to AR? Information asymmetry? Voyuerism? Cool factor?

How are you thinking about the design of AR artifacts?

What is the thinking behind certain decisions you make in designing these products?

What are the important priorities for design, in terms of how you think users will want certain features and the tradeoffs that you have to make?

IMAGINED FUTURES

Where do you see AR going in the next 5 years?

What would you say would be the best case scenario for AR development?

What kind of society would that look like?

What's the worst case scenario for AR in the future?

If you had the ear of the AR community, what would you try to persuade them?

Do you think people will be persuaded by technical promises or practical promises?

How might Sci-fi visions affect imaginations of the future?

OGC

How does the OGC view the burgeoning industry of augmented reality and how the OGC can help in its development?

From your perspective, how would you characterize the current state of augmented reality – what challenges are facing the technology, the industry?

How do you think the OGC process could be beneficial to the field, to a burgeoning AR industry?

What sorts of stakeholders are emerging around AR, and what is their level of participation in the standards process?

What might you say to someone in AR that says, the industry is still maturing, its too early to standardize?

What is the worst case scenario, if the field does not standardize – is there a cautionary tale that comes to mind?

The opposite – is there some sort of technology that you can recall that might be analogous to where AR is right now, where their openness to standards was a success?

What might that look like for AR?

These are the standards listed as OGC activities relevant to AR Standards – would you mind if we sort of ran through each of them, how they got started, and what they are designed to address?

REFERENCES

- Abbate, J. (1999). *Inventing the Internet*. Cambridge, Mass: MIT Press.
- Abrash, M. (2012). Why you won't see hard AR anytime soon. *Ramblings in Valve Time*. Retrieved from <http://blogs.valvesoftware.com/abrash/why-you-wont-see-hard-ar-anytime-soon/>
- Acar, G., Eubank, C., Englehardt, S., Juarez, M., Narayanan, A., & Diaz, C. (2014). The Web never forgets: Persistent tracking mechanisms in the wild. In *Proceedings of Computer and Communications Security (CCS)*,
- Akrich, M. (1992). [The De-Description of Technical Objects](#). In Bijker, W. E., & Law, J. (Eds), *Shaping technology/building society: Studies in sociotechnical change* (pp. 205-244). Cambridge, Mass: MIT Press.
- Akrich, M. and Latour, B. (1992). A summary of a convenient vocabulary for the semiotics of human and nonhuman assemblies. In Bijker, W. E., & Law, J. (Eds). *Shaping technology/building society: Studies in sociotechnical change*. Cambridge, Mass: MIT Press.
- Anderson, S. P., & Gabszewicz, J. J. (2006). The Media and Advertising: A Tale of Two-Sided Markets. *Handbooks in Economics*, 25, 567-614.
- Azuma, R. (1997). "A Survey of Augmented Reality." *Presence*, 6(4), 355-385.
- Barba, E., Mynatt, E. D., & MacIntyre, B. (2012). Here we are! Where are we? Locating mixed reality in the age of the smartphone. *Proceedings of the IEEE*. doi:10.1109/JPROC.2011.2182070
- Barry, A. (2001). *Political Machines: Governing a Technological Society*. London: Athlone Press.
- Bazerman, C. (1999). *The languages of Edison's light*. Cambridge: MIT Press.
- Berkhout, F. (2006). Normative expectations in systems innovation. *Technology Analysis & Strategic Management*, 18, 299-311.
- Berlin, L. (2009). Kicking reality up a notch. *New York Times*. Retrieved from http://www.nytimes.com/2009/07/12/business/12proto.html?_r=2&scp=15&sq=Sportvisi on&st=cse&
- Biermann, B.C. (2012). The AR AD Takeover: Augmented Reality and the Reappropriation of Public Space. Retrieved from <http://www.theheavyprojects.com/documents/adtakeover.pdf>
- Bijker, W. (1995). *Of bicycles, bakelites, and bulbs: Toward a theory of sociotechnical change*. Cambridge, Mass: MIT Press.
- Bijker, W. E., Hughes, T. P., & Pinch, T. J. (1987). *The Social construction of technological systems: New directions in the sociology and history of technology*. Cambridge, Massachusetts: MIT Press.
- Bijker, W. E., & Law, J. (1992). *Shaping technology/building society: Studies in sociotechnical change*. Cambridge, Massachusetts: MIT Press.
- Blind, K. (2004). *The economics of standards: Theory, Evidence, and Policy*. Northampton, Massachusetts: Edward Elgar Publishing, Inc.
- Boczkowski, P. J. (2004). *Digitizing the news: Innovation in online newspapers*. Cambridge, Massachusetts: MIT Press.
- Boczkowski, P. J. (2010). *News at work: Imitation in an age of information abundance*. Chicago:

- The University of Chicago Press.
- Boczkowski P. & Siles I. (2014). "Steps Toward Cosmopolitanism in the Study of Media Technologies: Integrating Scholarship on Production, Consumption, Materiality, and Content." In Gillespie, T., In Boczkowski, P. J., & In Foot, K. A. (Eds). *Media technologies: Essays on communication, materiality, and society*.
- Boden, D. (1994). *Business of Talk*. Cambridge: Polity Press.
- Bolter, J. D., & Grusin, R. (1999). *Remediation: Understanding new media*. Cambridge, Massachusetts: MIT Press.
- Borup, M., Brown, N., Konrad, K. and Van Lente, H. (2006). The sociology of expectations in science and technology. *Technology Analysis & Strategic Management* 18(3-4): 285–298.
- Bowker, G. C., & Star, S. L. (1999). *Sorting things out: Classification and its consequences*. Cambridge, Mass: MIT Press.
- Bresnahan, T.F. and Yin, P.L. (2007). Standard setting in markets: the browser war. In Greenstein, S. M., & Stango, V. (Eds). *Standards and public policy*. Cambridge: Cambridge University Press.
- Brown, J.S. & Duguid, P. (2001). Knowledge and Organization: A Social-Practice Perspective. *Organization Science* 12(2). 198-213.
- Brown, N. and Kraft, A. (2006). Blood Ties: Banking the stem cell promise, *Technology Analysis and Strategic Management*, 18(3/4), 313-327
- Brown, N., Rappert, B. and Webster, A. (Eds.) (2000). *Contested Futures: A Sociology of Prospective Technoscience*, Aldershot: Ashgate.
- Brynjolfsson, E., & Saunders, A. (2010). *Wired for innovation: How information technology is reshaping the economy*. Cambridge, Mass: MIT Press.
- Buchart, B. (2011). "Augmented Reality for Smartphones: A Guide for Developers and Content Publishers." Techwatch Report, JISC Observatory. Retrieved from http://observatory.jisc.ac.uk/docs/AR_Smartphones.pdf
- Busch, L. (2011). *Standards: Recipes for reality*. Cambridge, Mass: MIT Press.
- Cakmakci O. and Rolland. J.P. (2006). Head-Worn Displays: A Review. *Journal of Display Technologies* 2, 199-216.
- Callon, M. (1987). Society in the making: the study of technology as a tool for sociological analysis. In Bijker, W. E., Hughes, T. P., & Pinch, T. J. (Eds). *The Social construction of technological systems: New directions in the sociology and history of technology*. Cambridge, Mass: MIT Press.
- Callon, M. (1998). "An Essay on Framing and Overflowing: Economic Externalities Revisited by Sociology." In Michel Callon, (Ed.), *The Laws of the Markets* (pp. 244-269). Malden, Mass.: Blackwell.
- Cargill, C. F. (1997). *Open systems standardization: A business approach*. Upper Saddle River, NJ: Prentice Hall PTR.
- Cargill, C. F., and Bolin, S. (2007). Standardization: a failing paradigm. In Greenstein, S. M., & Stango, V. (Eds). *Standards and public policy*. Cambridge: Cambridge University Press.
- Carr, N.G. (2010). *The shallows: What the Internet is doing to our brains*. New York: W.W. Norton.
- Castells, M. (2000). *The rise of the network society*. Malden, MA: Blackwell.
- Castells, M. (2001). *The Internet galaxy: Reflections on the Internet, business, and society*. Oxford: Oxford University Press.

- Caudell, T. P., & Mizell, D. W. (1992). Augmented Reality: An Application of Heads-Up Display Technology to Manual Manufacturing Processes. *Proceedings of the Hawaii International Conference on System Sciences*, 25, 659.
- Charmaz, K (2002). *Constructing grounded theory: A practical guide through qualitative analysis*. New York: Sage.
- Cooke, P. (2006). New media and new economy cluster dynamics. In Lievrouw, L. A., & Livingstone, S. M. (Eds). *Handbook of new media: Social shaping and social consequences of ICTs*. London: SAGE.
- Cox, A. (2005). What are communities of practice? A comparative review of four seminal works. *Journal of Information Science*, 31(6), 527-540
- Crain, M. (2014). Financial markets and online advertising: reevaluating the dotcom investment bubble. *Information, Communication & Society*, 17(3), 371-384.
- Cusumano, M. A., Mylonadis, Y., & Rosenbloom, R. S. (1992). Strategic Maneuvering and Mass-Market Dynamics: The Triumph of VHS over Beta. *The Business History Review*, 66(1), 51-94.
- Cusumano, M. A., & Yoffie, D. B. (1998). *Competing on Internet time: Lessons from Netscape and its battle with Microsoft*. New York, NY: Free Press.
- Das, S. & Kramer, A. (2013). Self Censorship on Facebook, in *Proceedings of the Seventh International AAAI Conference on Weblogs and Social Media*, 120-127.
- David, Paul A. and Shane Greenstein (1990), The Economics of Compatibility Standards: An Introduction to Recent Research.” *Economics of Innovation and New Technology*, 1(1-2), 3-43.
- Denning, T., Kohno, T., & Dehlawi, Z., (2014). In situ with bystanders of augmented reality glasses: Perspectives on recording and privacy-mediating technologies. In *Proceedings of ACM Conference on Human Factors in Computing Systems (SIGCHI)*, 2377-2386.
- Denzin, N. K. and Lincoln, Y.S. (2005). *The Sage Handbook of Qualitative Research*. (3rd ed.). Thousand Oaks: Sage.
- Deuten, J., Rip, A., & Jelsma, J. (1997). Societal embedding and product creation management. *Technology Analysis & Strategic Management*, 9(2), 131-148.
- Douglas, S. (1987). *Inventing American Broadcasting, 1899-1922*. Baltimore, MD: Johns Hopkins University Press.
- Duguid, P. (2005). The art of knowing: Social and tacit dimensions of knowledge and the limits of the community of practice. *The Information Society*. 21(2), 109-118.
- Edwards, P. N. (1997). *The closed world: Computers and the politics of discourse in Cold War America*. Cambridge, Mass: MIT Press.
- Emerson, R. Fretz, R. & Shaw, L. (1995). *Writing Ethnographic Fieldnotes*. Chicago: University of Chicago Press.
- Eppler, M., & Mengis, J. (2004). The Concept of Information Overload: A Review of Literature from Organization Science, Accounting, Marketing, MIS, and Related Disciplines. *The Information Society*, 20(5), 325-344.
- Epstein, D. (2010). Constructing the Information Society: The Binding Nature of Nonbinding Debates About Internet Governance, in *38th Annual Telecommunication Policy Research Conference (TPRC)*.
- Evans, D. S. (2009). The Online Advertising Industry: Economics, Evolution, and Privacy *Journal of Economic Perspectives*, 23(3), 37-60.
- Epps, S.R., Parrish, M., Hammond, J., O’Connell, J., VanBoskirk, S., & Vokshi, A. (2013).

- Google glass: What marketers need to know. *Forrester Research*. Cambridge, MA: Forrester Publishing.
- Feiner S. (1999). The Importance of Being Mobile: Some Social Consequences of Wearable Augmented Reality Systems. *In Proc. of International Workshop on Augmented Reality, San Francisco, CA, October 20-21*, 145-148.
- Feiner, S., MacIntyre, B., Hollerer, T., and Webster, A. (1997). "A touring machine: prototyping 3D mobile augmented reality systems for exploring the urban environment." *In Proceedings of ISWC '97 (First International Symposium on Wearable Computers)*. Cambridge: MA; 74–81.
- Flew, T. and McElhinney, S. (2006). Globalization and the structure of new media industries. In Lievrouw, L. A., & Livingstone, S. M. (Eds). *Handbook of new media: Social shaping and social consequences of ICTs*. London: SAGE.
- Flichy, P. (2007). *Understanding technological innovation: A socio-technical approach*. Cheltenham, Glos, UK: Edward Elgar.
- Fontaine, M. A., & Millen, D. R. (2002). Measuring effectiveness: understanding the value of communities. *Knowledge Management Review*, 5(3), 24-27.
- Fontana A. & Frey, J.H. (2005) The interview: from structured questions to negotiated text. In Denzin, N. K., & Lincoln, Y. S. (Eds). *Handbook of qualitative research*. Thousand Oaks, Calif: Sage Publications.
- Freeman, C. and Soete, L. (1997). *The Economics of Industrial Innovation*. (3rd ed.), Cambridge: The MIT Press.
- Fortun, M. (2008). *Promising Genomics: Iceland and deCODE Genetics in a World of Speculation*. Berkeley: University of California Press.
- Fox, S. (2000). Communities Of Practice, Foucault And Actor-Network Theory. *Journal of Management Studies*, 37(6), 853-868.
- Frank, R. H., & Cook, P. J. (1995). *The winner-take-all society: How more and more Americans compete for ever fewer and bigger prizes, encouraging economic waste, income inequality, and an impoverished cultural life*. New York: Free Press.
- Geels, F. W. (2004). From sectoral systems of innovation to socio-technical systems. *Research Policy*, 33, 897-920.
- Geels, F. W., & Smit, W. A. (2000). Failed technology futures: pitfalls and lessons from a historical survey. *Futures*, 32(9), 867-885.
- Gieryn, T. F. (1983). Boundary-Work and the Demarcation of Science from Non-Science: Strains and Interests in Professional Ideologies of Scientists. *American Sociological Review*, 48(6), 781-795.
- Gieryn, T. F. (1999). *Cultural boundaries of science: Credibility on the line*. Chicago: University of Chicago Press.
- Gillespie, T. (2007). *Wired shut: Copyright and the shape of digital culture*. Cambridge, Mass: MIT Press.
- Greenstein, S., & Rysman, V. (2007). Introduction. In Greenstein, S. M., & Stango, V. (Eds). *Standards and public policy*. Cambridge: Cambridge University Press.
- Greenstein, S. M., & Stango, V. (2007). *Standards and public policy*. Cambridge: Cambridge University Press.
- Groten, D. (2014). First Step to Interoperability, Layar News. Retrieved from <https://www.layar.com/news/blog/2014/02/11/first-step-to-interoperability/>
- Grover, D. (2014). Augmented reality history, background, and philosophy. Retrieved from

- <https://wiki.mq.edu.au/display/ar/Augmented+reality+history%2C+background+and+philosophy>
- Grunwald, A. (2004). Vision Assessment as a new element of the FTA toolbox. Paper presented at the *EU-US seminar: New Technology Foresight, Forecasting & Assessment Methods*. Seville 13-14, May 2004
- Gudis, C. (2004). *Buyways: Billboards, automobiles, and the American landscape*. New York: Routledge.
- Gupta, A., and Ferguson, J. (1997). Discipline and practice: “The field” as site, method, and location in anthropology. In Gupta, A., and Ferguson, J. (Eds.) *Anthropological locations: Boundaries and grounds of a field science*, (pp. 1-46). Berkeley and Los Angeles: University of California Press.
- Haas, P. M. (1992). Introduction: Epistemic Communities and International Policy Coordination. *International Organization*, 46(1), 1-35.
- Hedgecoe, A. and Martin, P. (2003). The drugs don't work: Expectations and the shaping of pharmacogenetics. *Social Studies of Science*, 33, 327–364.
- Heggestuen, J. (2013). One In Every 5 People In The World Own A Smartphone, One In Every 17 Own A Tablet. Retrieved from <http://www.businessinsider.com/smartphone-and-tablet-penetration-2013-10#ixzz359OkI150>
- Henderson, S., & Feiner, S. (2007). Augmented reality for maintenance and repair (ARMAR). *Columbia University, New York Department of Computer Science*.
- Henderson, S., & Feiner, S. (2009). Evaluating the benefits of augmented reality for task localization in maintenance of an armored personnel carrier turret. *Mixed and Augmented Reality, ISMAR 2009, 8th IEEE International Symposium*.
- Henderson, S., & Feiner, S. (2011). Exploring the benefits of augmented reality documentation for maintenance and repair. *IEEE Transactions on Visualization and Computer Graphics*, 17(10) 1355-68.
- Hesmondhalgh, D. (2002). *The cultural industries*. London: SAGE.
- Hildreth, P. M., & Kimble, C. (2004). *Knowledge networks: Innovation through communities of practice*. Hershey PA: Idea Group Pub.
- Hill, A., MacIntyre, B., Gandy, M., Davidson, B., & Rouzati, H. (2010) – KHARMA: KML/HTML Augmented Reality Mobile Architecture, in *IEEE Proc. of the 9th International Symposium on Mixed and Augmented Reality*.
- Hoang, T., Smith, B., & Thomas, B. H. (2012). The Augmented Reality Space Continuum. Presented at *ISMAR2012 Workshop for Classifying the AR Presentation Space*, Atlanta, GA
- Hughes, T. P. (1983). *Networks of power: Electrification in Western society*. Baltimore: Johns Hopkins University Press.
- Hughes, T.P (1989). *American Genesis: A Century of Invention and Technological Enthusiasm*. New York, NY: Viking.
- Hughes, T. P. (1987). The evolution of large technological systems. In W. Bijker, T. Hughes, and T. Pinch (Eds.) *The social construction of technological systems: New directions in the sociology and history of technology* (pp. 51-82). Cambridge, MA: MIT Press
- Hugues, O., Fuchs, P., & Nannipieri, O. (2011). New Augmented Reality Taxonomy: Technologies and features of augmented environment. In B. Furht (Ed.), *Handbook of Augmented Reality*, 47-63.

- Jasanoff, S. (2005). *Designs on Nature: Science and Democracy in Europe and the United States*. Cambridge: Harvard University Press.
- Jenkins, H. (2006). *Convergence culture: Where old and new media collide*. New York: New York University Press.
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, 17(11), 4302-11.
- Katz, M. L., & Shapiro, C. (1994). Systems Competition and Network Effects. *Journal of Economic Perspectives*, 8(2), 93-115.
- Katz, J.E. & Sugiyama, S. (2005). Mobile phones as fashion statements: the co-creation of mobile communication's public meaning. In Ling, R. S., & Pedersen, P. E. (Eds). *Mobile communications: Re-negotiation of the social sphere*. London: Springer.
- Kelty, C. (2014). The Fog of Freedom. In Gillespie, T., Boczkowski, P. J., & Foot, K. A. (Eds). *Media technologies: Essays on communication, materiality, and society*.
- Kendall, L. (1999). Nerd Nation: Images of Nerds in US Popular Culture. *International Journal of Cultural Studies*, 2(2), 260-283.
- Kilpatrick, M. (2011). Startups Aim to Make Augmented Reality Easy and Cheap. Read/Write Web Research, Retrieved from <http://readwrite.com/2011/02/25/augmented-reality-publishing-i#awesm=~oHHMrDtgpOzG8S>
- Kim, G. J., Kim, Y., & Perey, C. (2011). A Proposal for Reference AR System Architecture Based on Standardized Content Model. International AR Standards Meeting. Retrieved from [http://www.perey.com/ARStandards/\[Kim,Perey\]AR_Architecture_Survey_ISO-JTC1.pdf](http://www.perey.com/ARStandards/[Kim,Perey]AR_Architecture_Survey_ISO-JTC1.pdf)
- Kirsch, D., Goldfarb, B., & Gera, A. (2009). Form or substance: the role of business plans in venture capital decision making. *Strategic Management Journal*, 30(5), 487-515.
- Klein, H. K., & Kleinman, D. L. (2002). The Social Construction of Technology: Structural Considerations. *Science, Technology, & Human Values*, 27, 1, 28-52.
- Kline, R. (2003). Resisting consumer technology in rural America: the telephone and electrification. In Oudshoorn, N., & Pinch, T. J. (Eds). *How users matter: The co-construction of users and technologies*. Cambridge, Mass: MIT Press.
- Kline, R., & Pinch, T. (1996). Users as Agents of Technological Change: The Social Construction of the Automobile in the Rural United States. *Technology and Culture*, 37(4), 763.
- Koskela, H. (2002). Video Surveillance, Gender, and the Safety of Public Urban Space: "Peeping Tom" Goes High Tech?. *Urban Geography*, 23, 257-278.
- Lampland, M., & Star, S. L. (2009). *Standards and their stories: How quantifying, classifying, and formalizing practices shape everyday life*. Ithaca: Cornell University Press.
- Latour, B. (1987). *Science in action: How to follow scientists and engineers through society*. Cambridge, Mass: Harvard University Press.
- Latour, B. (1988). Mixing humans and nonhumans together: The sociology of a door-closer. *Social Problems* 35, 298-310.
- Latour, B. (1992). Where are the Missing Masses? The Sociology of a Few Mundane Artifacts." In W.E. Bijker & J. Law (Eds.) *Shaping Technology/Building Society: Studies in Sociotechnical Change*, (pp. 225-258). Boston: MIT Press.

- Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge: Cambridge University Press.
- Law, J. (1987). Technology and Heterogeneous Engineering: the Case of the Portuguese Expansion. In W. Bijker, T. Hughes and T. Pinch (Eds.), *The Social Construction of Technological Systems* (pp. 111-134). Cambridge, Mass., M.I.T. Press.
- Lechner, M. and Tripp, M. (2010) – ARML - An Augmented Reality Standard. *Mobile Augmented Reality Summit*, Barcelona.
- Lemordant, J. (2010). Summary International AR Standards Workshop. http://www.perey.com/ARStandards/International_AR_Standards_Meeting_Summary_Oct%2015.pdf
- Lessig, L. (1999). *Code and other laws of cyberspace*. New York: Basic Books.
- Levy, S. (2011). *In the plex: How Google thinks, works, and shapes our lives*. New York: Simon & Schuster.
- Liao, T. (2012). A Framework for Debating Augmented Futures: Classifying the Visions, Promises and Ideographs Advanced About Augmented Reality. In *IEEE Proc. of the 11th International Symposium on Mixed and Augmented Reality (ISMAR)*.
- Liao, T., & Humphreys, L. (2014). Layar-ed places: Using mobile augmented reality to tactically reengage, reproduce, and reappropriate public space. *New Media & Society*.
- Liebowitz, S. J., & Margolis, S. (1999). *Winners, losers & Microsoft: Competition and antitrust in high technology*. Oakland, Calif: Independent Institute.
- Lievrouw, L. A., & Livingstone, S. M. (2002). *Handbook of new media: Social shaping and consequences of ICTs*. London: SAGE.
- Lindsay, C. (2003). From the shadow: users as designers, producers, marketers, distributors, and technical support. In Oudshoorn, N., & Pinch, T. J. (Eds). *How users matter: The co-construction of users and technologies*. Cambridge, Mass: MIT Press.
- Lofland, J., Snow, D., Anderson, L., and Lofland, L.H. (2006) *Analyzing Social Settings: A Guide to Qualitative Observation and Analysis*. Belmont, CA: Wadsworth.
- Lombard, M., & Ditton, T. (1997). At the Heart of It All: The Concept of Presence. *Journal of Computer-Mediated Communication*, 3(2).
- Lösch, A. (2006). “Anticipating the Futures of Nanotechnology: Visionary Images as Means of Communication.” *Technology Analysis & Strategic Management*, 18(3/4), 393-409.
- Lu, W. & Duh, H.B (2012). Designing Augmented Reality Visualizations based on a Model of Human Attention. Presented at *ISMAR2012 Workshop for Classifying the AR Presentation Space, Atlanta, GA*.
- Mann, S. (1997). Wearable computing: a first step toward personal imaging, *IEEE Computer*. 30(2) 25-32.
- Mann, S. (2013). Vision 2.0: Instead of carrying around smartphones: we might soon be donning smart eyewear. A pioneer of computerized vision shares some tips. *IEEE Spectrum*, 50(3), 40.
- Mann, S., Nolan, J., & Wellman, B. (2003). Sousveillance: Inventing and Using Wearable Computing Devices for Data Collection in Surveillance Environments. *Surveillance & Society*, 1(3), 331-355.
- Mansell, R. (2004). Political Economy, Power and New Media. *New Media & Society*, 6, 1, 96-105.
- Marner, M., Irlitti, A., & Thomas, B. H. (2013). Improving Procedural Task Performance with

- Augmented Reality Annotations. In *IEEE Proc. of the International Symposium on Mixed and Augmented Reality (ISMAR13)*.
- Marvin, C. (1988). *When Old Technologies Were New*. Oxford: Oxford University Press.
- Marwick, A. (2013). *Status Update: Celebrity, Publicity and Branding in the Social Media Age*. New Haven: Yale University Press.
- McPhee, R.D. & Iverson, J. (2009). Agents of constitution in comunidad: Constitutive processes in organizations. In L.L. Putnam and A.M. Nicotera (Eds.), *Building theories of organization: The constitutive role of communication*. New York: Routledge
- McPhee, R.D. & Zaug, P. (2000). The communicative constitution of organizations: A framework for explanation. *Electronic Journal of Communication* 10(1/2): 1–16.
- McStay, A. (2010). *Digital advertising*. Houndmills, Basingstoke, Hampshire: Palgrave Macmillan.
- Metz, R. (2012). Augmented Reality is Finally Getting Real. MIT Technology Review, Retrieved from <http://www.technologyreview.com/news/428654/augmented-reality-is-finally-getting-real/>
- Michael, M. (2000). Futures of the Present: From Performativity to Prehension. In N. Brown, B. Rappert, & A. Webster (Eds.), *Contested Futures: A Sociology of Prospective Technoscience*, (pp. 21-39). Aldershot: Ashgate.
- Michael, M. and Brown, N. (2004) ‘The Meat of the Matter: Grasping and Judging Xenotransplantation’, *Public Understanding of Science*, 13, 379-397.
- Milgram, P., & Kishino, F. (1994). A taxonomy of mixed reality displays. *IEICE Transaction on Information Systems*, E77-D(12)
- Miller, C. (2014). Technology’s Man Problem. New York Times. Retrieved from http://www.nytimes.com/2014/04/06/technology/technologys-man-problem.html?_r=0
- National Research Council. (1999). *Funding a Revolution: Government Support for Computing Research*. Washington, DC: The National Academies Press.
- Neff, G. (2012). *Venture labor: Work and the burden of risk in innovative industries*. Cambridge, Mass: MIT Press.
- Neff, G., Wissinger, E., & Zukin, S. (2005). Entrepreneurial Labor among Cultural Producers: ‘Cool’ Jobs in ‘Hot’ Industries. *Social Semiotics*, 15(3), 307-334.
- Nielsen (2014). Multi-platform Gaming: For the Win! Retrieved from <http://www.nielsen.com/us/en/newswire/2014/multi-platform-gaming-for-the-win.html>
- Normand, J.-M., Servieres, M., Moreau, G., & 3rd Augmented Human International Conference, AH’12. (2012). A new typology of Augmented Reality applications. *Acm International Conference Proceeding Series*.
- Olsson T. and Salo. M. (2011). Online user survey on current mobile augmented reality applications. In *Proc. of International Symposium on Mixed and Augmented Reality ‘11*, Basel, Switzerland. (pp.75-84).
- Orlikowski, W. (2000). Using Technology and Constituting Structures: A Practice Lens for Studying Technology in Organizations. *Organization Science*, 11(4): 404-428
- Oudshoorn, N., & Pinch, T. J. (2003). *How users matter: The co-construction of users and technologies*. Cambridge, Mass: MIT Press.
- Pachal, P. (2013). Why Google Glass Isn’t Segway. Mashable, Retrieved from <http://mashable.com/2013/07/29/why-google-glass-isnt-the-segway/>
- Papagiannakis, G., Singh, G., and Magnenat-Thalmann, N. (2008). A survey of mobile and

- wireless technologies for augmented reality systems. *Journal of Computer Animation and Virtual Worlds*, 19(1), 3-22.
- Pavlik, J. V., & Bridges, F. (2013). The Emergence of Augmented Reality (AR) as a Storytelling Medium in Journalism. *Journalism & Communication Monographs*.
- Pfaffenberger, B. (1992). Technological Dramas. *Science, Technology, & Human Values*, 17(3), 282-312.
- Pinch, T.J. and Bijker, W.E. (1987). The social construction of facts and artifacts: or how the sociology of science and the sociology of technology might benefit each other. In Bijker, W. E., Hughes, T. P., & Pinch, T. J. (Eds). *The Social construction of technological systems: New directions in the sociology and history of technology*. Cambridge, Mass: MIT Press.
- Pollock, N. and Williams, R. (2010). The business of expectations: How promissory organizations shape technology and innovation. *Social Studies of Science* 40(4), 525-548.
- Powell, W.W. (2007). The capitalist firm in the twenty-first century: emerging patterns in western enterprise. In Wharton, A. S. (Ed). *The sociology of organizations: An anthology of contemporary theory and research*. Los Angeles, Calif: Roxbury Pub. Co.
- Pratt, A. (2000). New media, the new economy and new spaces. *Geoforum*, 31(4), 425-436.
- Putnam L.L. & Nicotera A.M. (2009) (Eds.), *Building theories of organization: The constitutive role of communication*. New York: Routledge
- Renhart, G., and Patron, C. (2003). Integrating augmented reality in the assembly domain – fundamentals, benefits, and applications. *CIRP Annual Manufacturing Technologies*, 52, 5-8.
- Rheingold, H. (1991). *Virtual reality: The revolutionary technology of computer-generated artificial worlds - and how it promises to transform society*. New York: Simon & Schuster.
- Rip, A. & Schot, J., (2002). Identifying loci for influencing the dynamics of technological Development. In: Sorensen, K.H., Williams, R. (Eds.), *Shaping technology, guiding policy: concepts, spaces and tools* (pp. 158-176). Edward Elgar, Cheltenham, UK.
- Rochet, J.-C., & Tirole, J. (2003). Platform Competition in Two-Sided Markets. *Journal of the European Economic Association*, 1(4), 990-1029.
- Roesner, F., Kohno, T., & Molnar, D. (2014). Security and privacy for augmented reality systems. *Communications of the ACM*, 57(4), 88-96.
- Russell, S. (1986). The social construction of artefacts: a response to Pinch and Bijker, *Social Studies of Science*, vol. 16, 331–46
- Schein, E. H. (1996). Culture: The Missing Concept in Organization Studies. *Administrative Science Quarterly*, 41(2), 229.
- Schot, J. & de la Bruheze, A.A. (2003). The mediated design of products, consumption, and consumers in the twentieth century. In Oudshoorn, N., & Pinch, T. J. (Eds). *How users matter: The co-construction of users and technologies*. Cambridge, Mass: MIT Press.
- Schramm, W. (1988). *The story of human communication: Cave painting to microchip*. New York: Harper & Row.
- Shapin, S. (2008). *The Scientific Life: A Moral History of a Late Modern Vocation*. Chicago: Chicago University Press.
- Shapiro, C., & Varian, H. R. (1998). *Information Rules: A Strategic Guide to the Network Economy*. Boston: Harvard Business School Press.

- Silverstone, R. (1999). *Why Study the Media*. London: Sage.
- Silverstone, R., & Haddon, L. (1996). Design and the domestication of information and communication technologies: Technical change and everyday life. In R. Silverstone & R. Mansell (Eds.), *Communication by design: The politics of information and communication technologies* (pp. 44-74). Oxford: Oxford University Press.
- Simcoe, T. (2007). Delay and de jure standardization: exploring the slowdown in Internet standards development. In Greenstein, S. M., & Stango, V. (Eds). *Standards and public policy*. Cambridge: Cambridge University Press.
- Slack, J. D. & Wise, J.M. (2006). Cultural Studies and Communication Technology. In Lievrouw L. & Livingstone, S (eds) *The Handbook of New Media: Updated Student Edition*. London: Sage. 141-162
- Smith, A. (2013). Smartphone Ownership 2013. Pew Research Internet Project. Retrieved from <http://www.pewinternet.org/2013/06/05/smartphone-ownership-2013/>
- Solove, D. J. (2007). *The future of reputation: Gossip, rumor, and privacy on the Internet*. New Haven: Yale University Press.
- Spurgeon, C. (2008). *Advertising and new media*. New York, NY: Routledge.
- Stango, V. (2004). The Economics of Standards Wars, *Review of Network Economics* 3(1), 2004: 1-19.
- Sterne, J. (2012). *MP3: The meaning of a format*. Durham: Duke University Press.
- Sterne, J. (2014). "What Do We Want?" "Materiality!" "When Do We Want It?" "Now!" In Gillespie, T., In Boczkowski, P. J., & In Foot, K. A. (Eds). *Media technologies: Essays on communication, materiality, and society*.
- Streeter, T. (2011). *The Net Effect: Romanticism, Capitalism and the Internet*. New York: New York University Press.
- Sturken, M. Thomas D. & Ball-Rokeach S. J. (Eds) (2004). *Technological Visions. The Hopes and Fears that Shape New Technologies*. Philadelphia: Temple University Press.
- Suchman, L.A. (2006). *Human-Machine Reconfigurations: Plans and Situated Actions*. Cambridge, UK: Cambridge University Press.
- Sutherland, I. E. (1965). "The Ultimate Display." *Proceedings of IFIP Congress*, 506-508.
- Sutherland, I. E. (1968). A Head Mounted Three-Dimensional Display. *Proceedings of AFIPS* 68, 757-764
- Suomela, R., & Lehikoinen, J. (2004). Taxonomy for visualizing location-based information. *Virtual Reality*, 8(2), 71-82.
- Tamarjan, D. (2012). Mobile AR browsers and what stops them from taking over the world. Retrieved from <http://augmentedtomorrow.com/mobile-ar-browsers-and-what-stops-them-from-taking-over-the-world/>
- Tang, A., Owen, C., Biocca, F., & Mou, W. (2003). Comparative Effectiveness of Augmented Reality in Object Assembly. *Chi Conference*, 73-80.
- Taylor, J. R., & Van, E. E. J. (2000). *The emergent organization: Communication as its site and surface*. Mahwah, N.J: Lawrence Erlbaum Associates.
- Tönnis, M., & Plecher, D. A. (2011). Presentation Principles in Augmented Reality – Classification and Categorization Guidelines. Techreport, Technische Universität München, 2011. Version 1.0.
- Tsang, M.M., Ho, S.C., & Liang, T.P. (2004). Consumer Attitudes Toward Mobile Advertising: An Empirical Study. *International Journal of Electronic Commerce*, 8(3), 65-78.
- Turkle, S. (2011). *Alone together: Why we expect more from technology and less from each*

- other*. New York: Basic Books.
- Turow, J. (2005). Audience Construction and Culture Production: Marketing Surveillance in the Digital Age. *The Annals of the American Academy of Political and Social Science*, 597(1), 103-121.
- Turow, J. (2011). *The daily you: How the new advertising industry is defining your identity and your worth*. New Haven: Yale University Press.
- Tutton, R. (2011). "Promising Pessimism: The Rise of Membrane Technology." *Social Studies of Science* 28(2), 221-254.
- Van Krevelen, D.W.F. & Poelman, R. (2010). A Survey of Augmented Reality Technologies, Applications, and Limitations. *The International Journal of Virtual Reality*, 9(2), 1-20.
- Van Lente (2000). Forceful futures: from promise to requirement. In Brown, N., Rappert, B., & Webster, A. (Eds). *Contested futures: A sociology of prospective techno-science*. Aldershot, England: Ashgate.
- van Lente, H. and Rip, A. (1998). "The Rise of Membrane Technology: From Rhetorics to Social Reality," *Social Studies of Science*, 28(2), 221-254.
<http://sss.sagepub.com/cgi/content/abstract/28/2/221>
- Van Maanen, J. and S. Barley (1984). "Occupational Communities: Culture and control in organizations." *Research in Organizational Behavior*, 6, 287-365.
- Visser, A. (2011). Survey of XML Languages for Augmented Reality Content, *Proc. of AR Standardization Forum*, Barcelona.
- Wacjman, J. (2000). Reflections on Gender and Technology Studies: In What State is the Art? *Social Studies of Science*, 30(3), 447-464.
- Wagner, D. (2007). Handheld Augmented Reality, Ph.D. Dissertation, Institute for Computer Graphics and Vision, Graz University of Technology
- Wassom, B. (2012). Augmented Legality 1.0. <http://www.wassom.com/wp-content/uploads/Augmented-Legality-1.0-E-book.pdf>.
- Wassom, B. (2014). IP in an Augmented Reality. *Landslide* 6(3).
http://www.americanbar.org/publications/landslide/2013-14/january-february/ip_an_augmented_reality.html
- Weick, K. E. (1995). *Sensemaking in organizations*. Thousand Oaks: Sage Publications.
- Wenger, E. (1998). *Communities of practice: Learning, meaning, and identity*. Cambridge: Cambridge University Press.
- Wenger, E. (2000). Communities of practice and social learning systems. *Organization*, 7(2), 225-246. doi: 10.1177/135050840072002
- Williams, R. (1990). *Television: Technology and cultural form*. New York: Schocken Books.
- Williams, R., & Edge, D. (1996). The social shaping of technology. *Research Policy*, 25(5), 865.
- Winner, L. (1993). Upon opening the black box and finding it empty: social constructivism and the philosophy of technology. *Science, Technology, & Human Values*, 18(3), 362-78
- Wellman, B. (2004). The Three Ages of Internet Studies: Ten, Five and Zero Years Ago. *New Media & Society*, 6(1), 123-129.
- Wenger, E., McDermott, R. A., & Snyder, W. (2002). *Cultivating communities of practice: A guide to managing knowledge*. Boston, Mass: Harvard Business School Press.
- Woolgar, S. (1991). Configuring the user: the case of usability trials. In Law, J. (Eds). *A Sociology of monsters: Essays on power, technology, and domination*. London: Routledge.
- Wu, T. (2010). *The Master Switch: The Rise and Fall of Information Empires*. New York:

- Random House.
- Wyatt, S. (2003). Non-users also matter: the construction of users and non-users of the Internet. In Oudshoorn, N., & Pinch, T. J. (Eds). *How users matter: The co-construction of users and technologies*. Cambridge, Mass: MIT Press.
- Wyatt, S. (2004). Danger! Metaphors at Work in Economics, Geophysiology, and the Internet. *Science, Technology, & Human Values*, 29(2), 242-261.
- Yuan, M. L., Ong, S. K., & Nee, A. Y. C. (2008). Augmented reality for assembly guidance using a virtual interactive tool. *International Journal of Production Research*, 46(7), 1745-1767.
- Zider, B. (1998). How venture capital works. *Harvard Business Review*, 76, 6.
- Zook, M. A. (2005). *The geography of the Internet industry: Venture capital, dot-coms, and local knowledge*. Malden, MA: Blackwell Pub.
- Zuboff, S. (1988) *In the age of the smart machine : the future of work and power*. New York: Basic Books.