

Ethical Implications of Implicit Bias in AI: Impact for Academic Libraries

Kim Paula Nayyer and Marcelo Rodriguez

Introduction

Academic libraries are exploring artificial intelligence (AI) applications that have the potential to create new or improved user experiences, streamline ways of working, and deliver new insights to their activities. Nevertheless, it is now clear that AI applications are not neutral technological solutions. They can embed and magnify prejudices and stereotypes, and they can perpetuate errors and limitations in training and accumulated datasets. At the same time, academic libraries abide by ethical considerations of social responsibility. If datasets and algorithmic *black boxes* in AI systems replicate or aggravate inappropriate discrimination in their use of information, or if they simply lack or ignore data, they can produce distorted outcomes. The ethical implications for academic libraries and end-users can be profound.

This chapter examines these issues, illustrates problematic outcomes, and identifies both the need for caution and some paths to the ethical use of AI applications in academic libraries. After a general exploration of the essence of machine learning (ML), this chapter explains what implicit bias is, how it enters ML applications, and why the problem is insidious and challenging. The authors present an illustrative review of the ethical foundations of the work of academic libraries and draw analogies to other professional interfaces with AI and implicit bias. Possible scenarios of ethically problematic outcomes in academic libraries are explored.

What Is Implicit Bias and How Does It Enter AI Applications?

To understand how AI applications can embody implicit bias, one must understand how modern AI works. Until a few decades ago, the phrase artificial intelligence conjured images of futuristic technologies ranging from friendly robotic companions to mission-fulfilling supreme computers like HAL 9000 of the classic film and book *2001: A Space Odyssey*.¹ Now, rather than being fanciful or theoretical, AI applications are present everywhere in modern life. Smartphones learn to spell friends' and family members' names, streaming video services seem to know what viewers want to watch before they do, and popular search engines predict what users are searching for after a few keystrokes.

The advent of powerful computing and vast amounts of usable data in digital form continue to push the proliferation of ML applications, and they are now commonplace in environments in which people provide and use services.² Generally, in library environments and elsewhere, AI is seen not as omniscient standalone operators but rather as components of larger products and processes.³ AI tools perform specific tasks in service of the goals of the system. To a greater or lesser degree and with varied effectiveness, the AI component replicates, or at least replaces, human thought in the fulfillment of that task.

Machine Learning: An Algorithm, Its Training Data, Iterative Learning, and Algorithm Self-Revision

Early AI tools were developed using an expert-systems approach sometimes called “Good Old-Fashioned AI” or GOFAI.⁴ The algorithm is a complex decision tree, coded to draw from the deep knowledge of human experts. The algorithm is fed information it needs to make decisions that would replicate the decisions of someone with that realm of knowledge.⁵ Any errors, by miscalculation, from bias or otherwise, derive only from the decision instructions and rules the algorithm is given. The GOFAI tools made calculated decisions, but they did not learn. Their knowledge was static, based on the expertise and rules supplied to them.

Conversely, the powerful AI tools and processes of today and the near future have machine learning at their core. The machine learning-driven AI algorithm differs from simple executable algorithms. It does not merely draw from an encyclopedic body of knowledge and execute a task its algorithm instructs it to do. Rather, the ML algorithm is written to cause the tool to train itself from data supplied to it and to learn from its own executions. The learning is iterative: the machine learns first from training datasets to perform its function and usually to make simple decisions of a predictive nature. These decisions follow outcomes in the dataset the tool is fed and then add to the content of the dataset. The machine learns from its previous predictions and activity grows its data. From a large set of data and outcomes, the machine predicts what may be similar to previous choices.

The Implicit Bias Problem: Humans, Data, and Bias

The term “bias” is associated with prejudice, reliance on irrelevant factors, negatively imbalanced outcomes, and often connotes a conscious exercise of moral unfairness. For example, any bias in the output of a GOF AI tool closely derives from bias in its algorithm or the expert knowledge in its encoding. This kind of explicit bias certainly occurs and can infect a machine learning process, but it isn’t the whole story.⁶ Biases can also enter implicitly into ML as humans subconsciously form schemas or put things in categories to manage vast amounts of information. This process may help organize the environment, but it can also have more serious impacts. A variable can come to represent a class that shares certain characteristics; for example, someone may inappropriately ascribe to a single individual a host of other features sometimes associated with that class. Implicit bias can cause superficially or facially neutral factors to produce imbalanced outcomes. This kind of bias can enter implicitly into ML applications and can give rise to subtle but problematic outcomes.⁷

Machine learning algorithms learn from data and coding produced by flawed human systems, data that derives from past bad decisions, and data that derives from systemic and societal injustices.⁸ Sources of implicit bias in AI tools are multifold, and the impact of these sources are not equal. Neither are they all equally well understood. Some are identifiable and, through some work and attention, can be addressed, while others are subtle and more difficult to resolve: bias can enter an ML process insidiously, through pervasive and deeply ingrained systems that govern our societies.

Some literature on AI and libraries points to human coders as a key source of bias. The suggestion is that human biases imbue biased algorithms and are thus embedded into AI tools via those algorithms.⁹ They assert that the source of AI bias is human error or the biased human’s role in writing an algorithm.¹⁰ To be sure, subconscious and conscious biases accompany everything humans do; indeed, early AI proponents cited this human fallibility as a reason to let the machines make the decisions.¹¹ The bulk of the implicit bias is the data from which the algorithm learns and with which it operates. To grasp this requires an understanding of how current AI tools work.

A human may identify a problem, and then other humans proceed to design and code an algorithm to address that problem. The algorithm is trained to solve this by generating output based on an initial set of data. That dataset allows the algorithm to generate an analysis and a prediction of an outcome based on coded or labelled data and statistical probabilities. The outcome of the analysis supplements the dataset, and the algorithm continues to learn and change. The trained algorithm can then be applied to larger datasets, executing predictions from the new data, and continuing to learn or train itself from this data. The process continues iteratively, with more data added and the algorithm continuing to learn and change.¹²

Unsupervised Machine Learning and Supervised Machine Learning

In unsupervised machine learning, the algorithm runs on the dataset and produces its predictions or outputs on its own. There is no human intervention to reduce distortion or to evaluate results during the iterative learning process. This means any problematic results in the outputs can continue to feed the iterative process. If, for example, a dataset used to train a tool to predict challenging library patrons draws from big datasets of national criminal justice data, the tool's output likely will reflect historical societal and racial inequities that do not appropriately predict for the intended question. When this data is used to train a tool, serious implicit bias in that data can go undetected, all while the algorithm is adapting and refining itself from that data. The result is not only a replication but also an amplification of the initial biases in the data. The multiple iterations of the ML process cause the problems with the algorithm to become difficult to even identify, let alone resolve.

Supervised machine learning, on the other hand, incorporates some measure of human intervention and, along with it, a more ethics-driven approach. A person who is familiar with the dataset itself and the problem-solving algorithm, as well as with desired or predicted outcomes, will audit the dataset and the algorithm's initial outputs. They will assess likely accuracy and appropriateness of the data and will select and remove features with a goal to minimize distortions. The person will look for patterns and features the algorithm initially captured and will study the relationship between them and the initial outcome. With reference to an expected outcome, the supervisor will determine which variables to include and which to exclude in the model.

Ethical Foundations of Academic Libraries: Guidance for Addressing AI and Bias Issues

Over the last two decades, as AI has begun to see practical use cases, ethical guidelines seem to have proliferated. Universities, faculty, librarians, professional organizations, companies, governments, international organizations, and civil society all strive to provide frameworks, principles, or statements to guide solutions in their sectors. Overlaps even exist among the multiplicity of individual and localized efforts. Whereas academic libraries do not yet have explicit guidance for addressing ethical issues arising from AI, they can use both their existing framework and guidance from other sectors as resources to guide decisions and to support concrete steps—and perhaps to guide the creation of a statement. Further, given the multitude of actors who interact with academic libraries—faculty, students, vendors, government, the general public, or others—benefits can be drawn from awareness of ethical considerations that might guide other entities.

Academic Libraries and the ACRL Framework

An exploration of ethical guidance for academic libraries must begin with the *Framework for Information Literacy for Higher Education*, adopted in 2016 by the Board of the Association of College & Research Libraries (ACRL).¹³ The ACRL *Framework's* six concepts represent core values for higher education institutions to guide their teaching and researching tasks. The ACRL *Framework* strives to help academic institutions frame their own mission statements, scholar goals, and outcomes under shared concepts and ideas. Its conceptual understandings aim to provide the philosophical stepping stones that can assist academic institutions to develop their own tools, policies, and benchmarks for teaching and research.

Despite no explicit appearance of the word *ethics* in the text of the *Framework*, a few of its six concepts may help academic libraries construct their own ethical considerations. Examples include Authority is Constructed and Contextual, Information Creation as a Process, Information Has Value, and Searching as Strategic Exploration. Taken together, these concepts and their underlying statements set out exactly what the *Framework* was intended to facilitate: an initial set of steps for academic libraries to frame their own approaches to challenges, both current and unforeseen.

The *Framework* encourages academic libraries to challenge the entire information environment. It challenges embedded cultural, social, and political assumptions and the biases and context that underlie data, self-described evidence, research questions, methodologies, and conclusions. As the uses of the *Framework* continue to evolve, academic libraries are prompted to ask questions and pursue answers. The imminent applications of AI in academic libraries call for the *Framework* to be more explicit regarding ethics and bias in AI.

Universities and Ethical AI Imperatives

Efforts in universities are relevant to academic libraries not only for the academic institutional environment but also because university research activities can launch AI tools useful for academic libraries and give guidance for their ethical, bias-free application. Since the early 2010s, university scholars have urged robotics and robotics engineers to develop codes of ethics, and ethical guidance for the initial research stages is a continued imperative.¹⁴ Some universities have created their own recommendations that build upon their own histories as well as current AI thought leadership, and they may serve as models for other universities.¹⁵ Universities also are at the forefront of conversations about AI ethics through research centers and think tanks dedicated to this issue. A few examples highlight the development of spaces for consideration of ethical uses of AI, and they also illustrate initiatives that strive to bring numerous scholars together under the umbrella of multidisciplinary approaches to ethics in AI applications. For example, Stanford University collaborated with the Machine Intelligence Research Institute (MIRI) in 2006. MIRI's conferences and national events set the stage for conversations exploring concepts such as "friendly AI" and "effective altruism." In 2017, the MIT Media Lab and the Harvard Berkman-Klein Center for Internet and Society launched the Ethics and

Governance of AI Initiative. This initiative strives to both build a national network of universities, companies, and civic organizations and to financially support AI research and projects for the public interest. In the same year, New York University founded the AI Now Institute. Through its advocacy, research, symposia, and expert testimony, AI Now has positioned itself as a key actor in a national and international conversation of social implications and applications of AI research. More recently, in 2019, Stanford University launched the Institute for Human-Centered Artificial Intelligence (HAI), intended to be a cross-disciplinary hub for all faculties within Stanford University and other universities and colleges.

AI, Implicit Bias, and Ethical Responsibilities and Opportunities for Academic Libraries

Responsibilities of AI Innovators in Academic Libraries

Component tasks that are the responsibility of AI technologies can serve different types of functions. Some AI functions are in use in libraries today, mainly in predictive or decision-support contexts, and more extended applications in the future are foreseen. By their iterative and cumulative learning nature, the developed operation of machine learning algorithms and the basis for the outcomes they produce are often unknown to even the initial coder.¹⁶ Even if inadvertently, developers are likely creating AI tools that magnify systemic biases and distort outcomes. Unfortunately, many innovators may not be alert to the reality that implicit bias becomes embedded this way. They may not be familiar with the extent to which implicit bias in datasets can drive ML or how readily and regularly they exacerbate historical patterns of discrimination. Innovators using ML algorithm programming skills may build exciting tools to carry out valuable functions in libraries. However, if innovators use weak training data tainted by bias, unsupervised learning processes, or even supervised learning processes without sufficient understanding of variables and appropriate outcomes, they can unwittingly produce outcomes tainted by and perpetuating implicit bias.

Bias Mitigation in Supervised Machine Learning Can Intercede and Minimize Errors

Researchers are working to create methods to mitigate bias.¹⁷ They are able to supply corrective data and decision-making paths to help train AI tools in different ways. They may work to fix datasets and alleviate concerns regarding embedded implicit bias. These

processes may require extensive testing and repeated interventions to give statistically balanced results. They also require knowledgeable individuals who have, in addition to ML expertise, a good understanding of potential problems and expected outcomes. Some working in this area also advocate for regulatory guidance to provide a framework for AI.¹⁸

A related concern is that the power to develop AI tools relates directly to access to, ownership of, or control over big datasets.¹⁹ Most amassed big data that could be useful to academic libraries and information workers are held by a few stakeholders, and the masses are generally unaware of what is in that data or the details of the algorithms. Developers themselves, privately and within academic libraries, may never fully understand the nuances of AI-driven library solutions which, by their self-developing nature, become unknown to either the original human coder or the user. Ethical adoption of AI-driven solutions requires academic libraries to fully evaluate this reality.

A challenge to the development of locally created AI applications for use in academic libraries is likely to be a shortfall in large, representative amounts of high-quality usable data pertinent to the local community and the purpose. Many sources of openly available training datasets exist, but these will only assist if the focus of the data and their presentation are suitable to the desired machine learning goal and serve the project.²⁰ Even then, there are no assurances that the available data have not been drawn from sources that might have been subject to decades of systemic, institutional, and unconscious (or even explicit) bias.

We know that uncertainty exists in our ability to assess the fairness and reliability of AI tool outputs. Users without an ability to break down the results from an AI *black box* may have continued uncertainty about whether a prediction, a recommendation, or a decision is fair or embeds and perpetuates implicit biases. This raises the concern of whether use of the tool is a violation of academic library professional standards, patron respect, or ethical standards.

Supervision allows a human to step back and assess the value of the variables and the propriety of including any variable in the ML process. However, the availability of expertise can be a challenge: an effective supervisor must have broad and deep knowledge of the issue and model or problem-solving algorithm, the range of variables, and likely predicted outcomes. Academic librarians can be ideal partners for developers searching for clean datasets for their models. Because of the contribution of their longstanding and novel expertise, academic libraries are well-positioned to explore, experiment, and work with researchers, scholars, and practitioners to provide meaningful, creative, and ethical solutions to the problems of AI and implicit bias.

Conclusion

As machine learning technologies advance and data sources become more unwieldy and opaque, concerns about embedded and inextricable implicit bias are both real and increasingly widespread. At the same time, computing power and the abundance of datasets further the proliferation of AI applications that can be used in or by academic libraries.

Simultaneously, communities of practice and professions are advancing valuable ethical assessments that can assist in clarification and guidance for academic libraries.

As centers of information, knowledge, and creation, academic libraries can play a pivotal role in clarifying issues with data, AI, and implicit bias. The ethical implications of implicit bias in the creation and use of AI in academic libraries is an ongoing conversation, and a range of disciplines and communities offer valuable ethical guides. These resources can help inform a framework that is current, relevant, and robust in the face of AI challenges. Academic libraries can be increasingly adept at ethically addressing issues of AI and implicit bias for their own work and that of the wider AI development community.

Endnotes

1. David G. Stork, "The Best-Informed Dream: Hal and the Vision of 2001," in *Hal's Legacy: 2001's Computer as Dream and Reality*, ed. David Stork (Cambridge, MA: MIT Press, 1997), 7–15; HAL 9000 was the name of the spaceship computer in the film and book. Interestingly, the fictional HAL 9000 of 1968 bore the power—and ultimately the weight—of ethical decision-making. HAL 9000 is often seen as the antagonist or villain of the story, as he determines the way to adhere to his programmed rules is to kill his human space travel companions; see Daniel C. Dennett, "When HAL Kills, Who's to Blame? Computer Ethics," in *Hal's Legacy: 2001's Computer as Dream and Reality*, ed. David Stork (Cambridge, MA: MIT Press, 1997), 203–14.
2. Elizabeth A. Waraksa, ed., "Research Library Issues. RLI 299: Ethics of Artificial Intelligence," Association of Research Libraries (2019): 5.
3. AI of fiction and dystopian futures is often generalized AI, whose development is not yet within modern technological reach.
4. Charles Morgan, "Responsible AI: A Global Policy Framework," *ITechLaw* (November 18, 2019): 20, <https://www.itechlaw.org/ResponsibleAI>.
5. Morgan, "Responsible AI," 20–21; text accompanying endnotes, 24–29.
6. Ayanna Howard and Jason Borenstein, "The Ugly Truth About Ourselves and Our Robot Creations: The Problem of Bias and Social Inequity," *Science and Engineering Ethics* 24, no. 5 (2018): 1521–36, <https://doi.org/10.1007/s11948-017-9975-2>.
7. And sometimes biases such as stereotyping produce a nefarious circumstance where a decision-maker renders a biased judgment when fully aware of and intending to exercise that bias. This kind of explicit bias is detrimental and no doubt occurs, although it is not the focus of this chapter. For a thorough study of this, see Safiya Umoja Noble, *Algorithms of Oppression: How Search Engines Reinforce Racism* (New York, NYU Press, 2018), <http://ebookcentral.proquest.com/lib/cornell/detail.action?docID=4834260>.
8. Noble, *Algorithms of Oppression*; a good encapsulation of the problem is presented by Joy Buolamwini, founder of the Algorithmic Justice League: Joy Buolamwini, "PBS NewsHour: America in Black and Blue 2020: PBS NewsHour Weekend Special," WCNY Video, 2020, <https://video.wcny.org/video/america-in-black-blue-2020-a-pbs-newshour-weekend-special-4lpnmp/>; see also Cathy O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (New York: Broadway Books, 2016), <http://search.ebscohost.com/login.aspx?direct=true&db=nlebk&AN=1109940&site=ehost-live>.
9. See, e.g., Ellyssa Kroski and Robert J. Ambrogi, "The Future of AI in Law Libraries," in *Law Librarianship in the Age of AI* (Chicago, IL: ALA Editions, 2020).
10. See, e.g., Jamie J. Baker and Robert J. Ambrogi, "AI in Legal Research," in *Law Librarianship in the Age of AI* (Chicago, IL: ALA Editions, 2020).
11. See, e.g., John McCarthy, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon, "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955," *AI Magazine* 27, no. 4 (2006): 12, <http://www-formal.stanford.edu/jmc/history/dartmouth.pdf>.
12. A human, no doubt, usually will select an initial dataset and be involved in labelling but usually does not have control of each data point.

13. *Framework for Information Literacy for Higher Education*, American Library Association, February 9, 2015, accessed September 16, 2020, <http://www.ala.org/acrl/standards/ilframework>.
14. Illah Reza Nourbakhsh, *Robot Futures* (Cambridge, MA: MIT Press, 2013); Brandon Ingram, Daniel Jones, Andrew Lewis, Matthew Richards, Charles Rich, and Lance Schachterle, “A Code of Ethics for Robotics Engineers,” in *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (IEEE, 2010), 103–104.
15. Julie A. Shah and Melissa Nobles, “Working Group on Social Implications and Responsibilities of Computing Final Report” (Cambridge, MA: MIT Schwarzman College of Computing Task Force, 2019), 1–22, <http://web.mit.edu/comptfreport/sirc.pdf>.
16. Frank Pasquale, *The Black Box Society: The Secret Algorithms that Control Money and Information* (Cambridge: Harvard University Press, 2015), 320, <https://doi.org/10.4159/harvard.9780674736061>; the issue in this context is summarized by Ellyssa Kroski and James M. Donovan, “Benefits, Drawbacks, and Risks of AI,” in *Law Librarianship in the Age of AI* (Chicago, IL: ALA Editions, 2020).
17. “AI Fairness 360,” IBM Research Trusted AI, 2020, <http://aif360.mybluemix.net/>.
18. Laila Paszti, “Did the OPC Get It Right? A Proposed Approach for Regulating Artificial Intelligence,” Canada | Global law firm | Norton Rose Fulbright, 2020, <https://www.nortonrosefulbright.com/en-ca/knowledge/publications/33076187/did-the-opc-get-it-right-a-proposed-approach-for-regulating-artificial-intelligence>.
19. Pasquale, *The Black Box Society*.
20. Joaquin Vanschoren, “Democratizing Machine Learning,” OpenML, 2020, <https://www.openml.org/>; “Public Datasets | Google Cloud,” Google Cloud Public Datasets, Google (Google, 2020), <https://cloud.google.com/public-datasets/>; “Open Datasets: Microsoft Azure,” Microsoft Azure, 2020, <https://azure.microsoft.com/en-us/services/open-datasets/>.

Bibliography

- American Library Association. *Framework for Information Literacy for Higher Education*. February 9, 2015. Accessed September 16, 2020. <http://www.ala.org/acrl/standards/ilframework>.
- Baker, Jamie J., and Robert J. Ambrogi. “AI in Legal Research.” In *Law Librarianship in the Age of AI*. Chicago, IL: ALA Editions, 2020.
- Buolamwini, Joy. “PBS NewsHour: America in Black and Blue 2020: PBS NewsHour Weekend Special.” WCNY Video. 2020. <https://video.wcny.org/video/america-in-black-blue-2020-a-pbs-newshour-weekend-special-4lpnmp/>.
- Dennett, Daniel C. “When HAL Kills, Who’s to Blame? Computer Ethics.” In *Hal’s Legacy: 2001’s Computer as Dream and Reality*, edited by David Stork, 203–14. Cambridge, MA: MIT Press, 1997.
- Google. “Public Datasets | Google Cloud.” Google Cloud Public Datasets. 2020. <https://cloud.google.com/public-datasets/>.
- Howard, Ayanna, and Jason Borenstein. “The Ugly Truth About Ourselves and Our Robot Creations: The Problem of Bias and Social Inequity.” *Science and Engineering Ethics* 24, no. 5 (2018): 1521–36. <https://doi.org/10.1007/s11948-017-9975-2>.
- IBM Research Trusted AI. “AI Fairness 360.” 2020. <http://aif360.mybluemix.net/>.
- Ingram, Brandon, Daniel Jones, Andrew Lewis, Matthew Richards, Charles Rich, and Lance Schachterle. “A Code of Ethics for Robotics Engineers.” In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2010, 103–104.
- Kroski, Ellyssa, and Robert J. Ambrogi. “The Future of AI in Law Libraries.” In *Law Librarianship in the Age of AI*. Chicago, IL: ALA Editions, 2020.
- Kroski, Ellyssa, and James M. Donovan. “Benefits, Drawbacks, and Risks of AI.” In *Law Librarianship in the Age of AI*. Chicago, IL: ALA Editions, 2020.
- McCarthy, John, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon. “A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955.” *AI Magazine* 27, no. 4 (2006): 12. <http://www-formal.stanford.edu/jmc/history/dartmouth.pdf>.
- Microsoft Azure. “Open Datasets: Microsoft Azure.” 2020. <https://azure.microsoft.com/en-us/services/open-datasets/>.
- Morgan, Charles. “Responsible AI: A Global Policy Framework.” *ITechLaw* (November 18, 2019): 20. <https://www.itechlaw.org/ResponsibleAI>.

- Noble, Safiya Umoja. *Algorithms of Oppression: How Search Engines Reinforce Racism* (New York, NYU Press, 2018). <http://ebookcentral.proquest.com/lib/cornell/detail.action?docID=4834260>.
- Nourbakhsh, Illah Reza. *Robot Futures*. Cambridge, MA: MIT Press, 2013.
- O’Neil, Cathy. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Broadway Books, 2016. <http://search.ebscohost.com/login.aspx?direct=true&db=nlebk&AN=1109940&site=ehost-live>.
- Pasquale, Frank. *The Black Box Society: The Secret Algorithms that Control Money and Information*. Cambridge: Harvard University Press, 2015, 320. <https://doi.org/10.4159/harvard.9780674736061>.
- Paszi, Laila. “Did the OPC Get It Right? A Proposed Approach for Regulating Artificial Intelligence.” Canada | Global law firm | Norton Rose Fulbright. 2020. <https://www.nortonrosefulbright.com/en-ca/knowledge/publications/33076187/did-the-opc-get-it-right-a-proposed-approach-for-regulating-artificial-intelligence>.
- Shah, Julie A., and Melissa Nobles. “Working Group on Social Implications and Responsibilities of Computing Final Report.” Cambridge, MA: MIT Schwarzman College of Computing Task Force, 2019, 1–22. <http://web.mit.edu/comptfreport/sirc.pdf>.
- Stork, David G. “The Best-Informed Dream: Hal and the Vision of 2001.” *Hal’s Legacy: 2001’s Computer as Dream and Reality*, edited by David Stork, 7–15. Cambridge, MA: MIT Press, 1997.
- Vanschoren, Joaquin. “Democratizing Machine Learning.” OpenML. 2020. <https://www.openml.org/>.
- Waraksa, Elizabeth A., ed. “Research Library Issues. RLI 299: Ethics of Artificial Intelligence.” Association of Research Libraries (2019): 5.