

CIRCADIAN COMPUTING: SENSING AND STABILIZING BIOLOGICAL RHYTHMS

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

Saeed Abdullah

August 2017

This document is in the public domain.

CIRCADIAN COMPUTING: SENSING AND STABILIZING BIOLOGICAL
RHYTHMS

Saeed Abdullah, Ph.D.

Cornell University 2017

This dissertation lays the groundwork for *Circadian Computing* with a novel and broad vision of technologies that support and adapt to our innate biological rhythms. Similar to most terrestrial organisms, human physiology and behavior are shaped by a 24-hour periodicity known as circadian rhythm. Indeed, almost every neurobehavioral process including our sleep, metabolism, cognitive performance, and mood reflects circadian rhythms. These rhythms ensure synchronization across different processes and as such, are crucial for our health and well-being. Persistent circadian disruption increases risk for cancer, obesity, and cardiovascular diseases. It has been associated with occupational accidents and serious loss of productivity in the workplace as well. Recent findings have also started identifying links between circadian disruption and mental illnesses including bipolar disorder and schizophrenia.

However, in our modern world, circadian disruption is becoming increasingly widespread. The invention of artificial light fundamentally changed our ancestral sleep and wakeup patterns. Since then, we have gradually moved towards a 24-hour society. The recent development in entertainment and communication technologies has also resulted in an “always-on” ethos. The resulting trend is worrisome. Sleep pathologies are reaching an epidemic level with 70% of the population suffering from significant circadian disruptions.

As a result, recently there has been an increased focus on monitoring and

identifying disruptions in circadian rhythms. However, these methods and findings are often limited to controlled lab environments. As a result, they are not adequate for granular monitoring of circadian disruptions in the wild over a longitudinal period of time. As such, there is a need for novel pervasive technologies for tracking, monitoring, and modeling circadian disruptions and its impact in the real world. There is also an opportunity for developing intervention tools for maintaining circadian stability.

This dissertation is a leading step towards the broad vision of circadian-aware technologies for sensing, adapting to, and stabilizing our innate biological rhythms. In my PhD work, I have shown the feasibility of bringing a circadian-aware perspective across different application domains. Specifically, I have developed and evaluated methods for unobtrusively assessing circadian disruptions. I have also showed that behavioral and contextual data can be used for modeling and predicting alertness — a circadian process integral to our cognitive performance. I have also developed, deployed, and evaluated a data driven tool focusing on identifying circadian anomalies in patients with bipolar disorder.

With this groundwork in place, I believe that there is an exciting opportunity lying ahead for Circadian Computing. In particular, circadian-aware technologies can potentially reshape a number of application domains including education and learning, optimized scheduling, mental health care, and chronotherapy. I hope this dissertation motivates a circadian perspective in future technology development and contributes to the shared effort of improving our productivity, health, and well-being.

BIOGRAPHICAL SKETCH

Saeed Abdullah grew up in Bangladesh. He received his Bachelor of Science in Computer Science and Engineering from Bangladesh University of Engineering and Technology (BUET). His undergraduate thesis explored encoding schemes for representing artificial neural networks. Upon graduation, he worked as a software developer focusing on J2ME system development. He also participated in Google Summer of Code working on a Just-In-Time (JIT) compiler for Java.

Saeed then moved to University of Vermont as a graduate student. He worked with Xindong Wu focusing on Data Mining and Information Cascade. He received his Master of Science from the Computer Science department. Afterwards, he joined the PhD program in Information Science at Cornell University.

Throughout his PhD, Saeed developed novel data-driven technologies to improve health and well-being. His research is inherently interdisciplinary and he has collaborated with psychologists, psychiatrists, and behavioral scientists. His work has introduced assessment and intervention tools across a number of health related domains including sleep, cognitive performance, bipolar disorder, and schizophrenia. Going forward, Saeed aims to continue his work on developing novel technologies that will improve our productivity, health, and well-being.

ACKNOWLEDGEMENTS

There have been many individuals who provided support and motivation throughout my PhD. I thank each and every one of them. I am particularly appreciative of my committee who guided me in this journey. Tanzeem Choudhury, my chairperson, has been an invaluable resource in shaping my trajectory of intellectual contributions over the years. I will always appreciate how flexible and encouraging she has been in letting me pursue my own research interests. I have been fortunate to have Geri Gay as my committee member. She has been very supportive throughout my PhD. I will miss her guidance and wonderful sense of humor. My other committee member, Deborah Estrin, has provided valuable feedback and motivation to address issues with real-world implications.

Throughout my PhD, I have collaborated with a number of individuals with varied expertise and background. I am truly grateful to have these opportunities. My collaborators have shaped my research identity and greatly influenced this dissertation. I specifically want to thank Mark Matthews and Elizabeth Murnane. Mark has been a great friend and mentor. My PhD work would have been very different without his support, guidance, and encouragement. Elizabeth, as a collaborator, has been an invaluable resource. In particular, her enlightening feedback and support has helped me in every stage of my PhD.

I am also thankful to have the continued support from the Department of Information Science and Cornell University. The department has provided a caring, inclusive, and inspiring environment for me to grow as a young scholar. I have learned a lot from Steve Jackson about fostering and sustaining a close-knit community. I am also grateful to have received guidance and feedback from Dan Cosley on numerous occasions. My lab mates and fellow PhD students

have also been influential in shaping my research contributions in varied ways.

While living in Ithaca for six years, I have been blessed with some great friends. I particularly want to thank Elizabeth for her kindness and love. She was always there for me. Ranjit Singh has been a continuous source of fun and support. Samir Passi and Ritwick Ghosh have been a great help during my PhD as well. I also want to thank the members of Cornell Badminton Club for their support.

Finally, I want to thank my parents, sister, and family for their trust and help. Without their sacrifices and support, this journey would not have been possible. Thank you.

TABLE OF CONTENTS

Biographical Sketch	iii
Acknowledgements	iv
Table of Contents	vi
List of Tables	viii
List of Figures	ix
1 Introduction	1
1.1 Motivation	3
1.1.1 Research Questions and Contributions	6
1.2 Dissertation Overview	9
2 Background	11
2.1 Biological Rhythms	12
2.1.1 Sleep and Wakefulness Cycle	15
2.2 Circadian Disruption	16
2.2.1 Impact on Health	17
2.2.2 Impact on Cognitive Performance	18
2.2.3 Impact on Mental Illnesses	19
2.3 Assessment of Chronotype and Circadian Disruption	20
2.3.1 Physiological Marker Based Assessment	21
2.3.2 Survey Based Assessment	23
2.3.3 Ubiquitous Technology Based Assessment	25
2.4 Summary	30
3 Passive Sensing of Chronotype and Circadian Disruptions	31
3.1 Introduction	31
3.2 Methods	33
3.2.1 Data Collection	34
3.2.2 Sleep Algorithm	35
3.3 Findings	37
3.3.1 Sleep Algorithm Accuracy	37
3.3.2 Assessing Sleep Debt	39
3.3.3 Social jet lag	41
3.3.4 Inferring Sleep Inertia (SI)	43
3.4 Discussion	44
3.4.1 Limitations	45
3.4.2 Implications and Opportunities	48
3.5 Summary	51

4	Predicting Rhythms of Alertness	53
4.1	Introduction	53
4.2	Methods	56
4.2.1	Participants	56
4.2.2	Data	57
4.2.3	Momentary Assessments	57
4.2.4	Sleep and Chronotype Information	60
4.2.5	Phone Usage Data	61
4.3	Findings	61
4.3.1	Replicating and Extending Extant Lab-Based Findings	62
4.3.2	Passive assessment of alertness	70
4.4	Discussion	75
4.4.1	Limitations	76
4.4.2	Implications and Opportunities	78
4.5	Summary	81
5	Circadian Disruption and Bipolar Disorder	83
5.1	Introduction	83
5.2	Method	87
5.2.1	Participants	87
5.2.2	MoodRhythm	88
5.3	Findings	91
5.4	Discussion	95
5.4.1	Implications and Opportunities	95
5.4.2	Limitations	99
5.5	Summary	100
6	Discussion	102
6.1	Summary of Prior Chapters	102
6.2	Considerations	105
6.2.1	The Contradictory Roles of Technology	106
6.3	Future Opportunities	108
6.3.1	Novel Sensing for Granular Assessment	109
6.3.2	Fixing a Broken Clock: Sleep and Circadian Interventions	110
6.3.3	Optimizing Cognitive and Physical Performance	114
6.3.4	Clinical Applications	119
6.4	Concluding Remarks	122
	Bibliography	125

LIST OF TABLES

1.1	A summary of research questions and studies I have conducted to address them	7
2.1	Existing methods for assessing circadian phase markers and disruptions	20
3.1	Participant Information	34
3.2	Sleep midpoint across different phases.	41
4.1	Groupings of stimulants based on how they impact alertness . .	68
4.2	RRT differs significantly between self-assessed high and low alertness, indicating fatigued individuals are usually aware of their reduced capability.	69
4.3	Top ranking feature groups for modeling response time and RRT based on relative feature elimination (RFE)	74
5.1	Demographic of the patients	88
5.2	Correlation between sensor stream and trend of self-assessed energy scores computed as rolling average over 7 days (** : $p < .01$, *** : $p < .001$)	91
5.3	Ranking of feature importance for the stable and unstable status classification using recursive feature elimination (RFE). It also shows weights assigned to features in a support vector machine using a linear kernel.	93

LIST OF FIGURES

2.1	Sleep onset and duration depends on our circadian system, homeostatic sleep pressure, and social behavior.	15
3.1	Average sleep onset and duration across participants from phone and journal data. Phone non-usage coincides with sleep events; the trend is more stable on weekdays due to more data points. For most participants, sleep onset is delayed and duration is longer during the weekends while participant 4 — with an early chronotype — gets less sleep on weekends.	38
3.2	Average sleep duration across participants with 95% confidence interval	39
3.3	Duration of average social jet lag compared across chronotypes. 95% confidence interval has also been shown. Note that we have only one participant with chronotype 03:00, so interval estimation is set to zero in that case.	42
3.4	Inferred sleep inertia duration (with 95% CI) compared across early (N=1) and late chronotypes (N=8). Difference in sleep inertia duration from weekdays to weekends reflects the patterns of accumulated sleep debt across different chronotypes.	43
4.1	Relative response time (RRT) with standard error of mean (SEM) over the day. Positive and negative values of RRT indicate alertness higher and lower than individual baseline, respectively. . .	62
4.2	RRT of early chronotypes compared to late chronotypes across the day. Blue and red colors indicate higher RRT for early and late types, respectively. In the morning, early chronotypes display much higher alertness than late types, while the opposite is observed later in the day. Response time difference between early and late types across the day is also statistically significant: $F(5, 964) = 5.36, p < 0.001$	64
4.3	RRT with standard error of mean (SEM) with internal time on X-axis	65
4.4	RRT change for early and late chronotypes before and after Spring DST. While the DST transition negatively affects both early and late types, late types suffer more. The difference in RRT before and after DST for both early ($t = -2.37, p = 0.02$) and late ($t = -2.52, p = 0.01$) types is statistically significant.	66
5.1	Sample paper-based Social Rhythm Metric (SRM) form that is used as part of Interpersonal Social Rhythm Therapy (IPSRT). . .	85
5.2	MoodRhythm application used by the patients with bipolar disorder.	89

CHAPTER 1

INTRODUCTION

Whatever physiological variables we measure, we usually find that there is a maximum value at one time of day and a minimum value at another.

Jürgen Aschoff

Time is not just a modern social construct. Temporal structures are etched into every living creature. Our biological processes vary considerably and predictably over time. These biological rhythms are a consequence of living on a rotating planet. Indeed, throughout evolution, the environment has been ever-changing but with one constant: the earth turns on its axis every 24 hours resulting in periodic changes in light and warmth. As a result, we have evolved with biological clocks that can anticipate and exploit such periodic changes.

In particular, almost every biological process shows circadian¹ rhythms with a roughly 24 hour period. Our sleep and waking behaviors perhaps are the most obvious examples of such circadian processes. These rhythms enforce a “programmed regularity” [120] and ensure synchronization across different processes that are critical to our health and well-being. As a result, recently there has been an increased focus on monitoring and identifying disruptions in circadian rhythms. However, these recent methods and findings are often limited to controlled lab environments leaving a gap in knowledge and applicability when it comes to real-world settings.

Ubiquitous technologies can help to address this gap. Specifically, the rise of phones and wearables around the globe provides an unprecedented opportu-

¹The word circadian comes from *circa* (approximately) and *diem* (a day).

nity in data collection, modeling, and appropriate interventions. Furthermore, a circadian-centric perspective in technology development can fundamentally reshape our daily interactions with technologies. Imagine your devices and smart-home environment nudging you towards an optimal sleep schedule so that you do not suffer from jet lag after your upcoming travel next week or your calendar scheduling daily tasks with appropriate difficulties given your personalized rhythms of cognitive performance. Indeed, circadian-aware technologies can fundamentally reshape the domain of ubiquitous and personalized computation.

However, there is much to be done to make these possibilities a reality. Towards this goal, in this dissertation, I will introduce the concept of *Circadian Computing* — technologies that support and adapt to our innate biological rhythms. In concrete terms, the main goals of Circadian Computing are three-fold: i) developing low-cost and scalable methods for data collection that can complement and extend lab based sleep and circadian rhythms research findings, ii) modeling and predicting in-situ changes in circadian processes based on passively sensed data from real-world setups, iii) designing and developing interventions focusing on overall stability and optimizing performance (both cognitive and physical) *in accordance* with our circadian rhythms.

This dissertation will not be able to cover every aspect of Circadian Computing given its broad vision and goals. Instead, I have strived to lay the groundwork here through the development of novel methods and findings with a focus on showing the feasibility of bringing a circadian-aware perspective across different application domains. The leitmotif of this dissertation, thus, revolves around identifying limitations of existing tools given a particular context, de-

veloping a novel technical solution grounded in theories from chronobiology to address those limitations, and then deploying it in the real-world to evaluate its performance.

1.1 Motivation

Like every other terrestrial organism, human physiology follows circadian rhythms. Indeed, almost every neurobehavioral process ebbs and flows over a 24 hour period. Our body clocks also vary between individuals, from proverbial “early birds” (early types) to “night owls” (late types). These individual differences can significantly impact the timing of our biological processes. For example, the peak alertness or the optimal timing of sleep considerably varies from one person to another. In other words, when it comes to our physiology, there are two different dimensions of time — *internal* time and *external* (social) time.

When our body clocks are disrupted, it results in misalignment between internal and external time. For example, a late type being forced to wake up early will result in a misalignment between her internal and external time. Similarly, shift-work can also induce circadian misalignments. Persistent circadian disruptions can have devastating effects on our health and overall well-being. For example, circadian disruption over a long period of time has been associated with breast cancer [142]. Kubo et al. [122] also found that men doing rotating shift-work — which result in considerable misalignments — are three times more likely to develop prostate cancer. Indeed, based on existing evidence, the International Agency for Research on Cancer (IARC) has concluded that “shift-

work that involves circadian disruption is probably carcinogenic to humans” [215].

Circadian disruption has been associated with a range of mental illnesses including alcohol and substance abuses [192, 96], anxiety disorder [171], schizophrenia[241], and bipolar disorder [94]. In particular, the role of circadian disruption in schizophrenia and bipolar disorder is well-established. Sleep and circadian disruptions are one of the most common symptoms for patients with schizophrenia [173]. Similarly, circadian instability can also trigger relapse onset in patients with bipolar disorder [94]. A number of recent studies have also linked schizophrenia [26] and bipolar disorder [22, 134] with “clock genes” (genes involved in generation of circadian rhythms).

Our metabolic processes follow circadian rhythms as well. As a result, persistent circadian disruptions can have adverse impact on our metabolic systems. A number of studies have associated circadian misalignment with obesity and diabetes [200, 196]. Similarly, it can also increase the risk of cardiovascular diseases significantly [116]. For younger generation, circadian disruption can increase the risk of drug and alcohol use [238, 217]. It can also result in cognitive impairments and learning deficits [37].

However, in our modern world, circadian disruption is getting increasingly widespread. The invention of artificial light fundamentally changed our ancestral sleep and wakeup patterns [61]. Since then, we have gradually moved towards a 24-hour society. Indeed, around 15–20% of total workforce in most industrial societies are shift-workers [27]. The recent development in entertainment and communication technologies has also resulted in an “always-on” ethos, which has potentially exacerbated the situation. Based on a large scale

study, Roenneberg et al. [196] found that 70% of the population suffer from significant circadian disruptions. Sleep pathologies — often indicative of circadian misalignments — are reaching an epidemic level affecting 50–70 million people in the USA alone [164]. The Centers for Disease Control and Prevention (CDC) declared insufficient sleep to be a public health problem in 2015 [72]. Sleep and circadian issues also result in enormous economic burden. Hafner et al. [91] reported that the yearly economic loss caused by insufficient sleep amounts to a staggering \$411 billion in the USA.

Given this wide-ranging adverse effects of sleep and circadian disruptions on well-being and productivity, there has been an increased focus on monitoring and better understanding effects of circadian misalignments. While this has led to significant advances in untangling the biological underpinning of circadian disruptions, these studies are often done in artificial setups (e.g., keeping participants in artificial light-dark cycles where they are required to periodically provide blood and saliva samples). These methods are, understandably, not scalable to a large population. To address these issues, researchers have used subjective assessments and surveys including Munich ChronoType Questionnaire (MCTQ) [238]. However, these assessment methods introduce a different set of challenges regarding data quality and coverage. In particular, subjective assessment methods are difficult to deploy for a longitudinal study given the user burden. Furthermore, these methods are also not ideal when it comes to granular tracking of instantaneous changes. As such, chronobiologists have repeatedly point out the need for new methodologies that will allow *in-situ* and real-time data collection over a long period of time, across various time-zones and geographical locations [194].

Pervasive and ubiquitous technologies can help to address these issues. Phones and wearables are getting computationally more powerful with added sensing capabilities. As a result, these devices can provide new ways of in-situ monitoring of circadian disruptions. Furthermore, they have been adopted around the world, which can help to address issues related to large scale deployments. Indeed, around 3.9 billion people have access to phones today and this number is expected to rise to 6.8 billion by 2022 [63]. Widespread availability of these sensor-rich platforms opens up a compelling opportunity to not only bridge the gap between lab based chronobiological findings and real-world implications but also further untangle the relationships between circadian rhythms and behavioral cues. This dissertation is a leading step towards this vision.

Specifically, in this dissertation, I will introduce the concept of *Circadian Computing*. The goal of Circadian Computing is to design and develop pervasive and ubiquitous technologies that can understand and adapt to one’s underlying biological rhythms. For example, can we passively monitor and complement the ebb and flow of one’s alertness — a circadian process critical to our cognitive performance? Circadian Computing, thus, moves beyond the incomplete assumptions about steady capabilities and fixed requirements of its users throughout the day. In doing so, it significantly broadens the scope and focus of ubiquitous and personalized computing.

1.1.1 Research Questions and Contributions

With this broad vision in mind, throughout my PhD, I have pursued a number of research questions focusing on advancing the current state of technology

design and development. In this dissertation, I will particularly focus on the research questions summarized in Table 1.1. I believe these research questions adequately represent my overall contributions and also provide the groundwork for next steps of Circadian Computing.

#	Research Question	How The Dissertation Addressed It
1	Can we develop a scalable method for passively assessing and monitoring circadian disruptions over a long period of time?	Development of a phone usage based algorithm that can accurately infer circadian disruptions (Chapter 3)
2	How can we model and predict the rhythms of a circadian process in the real-world setup?	Development and deployment of an unobtrusive method for assessing alertness — a circadian process critical for cognitive performances (Chapter 4)
3	How can we design and develop tools that can help to maintain circadian stability?	Design, development and deployment of MoodRhythm — a phone based application for inferring stability in bipolar disorder (Chapter 5).

Table 1.1: A summary of research questions and studies I have conducted to address them

Contributions

To answer these research questions, I have undertaken a number of user studies. They involved designing and developing of novel technologies as well as evaluating their performance in real-world setups by deploying over appropriate study populations. These studies have resulted in novel artifacts, methodology, and empirical findings. In particular, these methodologies and findings bring a circadian-aware perspective to a number of different application domains. These studies, thus, provide the basis of Circadian Computing and also identify its potential scope and opportunities.

Specifically, in this dissertation, I have made the following contributions:

- I have designed and developed a phone usage based method to unobtrusively assess sleep and circadian disruptions (chapter 4). This method has been used and replicated by others in a large scale study [48].
- I have developed a phone based data collection framework for monitoring patterns of alertness — a circadian process crucial to our cognitive performance. Based on these data, I showed that a number of factors including time of the day, body clock type, sleep, and stimulant intake can significantly influence alertness. Furthermore, I have developed and validated a predictive model for inferring alertness using behavioral and contextual features (chapter 5).
- Focusing on maintaining circadian stability in mental health, I have developed MoodRhythm — a phone application that uses passive sensing data for continuous and unobtrusive monitoring of behavioral and contextual features that might indicate circadian disruptions in bipolar disorder (chapter 5).

Most of the work as presented in this dissertation has been published in a number of prominent conferences and journals. For example, I have published multiple papers in the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp) [3, 159], the International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI) [139, 160] (along with *a best paper award* for [160]), IEEE Computer [137], Assessment [138], and the Journal of the American Medical Informatics Association (JAMIA) [2]. I have also co-authored a book chapter on Circadian Computing [4].

1.2 Dissertation Overview

I have organized this dissertation into the following six chapters.

In chapter 1 (Introduction), I have provided the high-level motivation behind Circadian Computing. The chapter also details the research questions I have pursued and the specific contributions I have made throughout my PhD. Furthermore, I have explained how these contributions provide the basis for next steps in Circadian Computing.

In chapter 2 (Background), I will provide the theoretical underpinnings of my research. Drawing from chronobiology, I will introduce new vocabulary, definitions, and terminologies that will be used in the remainder of the dissertation. This chapter also details the current findings regarding the critical importance of circadian stability for our overall health and well-being. I will summarize the existing tools used for sensing circadian misalignments and identify their limitations as well.

In chapter 3 (Passive Sensing of Chronotype and Circadian Disruptions), I will describe a novel method for unobtrusively assessing chronotype and circadian disruptions using phone usage data. I will evaluate the performance of this method by using data from a study spanning 97 days over 9 participants. I will also explain how these method and findings broaden the scope of Circadian Computing.

In chapter 4 (Predicting Rhythms of Alertness), I will focus on alertness — a complex biological process that reflects circadian rhythms. For this, I will first detail a phone based data collection framework for monitoring patterns of

alertness in the real-world setup. Based on this data, I will then present findings about factors that influence alertness. Further, I will develop and evaluate models to predict alertness using behavioral and contextual data. I will also detail how being able to predict alertness can have significant impact across different application domains including occupational safety.

In chapter 5 (Circadian Disruption and Bipolar Disorder), I will describe how maintaining circadian stability is crucial for illness management of patients with bipolar disorder. I will also identify how existing clinical tools are inadequate for monitoring circadian stability of the patients over a long period of time. I will then describe MoodRhythm — a phone application focusing on overcoming the limitations of these existing tools. I will also explain how the sensing data collected by MoodRhythm can be used for automated and unobtrusive monitoring of circadian stability of the patients.

Chapter 6 (Discussion) is the concluding chapter. Here, I will revisit some of the key assumptions I have made throughout this dissertation as well as summarizing the contributions. I will also identify opportunities for coming up with new methodologies for better monitoring of circadian factors (e.g., by taking light exposure into account). I will conclude by identifying future opportunities with a specific focus on interventions grounded in theories from chronobiology. I will also describe potential applications of circadian-aware technologies in a number of new domains, thus further broadening the scope of Circadian Computing.

CHAPTER 2

BACKGROUND

In this chapter, I will provide the background for Circadian Computing. My vision for Circadian Computing is highly interdisciplinary — my PhD work has drawn from chronobiology¹, sleep research, mental health care, ubiquitous computing, and mobile sensing. As such, it is not feasible to provide an exhaustive review of related works from all these domains. Instead, my intention here is to set up the appropriate context that motivated this dissertation.

Towards this goal, I will start with a brief history of research into sleep and circadian rhythms. I will also point out that circadian rhythms provide an evolutionary advantage by enabling organisms to predict and adapt to periodical changes in external environment. I will then describe the endogenous and self-sustaining nature of our circadian system and how it uses external cues to get synchronized with the environment. Circadian rhythms are crucial for our cognitive and physical functioning. I will detail how circadian disruption can negatively impact our overall health and well-being. Specifically, I will review the current evidence regarding the association between circadian disruption and increased risk of cancer, diabetes, and obesity. I will also point out how sleep and circadian disruption is often the most common symptoms of a number of mental illnesses including schizophrenia.

As such, continuous assessment of circadian disruption is important for both individual well-being and public health. In this chapter, I will review the existing methods for assessing circadian markers and disruptions. I will group these assessment methods into three broad categories: i) physiological marker based

¹Chronobiology is an interdisciplinary field that studies biological rhythms.

assessment (e.g., melatonin from blood samples), ii) survey based assessment (e.g., Munich ChronoType Questionnaire or MCTQ), and iii) ubiquitous technology based assessment (e.g., actigraphy and phone). These existing methods are often inadequate when it comes to continuous and real-time assessment of circadian disruptions over a longitudinal period of time across a large population. Circadian Computing aims to bridge this gap by providing scalable sensing methods that can collect granular and real-time data about our circadian processes and disruptions. It also aims to provide effective intervention methods that will minimize circadian disruption.

2.1 Biological Rhythms

The earth's rotation around its axis results in geophysical cycles including daily, seasonal, and annual rhythms. The changes across these geophysical cycles provide temporal cues crucial for our biological survival and success [73]. In particular, nearly every organism takes advantage of the predictable environmental changes (e.g., light and warmth) resulting from daily cycles. Organisms maintain circadian clocks to anticipate periodic changes in their external environments. These circadian clocks provide an evolutionary advantage by enabling better adaptation to the periodic oscillations of external resources [172].

The association of biological processes and 24-hour periodicity has a long history. In 1729, de Mairan reported the circadian rhythms in leaf movement [52]. In this experiment, he kept a plant in constant darkness and found that even without external stimuli its leaf movement showed daily periodicity. This reflects the endogenous nature of circadian clocks. In last few decades, chronobiologists have identified circadian processes in almost every terrestrial organ-

ism.

Research on understanding circadian factors in human biology and behavior started relatively recently. Nathaniel Kleitman [115] (“the father of modern sleep research”) looked into the effects of living on a 21 hour and 28 hour schedules. In 1970s, Jürgen Aschoff [12] investigated how sleep and other biological processes changed if a participant was kept in constant condition and isolated from temporal cues. For this, the scientists built a bunker into a hill in Andechs as an isolation facility. Since then, a number of ground-breaking studies have helped us understand the biological clocks in humans.

These clocks result in endogenous and self-sustaining biological rhythms. That is, even without external stimuli (e.g., when kept in constant light or dark conditions), our biological processes can sustain a period of approximately 24 hours. This period length when our circadian system is not influenced by any external temporal cues is known as *free-running period* (or *tau*, τ). The free-running period for most human is slightly longer than 24 hours — Czeisler et al. [50] reported it to be 24.18 hours on average.

However, we do not live in constant light or dark conditions. Our circadian system is usually not free-running as it synchronizes with external environment using temporal cues. This process of synchronization with external time is known as *entrainment*². The environmental cues for entrainment are known as *zeitgeber*³. Light is the most important zeitgeber [51]. Other environmental factors can also act as zeitgebers. For example, meal timing [234] and exercise [239] can shift our circadian phase.

²The word entrainment came from the french term *entraîner* [227]. It means carrying over.

³Zeitgeber is a German word, which means time-giver (zeit: time, geber: giver).

Our circadian system is a hierarchical one. Almost every cell in mammals contains self-sustained circadian clocks [54]. Most peripheral tissues and organs also have their own circadian oscillators (*peripheral clocks*) [150]. All these clocks get synchronized by a central pacemaker — the *Suprachiasmatic Nucleus* (SCN) located in hypothalamus [54]. In other words, SCN acts “like the conductor of an orchestra, [to] keep the ensemble of human body beating to a collective time”. SCN plays a crucial role in entrainment to synchronize the internal and external (environmental) time.

The phase difference between internal and external time is known as the *phase of entrainment*. The difference in the phase of entrainment across individuals results in different *chronotypes* [199]. In other words, some of us are innately early types (“lark”), while others are late types (“night owl”). Chronotype is a phenotype, so it is determined by both genetics [229] and environmental factors [199]. For example, age and gender can influence chronotype [193]. Children usually are early chronotypes but becomes increasingly late types as they grow up with a peak around the age of 20 [197]. In other words, we reach maximum “lateness” around the age of 20, which has been suggested as the end marker of adolescence [198]. After that, we progressively become early chronotypes as we get older. On average, adults older than 60 years tend to have earlier chronotypes than the children. Similarly, gender can also influence chronotypes. Men tend to have later chronotypes compared to women [197]. However, the gender difference disappears after the age of 50 — the average age of menopause onset.

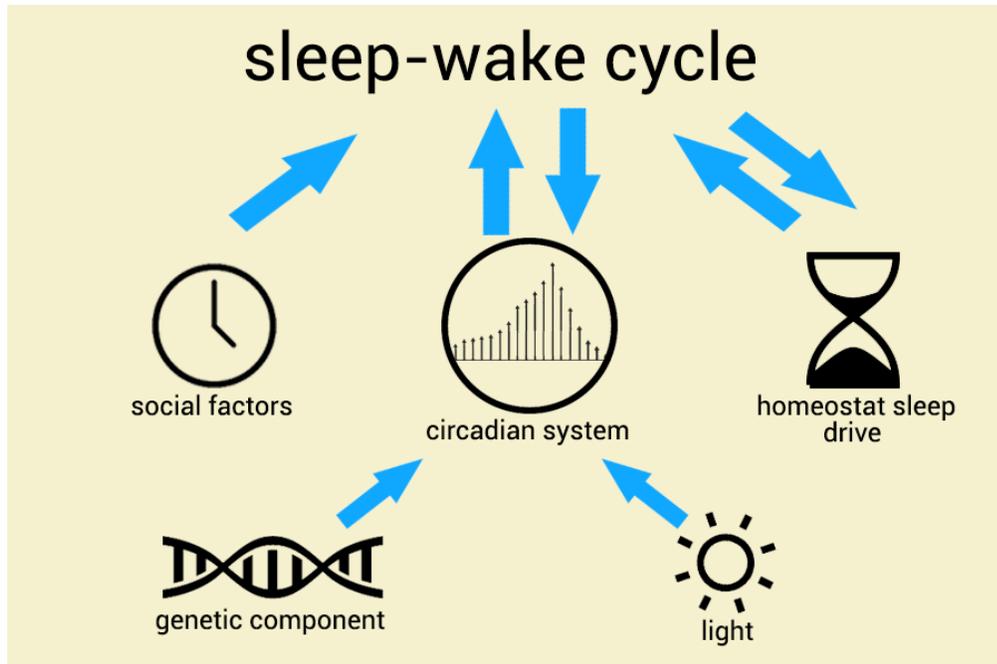


Figure 2.1: Sleep onset and duration depends on our circadian system, homeostatic sleep pressure, and social behavior.

2.1.1 Sleep and Wakefulness Cycle

Sleep and wakefulness cycle is perhaps the most ubiquitous and prominent example of biological rhythm. The regulation of sleep and wake behavior relies on two different biological mechanisms: i) the circadian system that determines the periodic and optimal time frame for sleep, and ii) the homeostatic sleep pressure that increases with the duration of time awake. In other words, the circadian process maintains the rhythm of sleep propensity with a 24 hour period, while the homeostatic sleep pressure depends on how long an individual has been awake. The interaction between these two processes determine when an individual goes to sleep and wakes up. Borbély [28] defined this as the *two process model of sleep regulation*.

However, we are also social animals and our behavior often gets influenced by our social relationships and responsibilities. This is true for sleep as well. Sleep onset and duration depends on our social behavior. In other words, as shown in Figure 2.1, sleep is determined by three complicated factors: our circadian system, homeostatic sleep pressure, and social behavior. These factors can be highly personalized depending on one's genetic makeup, age, gender, and environmental setups. As such, *not all of us can, or should, follow "early to bed and early to rise" lifestyle* since it might not align with innate our circadian systems.

2.2 Circadian Disruption

Our circadian system gets disrupted when our behavior and environmental cues are not aligned with our internal time. For example, travel across time-zones or imposed sleep-wake behaviors at inappropriate times can cause circadian disruptions. In our modern society, such disruptions are becoming quite common across the population. In particular, shift-workers are known to suffer from chronic circadian disruptions and around 15–20% of total workforce in most industrial societies are shift-workers [27].

Beyond shift-workers, circadian disruption is common among general population as well. Using data from around 55,000 participants, Roenneberg et al. [196] found that 70% of the population show sleep-wake behaviors indicative of circadian misalignments. Specifically, there is a significant discrepancy between sleep onset and duration across workdays and weekends. For example, late chronotypes suffer from "sleep debt" on weekdays as their social responsi-

bilities force them to wake up early. They compensate this accumulated sleep debt by sleeping significantly more during weekends. This is similar to having jet lag that would result from traveling across time zones westerly on Friday evening and then returning back on Monday morning. Roenneberg et al. [196] termed this circadian disruption as *social jet lag*.

2.2.1 Impact on Health

Chronic circadian misalignment can cause serious health and well-being issues [183]. Sleep and circadian disruption impairs our overall immune system [128] which can have serious clinical implications. Circadian disruption can increase the risk of cardiovascular disease. In particular, shift-work has been associated with myocardial infarction and ischaemic stroke [231].

Long term circadian disruption can also result in an increased risk of cancer. From a meta-analysis, Megdal et al. [142] concluded that long term night shift-workers are 48% more likely to get breast cancer (relative risk (RR) of 1.48 with 95% confidence interval 1.36–1.91). A number of studies have similarly associated circadian disruption with increased risk of prostate and colorectal cancer [233]. Indeed, International Agency for Research on Cancer (IARC) from World Health Organization has concluded that “shift-work that involves circadian disruption is probably carcinogenic to humans” based on existing evidence [215].

Circadian misalignment also disrupts our metabolic system. As such, it can result in higher risk for metabolic syndrome, obesity, and type 2 diabetes [209]. Roenneberg et al. [196] also found that social jet lag is associated with increased body mass index (BMI).

2.2.2 Impact on Cognitive Performance

Circadian disruption can also cause serious deficit in cognitive performance. In particular, alertness — a key biological process underlying our cognitive performance — reflects circadian rhythms [223]. Sleep deprivation and circadian disruption can result in poor alertness and physical reaction time [221]. Indeed, Lamond et al. [125] found that the decline in cognitive performance after 20–25 hours of wakefulness is equivalent to blood alcohol concentration (BAC) of 0.10%. To compare, in New York State, a BAC of more than 0.05% is considered as “impaired” and 0.08% is considered as “intoxicated” [165]. In other words, the effects of sustained sleep deprivation and circadian disruption on cognitive performance is similar (or worse) to being intoxicated.

The lack of alertness and fatigue issues have broader implications for well-being as well. Around 10 – 30% of fatal road accidents happen due to driver fatigue resulting in 1,550 deaths, 71,000 injuries, and \$12.5 billion in monetary losses per year in the US [74]. Sleep and circadian issues also result in serious productivity loss in workplace. Hafner et al. [91] reported that the yearly economic loss caused by insufficient sleep amounts to a staggering \$411 billion in the USA alone.

Sleep and circadian disruption also adversely impact memory and learning capabilities. In particular, hippocampal-dependent learning and memory forming strongly reflects circadian influence [210]. Gerstner et al. [84] hypothesized that long-term memory formation and circadian rhythmicity shares common mechanism. Based on data from airline flight crews, Cho [43] concluded that chronic jet lag causes temporal lobe atrophy, which might result in memory impairment and amnesia [30].

2.2.3 Impact on Mental Illnesses

Circadian disruption has been associated with a number of mental illnesses including alcohol and substance abuses [192, 96], anxiety disorder [171], schizophrenia [241], and bipolar disorder [94]. For example, sleep and circadian disruption impacts around 30–80% of patients with schizophrenia, making it one of the most common and consistent feature of the illness [46, 173]. Circadian disruption has been associated with bipolar disorder as well. In particular, sleep deprivation and circadian anomalies can trigger manic episode in patients with bipolar disorder [95]. A decreased need for sleep is also considered as a fundamental marker of mania phase [179].

A number of recent genetic studies have provided a deeper insight into the link between circadian clocks and mental illnesses. Specifically, these studies have associated schizophrenia with *clock genes* (genes involved in the generation of biological rhythms). Similar link between bipolar disorder and clock genes has also been established. For example, Benedetti et al. [22] reported association between a variant of clock gene *Per3* and early onset of bipolar disorder.

Furthermore, some researchers now hypothesize that sleep and circadian disruptions are not just limited to the pathology of mental illnesses, but might be more directly involved in disease etiology [106]. For example, Menet et al. [143] suggested that disruption in circadian system might be an underlying cause of the development of a number mental illnesses including bipolar disorder, depression, schizophrenia, and addiction. Lewis et al. [132] also suggested that sleep and circadian disruption might trigger postpartum psychosis in vulnerable women.

Methods	Example	Instrumentation	Considerations
Physiological markers	Core body temperature	Rectal thermometer [154]	Not scalable and very invasive
	Melatonin	Blood (serum and plasma), saliva, and urine [149]	Invasive and difficult to deploy at scale
	Cortisol	Blood, hair, saliva, urine, or feces [145]	Invasive and difficult to deploy at scale
	Gene expression	Blood sample [104], hair follicle [5]	Requires special hardware for data collection making it infeasible to deploy at scale
Self-assessment Surveys	Sleep Journal	Journaling of sleep onset and duration [247]	Can be unreliable due to potential non-adherence and unreliable recall
	Munich ChronoType Questionnaire (MCTQ)	MCTQ for assessing chronotype and social jet lag [199]	Unreliable recall can be a problem. Also, not suitable for tracking granular changes
Ubiquitous Technology	Heart rate	ECG (electrocardiography) [216], Wearables (e.g., Apple Watch [10])	Not a very reliable marker as external factors (e.g., carbohydrate intake) can influence it [119]
	Activity	Accelerometer based sleep and circadian rhythm monitoring [7]	Requires users to wear special devices all the time, so non-adherence can be a problem
	Mobile sensing	Smartphone usage patterns and sensors [3]	Participants must carry their phones consistently

Table 2.1: Existing methods for assessing circadian phase markers and disruptions

2.3 Assessment of Chronotype and Circadian Disruption

Given the crucial role of circadian rhythms in our health and well-being, there has been an increased focus on tracking circadian disruption. In this section, I will detail the most commonly used methods for assessing chronotypes and identifying circadian disruption.

2.3.1 Physiological Marker Based Assessment

Variation in our biochemical processes can be used as robust markers of circadian rhythms. For example, core body temperature (CBT) has been widely used to track circadian phase in individuals [50]. Humans, similar to other homeothermic organisms, maintains core body temperature through a complex thermoregulatory feedback system. CBT reflects a robust circadian trend [186] — it reaches maximum temperature during the day, begins to decrease at the onset of sleep, and drops to a minimum during major sleep phase (e.g., about 2 hours before waking up). CBT, as such, can be used as a reliable biomarker of circadian phase of an individual.

However, measuring CBT is very challenging. The most accurate method for assessing CBT is through rectal probing [50, 154]. Given its highly intrusive nature, a number of recent studies have focused on developing alternative methods for assessing CBT. For example, Niedermann et al. [162] used non-invasive wearable sensors tracking skin temperature, heat flux, and heart rate to predict CBT. However, the accuracy of these non-invasive approaches depend on environmental and physical conditions. For example, sweating can cause inaccurate assessment [243]. As such, using CBT as circadian phase marker in a scalable and unobtrusive way is still not feasible.

Melatonin and cortisol have also been used to track circadian phase. Melatonin is a hormone secreted by the pineal gland. Its level can indicate the onset of biological night [187] — melatonin concentration is low during the day and higher at night [44, 50]. Also, longer periods of darkness has been associated with longer duration of melatonin secretion [44]. Melatonin is considered as one of the most reliable markers of circadian phase. It is robust against a number of

external factors [149]. Cortisol is another hormone that reflects circadian trend [117]. Though, cortisol is a less robust marker than melatonin since a number of external factors including stress can influence it [147].

Melatonin concentration can be inferred from blood (serum and plasma), saliva, or urine [149]. Cortisol levels can also be measured from blood, hair, saliva, urine, or feces [145]. To identify circadian phase of an individual, researchers take frequent samples over the day. For example, Czeisler et al. [50] measured melatonin from blood samples every 20 to 60 minutes to calculate the intrinsic period of the human circadian pacemaker. However, outside of a laboratory, taking periodic blood, saliva, or urine samples poses serious practical challenges. Furthermore, storing these biological samples for later chromatography and/or mass spectrometry analysis can also be burdensome and costly. As a result, these methods are not suitable for in the wild studies.

A few recent studies have used gene expression for assessing circadian phase. Hughey et al. [104] developed ZeitZeiger that uses machine learning algorithms to predict circadian phase from genome-wide gene expression in human blood. However, given the difficulty in collecting blood samples, ZeitZeiger might be difficult to deploy at scale. Akashi et al. [5], on the other hand, used gene expression in hair follicle cells to infer circadian phase. In their study, it was sufficient to get 1–10 scalp hairs or 1–5 facial hairs to accurately detect circadian phase [136].

Compared to melatonin based methods using blood or saliva samples, this is significantly less intrusive. Furthermore, the required sampling frequency to accurately track circadian phase using this method is also significantly lower than melatonin based methods [136]. However, gene expression based methods

are often costly to deploy over large population. Measuring gene expression in the wild can be particularly challenging in terms of logistics. For example, these methods often employ polymerase chain reaction (PCR), which requires special hardware and devices.

2.3.2 Survey Based Assessment

Due to logistical limitations, methods using physiological markers for inferring chronotype and circadian disruption have been mostly limited to small-scale laboratory studies. To address these issues, recent studies have used survey based assessments. For example, sleep journal based methods have been widely used for inferring circadian stability. Zee et al. [247] used data from sleep diary to identify clinical cases of circadian disruptions. However, data from sleep journals can be unreliable due to recall error and non-adherence.

Horne et al. [99] developed Morningness-Eveningness Questionnaire (MEQ) to categorize individuals into morning or evening personality types. MEQ focuses on assessing subjective *daily preference*. As Roenneberg [195] pointed out, daily preference is different from chronotype, which is measured from phase of entrainment. To infer chronotype and social jet lag, Roenneberg et al. [199] developed Munich ChronoType Questionnaire (MCTQ).

Munich Chronotype Questionnaire (MCTQ)

MCTQ collects data about sleep and wake behavior for both work and free days. From this data, the mid-sleep point on free day (MSF) — the half-way time

point between sleep onset and waking up — is calculated as one’s chronotype. However, most people accumulate sleep debt on work days and compensate it by sleeping more on free days. For calculating chronotype, MCTQ corrects for the oversleep during free days and chronotype is assessed as the corrected mid-sleep point (MSF_{SC}):

$$MSF_{SC} = MSF - 0.5 (SD_F - (5 * SD_W + 2 * SD_F)/7)$$

Here, SD_F and SD_W are the average sleep duration on free days and work-days, respectively. The term $(5 * SD_W + 2 * SD_F)/7$ is the weighted average sleep duration per night. Given that age and gender can influence chronotype distribution, MSF_{SC} can be further corrected for these factors [197]. The assessment from MCTQ shows strong correlation with MEQ [246]. MCTQ has also been validated against sleep logs, actimetry, and biomarkers including melatonin and cortisol [197].

These survey-based methods are significantly easier to deploy beyond the controlled environment of a laboratory. However, they also pose a number of challenges in terms of data accuracy and user burden. Survey based methods are subject to recall errors. Also, longitudinal tracking of circadian phase using these surveys is difficult. Furthermore, these methods lead to infrequent data sampling. As such, instantaneous detection of sleep and circadian disruption based on survey data is often infeasible.

2.3.3 Ubiquitous Technology Based Assessment

Ubiquitous technology can enable in-situ and broad-data collection strategy for better tracking of circadian disruptions over a long period of time. Specifically, these technologies can enable reliable tracking of sleep-wake behavior, which in turn can be used for determining circadian patterns and disruptions. For example, actigraphy — a wearable device — has been widely used in clinical studies focusing on sleep and circadian disruption [7]. Actigraphy devices continuously collect accelerometry data, which can be used to determine sleep onset and wake up patterns. Actigraphy based sleep methods are robust and have been validated against polysomnography (PSG) [202]. Actigraphy is less invasive than most of the physiological marker based methods mentioned earlier.

However, the requirement of wearing a specialized device throughout all day and night is not ideal. Actigraphy devices can also be quite costly. Furthermore, retrieving data from actigraphy devices can be time-consuming and not ideal for real-time tracking. As such, using actigraphy for a large scale tracking of circadian patterns over a long period of time can be quite problematic due to device-burden and wear-compliance issues.

Heart rate also reflects circadian rhythm and can potentially be used to track circadian misalignment. For example, Taillard et al. [216] proposed to use heart rate as a biological marker of misalignment in major depression. Given that a number of recent wearables including Apple Watch [10] tracks heart rate, it can potentially be used for continuous and scalable assessment of circadian disruption in the field. However, heart rate based method is not very stable given that a number of external factors can influence it. For example, high carbohydrate intake can significantly elevate heart rate and as a result, can mask circadian

phase change [119].

Phone Based Assessment

Given the deep reach of phones in our daily lives and their ever increasing sensing capabilities, a number of recent studies have proposed phone based methods for assessing sleep. For example, Hao et al. [92] proposed iSleep framework that uses phone microphone to collect audio data for inferring sleep related events (e.g., moving and snoring). Based on data from 7 participants over 51 nights, their system achieved 90% accuracy in sleep-related event classification. Similarly, Krejcar et al. [121] developed wakeNsmile application, that uses audio data to infer sleep stages.

Chen et al. [42] used a number of phone usage features (e.g., recharging and screen unlocking) along with environmental cues (e.g., ambient sound and light) to predict sleep duration. They evaluated the proposed algorithm in a week long study with 5 graduate students and 3 visiting scholars. Based on this data, their model on average can estimate sleep duration within 42 minutes of ground-truth. Min et al. [148] used sound, light, movement, screen state, app usage, and battery status to classify sleep state and quality. Based on data from 27 participants over a month, their proposed methods can infer sleep state with 93.06% accuracy. Saeb et al. [203] developed a two-stage algorithm to infer sleep event from phone sensor data. They deployed their system over 207 participants for 6 weeks. From this data, their proposed algorithm achieved 81.8% accuracy.

While these studies have resulted in novel sleep detection algorithms, these methods often are computationally expensive and as a result can cause serious

battery drainage. As such, these methods can interfere with the normal daily use of phones. Some of these methods also use audio data, which can be privacy sensitive. More importantly, none of these studies considered circadian factors. Given that our circadian system underpins our sleep and wake behaviors, these studies at best get half the picture. For example, sleep issues might be symptoms of a misaligned biological system and methods focusing on sleep events alone will not be very effective.

As such, the existing methods for tracking sleep and circadian factors are not adequate. Specifically, they are not capable to facilitate long term and granular data collection critical to better understand the impact of circadian factors in our day to day lives. Till Roenneberg, as usual, eloquently conveys this point (emphasize mine):

“Researchers have made great advances in understanding which neurotransmitters and brain regions are involved in sleep, and how the timings of sleep and wakefulness are controlled by an internal (circadian) clock, among other things. *Yet we still do not have answers to the most basic questions.* It is not really understood, for instance, what sleep is for, how much is optimal, how sleep quality can be measured and predicted, or the role of genetic and environmental factors in determining ideal sleeping patterns. One reason for this lack of understanding is that most of what is known about sleep comes from laboratory studies [...]

Assessments of sleep are also often based on subjective responses to questions about how 'well' people feel after they have slept, or whether they think they experienced a good night's sleep. To learn about sleep in the real world, and to establish how to manage sleep to improve productivity, health and quality of life, we need a multidisciplinary 'human sleep project' [...]

A key goal of a human sleep project would be to identify simple, effective indicators that sidestep the need for cumbersome electrodes, or for blood or saliva samples, which are used to obtain conventional markers of circadian rhythms, such as the hormone melatonin". [194]

That is, there is an explicit need for technologies that will track behavioral, sleep, and circadian data in the wild. These technologies also have to be scalable and unobtrusive, so that they can be deployed across a large population over a long period of time. And, *this is where Circadian Computing comes in.*

Specifically, Circadian Computing aims to develop novel scalable methods for continuous and in-situ monitoring of sleep and circadian disruptions. By leveraging widely-used sensor rich platforms including phones and wearables, it also focuses on collecting data about relevant behaviors and contexts. Being able to collect in-situ data over a longitudinal period of time can help to bridge the gap between findings from the controlled environment of a laboratory and the implications in the real world. In chapter 4, I will describe such a scalable and data driven method for tracking sleep and circadian disruptions. Specifically, I will show how phone usage patterns can be used to accurately infer chronotype, social jet lag, and sleep inertia.

Circadian Computing is not limited to technologies that only monitor sleep.

Our circadian system influences biological processes beyond sleep as well. Indeed, almost every neurobehavioral process reflects circadian rhythm. Circadian misalignment, thus, can negatively impact these biological processes, which in turn can disrupt our daily behaviors and undermine our capabilities. However, the link between circadian disruption and its impact on our capabilities in the wild is under-explored. Circadian Computing can help to untangle this complex relationship by developing data driven methods that can model and predict the impact of disruptions on the outcomes of a circadian process. In chapter 5, I will describe such method for modeling and predicting alertness — a key circadian process underlying our physical and cognitive performance — in the wild.

Furthermore, Circadian Computing also enables development of technologies focusing on intervention. These intervention technologies will help to i) stabilize the circadian system to minimize the risk of disruption (i.e., “fixing the broken clock”), and ii) optimize our physical and cognitive performance *in accordance with* our individualized innate rhythms. In particular, given the importance of maintaining circadian stability in mental health patients, Circadian Computing can play a significant role in mental health care. In chapter 6, I will describe an example of a circadian-aware technology in this context. Specifically, I will describe the design, development, and evaluation of a data driven method that can help patients with bipolar disorder to monitor and maintain circadian stability.

2.4 Summary

In this chapter, I focused on the theoretical background underpinning Circadian Computing. Specifically, I describe the critical importance of circadian stability for our health and well-being. I have also pointed out that despite significant recent developments in our understanding of the underlying causes and effects of circadian disruptions, these methods and findings are often limited to a laboratory environments. As shown in Table 2.1, these methods are not adequate for granular monitoring of circadian disruptions in the wild over a longitudinal period of time. As a result, there is a need for novel pervasive technologies for tracking, monitoring, and modeling circadian disruptions and its impact in the real-world setup. There is also an opportunity for developing intervention tools for maintaining circadian stability. This dissertation aims to bridge this gap in pervasive technology focusing on circadian factors. In the following chapters, I will describe my overall PhD work, which I believe, is a leading step towards this broad vision of circadian-aware technologies for sensing, adapting to, and stabilizing our innate biological rhythms.

CHAPTER 3

PASSIVE SENSING OF CHRONOTYPE AND CIRCADIAN DISRUPTIONS

In the previous chapter, I have described the need for continuous and granular assessment of circadian disruptions. I have also noted that these methods need to be scalable so that they can be deployed over a long period of time across a large population. In this chapter, I will describe a smartphone based sensing method that addresses these challenges. In particular, I will show that phone usage pattern reflects our daily behavior and thus can be used for assessing sleep duration, chronotype, and circadian disruptions. I will also describe how these findings can have broad implications for future sensing and intervention technologies.

3.1 Introduction

As shown in Table 2.1, chronobiologists have used a wide range of tools for assessing chronotype and circadian disruptions. However, most of these methods are limited to laboratory setups only (e.g., taking periodic blood samples). On the other hand, survey based methods (e.g., MCTQ, sleep journals) are not suitable for longitudinal tracking. They also do not have enough data resolution to identify subtle changes. A number of other studies have used accelerometer based sleep and circadian rhythm monitoring (e.g., actigraphy). However, this requires a special device, which limits the ability to deploy over a large population.

Ubiquitous technologies can successfully address some of these issues. In

particular, phones and wearables can reshape the existing landscape of tools and methods for tracking circadian disruptions. The global widespread use of phone means that scalability and large scale deployment issues are relatively easier to solve. Phones now have deep reach in our modern lives, and thus can be successfully used to track daily behavior. Indeed, a number of recent studies have used phone based methods for detecting sleep patterns [148, 203]. However, these technologies often do not take factors related to circadian rhythms into into consideration (e.g., relationship between chronotype and social jet lag). As a result, these methods are at best getting half the picture. In this chapter, I will aim to address the limitations of these existing methods by developing a passive sensing framework that leverages phone usage patterns to assess chronotype, social jet lag, and sleep inertia.

In this work, I made the following contributions:

- I developed a low cost method to infer sleep onset and duration using phone usage patterns. The low computational and data requirements of this algorithm means that it is highly scalable.
- My collaborators and I also deployed this system over 97 days and 9 participants. Based on this data, I found that the phone usage based method can successfully infer sleep and chronotype. I also showed that phone usage data can be used to monitor markers of circadian disruptions including social jet lag, and sleep inertia.

3.2 Methods

For this study, my collaborators and I decided to focus on college students given that inadequate sleep and the resulting circadian disruptions are often a serious problem for this population [217]. Moreover, they tend to be habituated phone users [129], which makes them a great fit for this study. Participants were recruited through public mailing lists and snowball sampling. In total, the study had 9 participants (7 males, 2 females) with an age range of 19–25 years. All the participants were habituated phone users — they had been using phone in their daily lives consistently for at least six months prior to the study. Seven participants reported using their phones as alarm clocks. All the participants reported using their phones immediately after waking up for a number of activities including checking mails, interacting with social media apps and using browsers.

The study lasted for 97 days starting from November 22nd, 2013 and ending by February 26th, 2014. Our research team was interested to see how circadian stability and sleep schedule changes across time and between socially versus individually determined schedules. My collaborators and I, therefore, planned the data collection period to span three distinctive phases in undergraduate academic life: 5 weeks at the end of the Fall semester, 4 weeks of winter break, and 5 weeks of Spring semester. One participant was interning during the Fall and therefore was not attending any classes. Rest of the participants had standard class schedules throughout the study duration.

3.2.1 Data Collection

For assessing individual chronotype, I administered the Munich ChronoType Questionnaire (MCTQ) at the beginning of the study. As mentioned before, MCTQ calculates “corrected” mid-sleep point (MSF_{sc}) by taking sleep debts during the workdays and oversleep during the weekends into consideration. This corrected mid-sleep point (MSF_{sc}) is then used as individual’s chronotype. Table 3.1 shows the distribution of chronotypes of the participants. Overall, they had significantly late chronotypes compared to the general population, which is expected given their age range. However, participant 4 was a very early type with a chronotype of 03:02 AM. This wide distribution of chronotypes was particularly useful to compare social jet lag and sleep inertia across early and late types (discussed in a later section).

ID	Chronotype (MSF_{sc})	Age Range	Study Duration (Days)	Valid Journal Entries
1	06:59	20-21	97	93
2	06:41	20-21	96	94
3	06:12	18-19	95	28
4	03:02	18-19	93	66
5	06:38	18-19	93	78
6	05:18	18-19	91	66
7	04:54	22-24	87	80
8	04:41	18-19	92	46
9	05:43	20-21	76	74

Table 3.1: Participant Information

Each participant maintained a daily online sleep journal to record sleep onset and duration information. Participants also noted information about sleep disturbances. I set up an automated system that sent a daily mail to participants reminding them to complete the sleep journal. To ensure data quality and

prevent recall errors, I discarded any retrospective entries and only considered journal records for the previous day's sleep.

For collecting phone use data, I developed a custom Android application. It runs in the background and collects call, SMS, location, browser history, application usage, and screen on-off information. To maintain privacy, sensitive data (e.g., call, SMS and browser history) was one-way hashed. The application stored the data in the phone and periodically uploaded it to a secure server. We also collected social media logs of each participant (e.g., logs of Facebook status update). To get better insights about the phone use habit of the participants, we interviewed them three times during the study (beginning, middle, and at the end of the study). In this chapter, I will mostly focus on presenting the results based on phone use patterns. For findings regarding social media logs and interviews, please see the article by Murnane et al. [159].

Participants were compensated based on valid sleep journal entries and the duration of the study. The Cornell Institutional Review Board approved all procedures.

3.2.2 Sleep Algorithm

As mentioned earlier, one of the main goals of this study was to come up with a scalable and unobtrusive sleep algorithm. For this, I developed a rule-based algorithm to infer sleep events from phone usage data. In particular, it uses screen on and off information as input and returns sleep onset, duration, and mid-point as output. The pseudocode is listed as Algorithm 1.

During the first phase of processing, the algorithm identifies phone use sessions based on screen on and off events. Given that notifications from phone applications can briefly turn on the screen, I used a duration threshold (θ) to identify active user interaction. For this study, I set the threshold to 30 seconds. In other words, all usage sessions that lasted less than 30 seconds were discarded. I used these filtered sessions to calculate non-usage patterns. The sleep event, in this algorithm, is assumed to coincide with the longest non-usage block.

The algorithm makes two adjustments before calculating the final sleep onset and duration information. First, the algorithm only considers non-usage patterns *starting* between 10 PM to 7 AM for sleep detection. This is due to the fact that the participants were not shift workers and we were mostly interested about their night sleep behavior. Second, the algorithm adjusts the sleep duration by adding an individualized corrective term (δ) to the non-usage duration. For each participant, the corrective term is learnt using the first two weeks of data. This corrective term addresses individual differences in phone usage behavior. In other words, if there is a consistent difference between inferred sleep and ground-truth, the corrective term (δ) can help to minimize the error.

The algorithm can make sleep overestimation error (i.e., inferred sleep > ground truth) when a participant does not use her phone right before going to sleep or immediately after waking up. For example, if a participant sleeps from 11:45 PM to 7:00 AM, but does not use his phone from 11:30 PM to 7:15 AM, then the inferred sleep duration would be 30 minutes longer than the actual sleep. If there is a consistent overestimation error for an individual, then it can be minimized by using the corrective term (δ).

Similarly, the algorithm can also make sleep underestimation error (i.e., in-

ferred sleep < ground truth). This might happen when screen on and off events do not reflect active user interactions. In other words, if a non-usage duration gets interrupted by a false screen-on event, then the inferred sleep duration will be less than the actual sleep duration. To filter out these false positives, I used a time based threshold (θ). However, there might be cases for which this threshold based method might not be adequate. For example, if a user keeps snoozing the alarm and goes back to sleep, then the longest duration of non-usage would be smaller than the actual sleep. If such cases are consistent for an individual, the corrective term (δ) can also minimize the underestimation error.

3.3 Findings

In the following section, I will describe the findings based on the study data. I will first discuss the accuracy of the phone usage based algorithm to identify sleep onset, duration, and mid-point. I will also point out how these inferred values can be used to track sleep debts and social jet lag. Further, I will describe how phone usage patterns can be used for monitoring sleep inertia which might indicate circadian disruptions as well.

3.3.1 Sleep Algorithm Accuracy

As shown in Figure 3.1, the sleep algorithm accurately identifies sleep onset and duration for each participant. The average difference between inferred sleep duration and ground truth is less than 45 minutes for all participants. The accuracy of this algorithm is comparable with a more computationally expensive

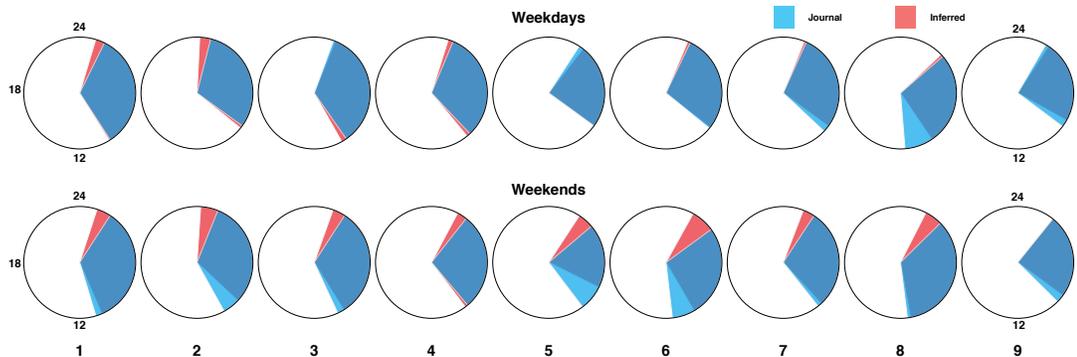


Figure 3.1: Average sleep onset and duration across participants from phone and journal data. Phone non-usage coincides with sleep events; the trend is more stable on weekdays due to more data points. For most participants, sleep onset is delayed and duration is longer during the weekends while participant 4 — with an early chronotype — gets less sleep on weekends.

method from Zhenyu et al. [42] that used a wide range of environmental sensing including light, sound, and user movement. Compared to this method, the proposed algorithm is significantly low-cost both in terms of computation and data requirements. It also accurately identifies sleep mid-point for all participants. The average difference between inferred sleep mid-point and ground truth across all participants is 23.8 minutes (95% CI: 11 mins). Being able to infer mid-sleep point accurately is important since this information can be used to identify circadian disruptions and social jet lag as discussed below.

Replication by Others

Furthermore, the accuracy of this algorithm has been replicated by Cuttone et al. [48] in a large study. They collected data from 126 participants over 2–4 weeks. They used wearable devices (SWR10 and SWR30 fitness tracking armbands from Sony Mobile) to track sleep data. In this study, they aimed to classify

between sleep and awake events. In other words, it was a binary classification problem. They found that the proposed algorithm is highly accurate in identifying sleep and awake events with a mean accuracy of 0.89 and mean F1-score of 0.83. This successful replication using data from a completely different population reaffirms the generalizability and efficacy of Algorithm 1.

3.3.2 Assessing Sleep Debt

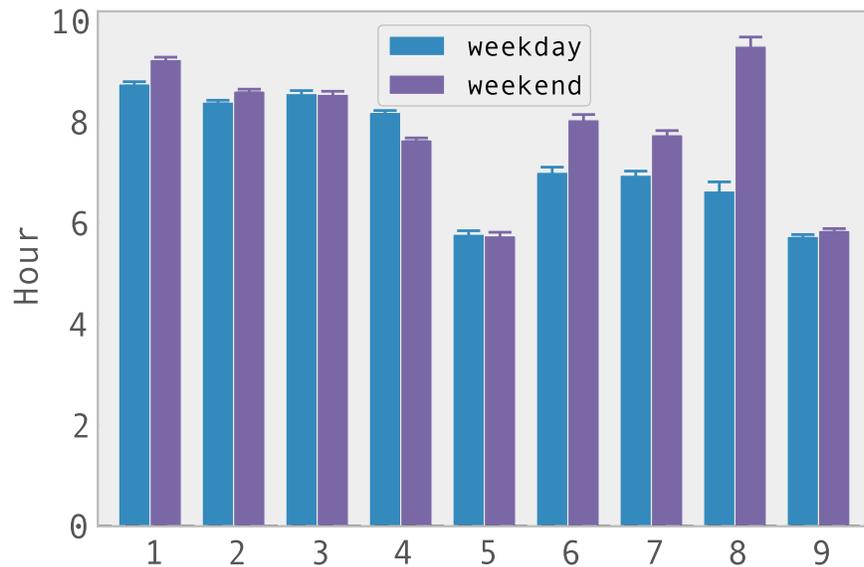


Figure 3.2: Average sleep duration across participants with 95% confidence interval

Given the high accuracy of the proposed algorithm, it can be used for continuous monitoring of sleep and circadian disruptions over time. In particular, I was interested to identify the patterns of sleep debts throughout the study. For this, I compared the inferred sleep duration across weekdays and weekends as shown in Figure 3.2. For all participants, there is a distinct difference in sleep

duration between weekdays and weekends. Specifically, all the participants except participant 4 and 5 slept more on weekends. This indicates that the participants accumulated sleep debts over weekdays and compensated it by sleeping more on weekends. Indeed, after excluding participant 8 (an outlier who sleeps significantly longer during the weekends) and participant 4 (the early chronotype), on average participants slept more than 20 minutes during the weekends compared to the weekdays.

Of the two participants who slept less on weekends, participant 5 is a late type. However, he often had external responsibilities on weekends. In other words, weekends were not necessarily “free days” for him; as a result, he did not have a consistent sleeping pattern over time. On the other hand, participant 4 — the early chronotype — consistently slept less on weekends. Indeed, participant 4 slept an average of 32.56 minutes less on weekends. I speculate that the sleep loss for participant 4 might happen due to the social pressure from his peers to stay late during weekends. Given the age range, most of his peers are going to be late types, so social activities with his peers would result in going to sleep late. However, his body clock causes him to wake up early in the morning resulting in consistent sleep deprivation during weekends. There was further evidence for such behavior in this dataset — the sleep onset of participant 4 gets delayed by 38 minutes on average during the weekends. This different sleep patterns of early and late chronotypes showing opposite trend between workdays and free days is a well-known phenomenon termed as the “scissors of sleep” [193].

When considering the different phases of the study, the sleep patterns show some interesting trends. In particular, during the Fall and Spring semesters, the

	Sleep Midpoint (AM) ($\pm 95\%$ CI) (Hr)
Weekday	05 : 24 \pm 0.02
Weekend	05 : 47 \pm 0.03
Weekday (fall)	05 : 06 \pm 0.04
Weekend (fall)	05 : 40 \pm 0.06
Weekday (winter break)	05 : 20 \pm 0.03
Weekend (winter break)	05 : 24 \pm 0.05
Weekday (spring)	05 : 30 \pm 0.02
Weekend (spring)	05 : 52 \pm 0.02

Table 3.2: Sleep midpoint across different phases.

sleep mid-point during weekdays are significantly earlier compared to weekends as shown in Table 3.2. This is due to the fact that during the Fall and Spring semesters, the participants often had imposed schedules (e.g., class schedules) which was inconsistent with their individualized rhythms (e.g., waking up early when one has a late chronotype). This resulted in an increased circadian instability, which is reflected by greater oscillations in sleep mid-point between weekdays and weekends. During the winter break, the schedule requirements were less stringent (i.e., participants could more freely choose their sleep timings). As a result, the circadian disruption was reduced for the participants during the winter break. Indeed, during this period, the average sleep midpoints differed only by 4 minutes between weekdays and weekends.

3.3.3 Social jet lag

To further quantify circadian disruptions across chronotypes, I calculated social jet lag from inferred sleep mid-points. Following MCTQ, social jet lag is the difference between mid-sleep point on free days (MSF) and mid-sleep point on

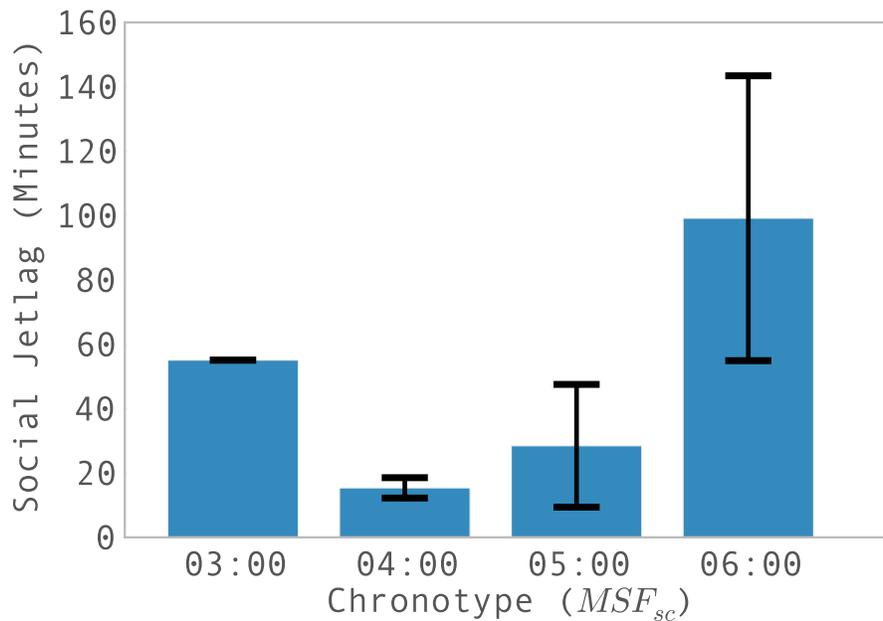


Figure 3.3: Duration of average social jet lag compared across chronotypes. 95% confidence interval has also been shown. Note that we have only one participant with chronotype 03:00, so interval estimation is set to zero in that case.

workdays (MSW)[238]:

$$\Delta MS = |MSF - MSW|$$

The distribution of social jet lag across chronotypes from this dataset is shown in Figure 3.3. Social jet lag is more pronounced among the late chronotypes. However, the participant with early chronotype also suffers from social jet lag given his consistent sleep deprivation during the weekends. Overall, the participants at the extreme ends of the chronotype spectrum suffer much more than those who are in the middle. This is consistent with the findings from a prior survey based study [238].

3.3.4 Inferring Sleep Inertia (SI)

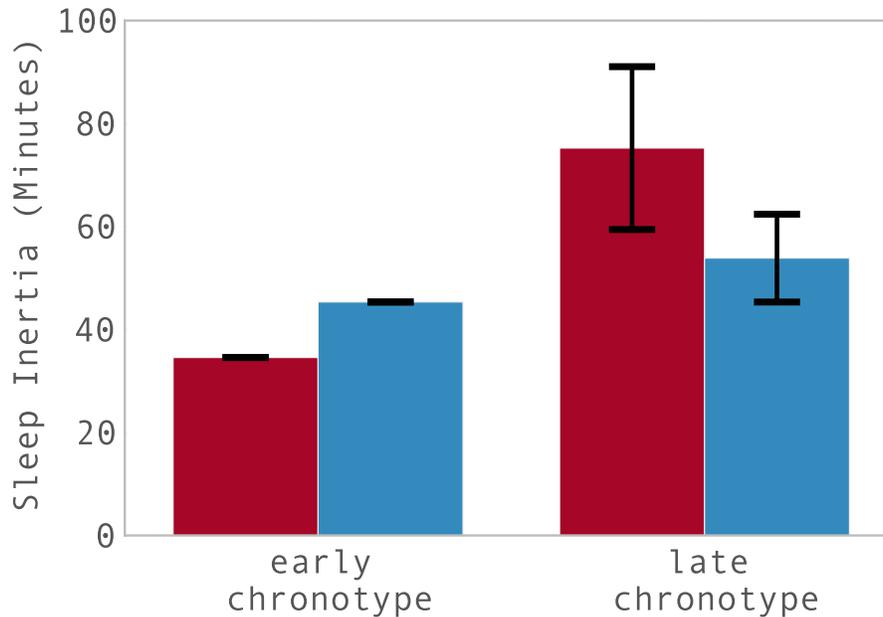


Figure 3.4: Inferred sleep inertia duration (with 95% CI) compared across early (N=1) and late chronotypes (N=8). Difference in sleep inertia duration from weekdays to weekends reflects the patterns of accumulated sleep debt across different chronotypes.

Sleep inertia — the time it takes to be fully awake — has also been used as a metric for assessing circadian instability [199]. To compute sleep inertia, MCTQ directly asks users how long they take to be fully awake. However, this might result in imprecise assessments due to recall errors from users.

Throughout this study, I noted that a majority of participants use their phones as a way to wake up. I decided to leverage this usage habit to compute sleep inertia. In particular, I hypothesized that an increased duration of sleep inertia would be reflected in longer duration of phone usage. I therefore defined sleep inertia (SI) as the total minutes of phone usage in the morning:

$$SI = \sum_{6 \text{ AM} < t < 12 \text{ PM}} \text{PhoneUsage}(t)$$

A comparison of calculated sleep inertia across weekdays and weekends between early and late chronotypes are shown in Figure 3.4. This data indicates that sleep inertia calculated from phone usage reflects circadian disruptions. That is, late chronotypes used phone much more as a waking up process during the weekdays as they suffer from sleep debts. However, for the early chronotype, the pattern was completely reversed, which again reflects the “scissors of sleep”.

3.4 Discussion

In this chapter, I described a passive sensing method that can be used for longitudinal monitoring of sleep and circadian disruptions over a large population. Specifically, I leveraged phone usage patterns for inferring sleep onset and duration. Compared to existing methods, this algorithm has significantly low overhead in terms of computational requirements and user burden. I also deployed the system over 9 participants for 97 days. The findings showed that the proposed algorithm can accurately infer sleep events. I further demonstrated that phone usage patterns can indicate circadian disruptions. Specifically, I established that late chronotypes suffer from sleep debt over workdays and compensate by sleeping more on weekends. The early chronotype, on the other hand, consistently slept less on weekends.

I also found that these discrepancies are more likely due to the mismatch be-

tween external (social) time and internal time. For example, these discrepancies are much lower during the winter break when the participants had less stringent schedules. Going beyond the sleep onset and duration, I also developed a new way to quantify sleep inertia (SI) as the total duration of phone use in the morning. I validated this method by showing that it reflects sleep debts over weekdays and weekends for both early and late types.

Given the widespread use of mobile phone around the globe, this phone usage based method can be potentially deployed over a large population for longitudinal tracking of sleep and circadian disruption patterns. Towards this goal, it is important to reflect on the limitations of the study and identify ways to extend it to address these limitations.

3.4.1 Limitations

The findings of the study are based on a small set of participants. Moreover, they are all college students resulting in a homogenous population. Given the age range of these participants, their phone usage and sleep behavior might be quite different from a broad population. As such, it will be important to investigate if these findings are generalizable to other populations. Specifically, it will be useful to check the accuracy of the sleep algorithm for the post-college population who have more regular work schedules. I speculate that the algorithm will be reasonably accurate for any population that habitually use phones. However, it might be necessary to come up with a more adaptive corrective term (θ) to address the individual differences in phone usage patterns for these populations.

Similarly, these participants were studying at a college in the USA. It will be useful to further validate the accuracy of the algorithm and the findings for people from different economic and cultural backgrounds. These participants often used phones for email checking and social media engagement, particularly right before going to bed and immediately after waking up. As a result, the accuracy of the algorithm is quite high. However, it will be crucial to ensure that the phone usage patterns are similar in countries with relatively scarce access to internet. In such cases, the algorithm might significantly overestimate the sleep duration. It should be noted that some of these generalizability issues (e.g., age range, population from different backgrounds) have been addressed by the replication study from Cuttone et al. [48].

The sleep algorithm is also not adaptive over time. It uses a corrective term (δ) to address sleep estimation error resulting from individual differences in phone use. However, δ is learnt using the first two weeks of data only. As such, the algorithm will not be able to address estimation errors if there is a change in phone use habit over time (after first two weeks). The algorithm should be adaptive in detecting changes in context (e.g., moving to a new place) to identify if there is any potential need for updating the corrective term. The current version of the algorithm also does not allow any user feedback. It will be very useful to extend the algorithm so that it can incorporate feedback from users. This will enable identifying consistent errors made by the algorithm and thus, a way to improve itself over time. One way of doing this would be providing a summarized representation of the decisions made by the algorithm and the user can make appropriate changes.

For this study, I did not take day-to-day schedule of a participant into con-

sideration. Such information can be particularly useful in further improving the accuracy of the algorithm. For example, an early class might mean significantly less use of phone in the morning and thus, the threshold parameter (θ) should be smaller for these days. Taking the daily schedule of the participants can also provide a better insight into the reasons and patterns for circadian instability. For example, if a participant consistently under-sleeps on Wednesday (e.g., for an early class), we might see the effects to continue over Thursday and Friday (e.g., sleep inertia). Similarly, I expect that the sleep and circadian issues to be significantly higher during the exam period for these participants. Being able to take advantage of such predictable changes can be very useful in both improving the accuracy of the algorithm and providing appropriate interventions.

Also, for calculating sleep debt and social jet lag, I assumed that the weekdays are workdays and weekends are free days for our participants. However, this assumption might not be true for a diverse population. For example, depending on class schedule, some students might consider Wednesday as a free day while the required preparation for Monday classes might result in working on Sundays for others. The distinction between workday and free day might also change over time. Future work should focus on being more adaptive when it comes to identifying workdays and free days of the participants.

Most of the participants in this study were late types, which is expected given their age range. It will be useful to extend these findings to a diverse population with a more balanced representation of chronotypes. I also considered chronotypes as discrete groups instead of a continuous spectrum. For example, I compared the sleep inertia pattern calculated from phone usage between early and late types in Figure 3.4. While this assumption works well for a small study,

it is important to go beyond this dichotomy of early and late types. Instead, it will be useful to see if these findings show consistent trends over the continuous distribution of chronotypes for a large population.

For this study, I did not use any environmental sensing. However, there are a number of external factors that might impact our circadian stability. In particular, light exposure can have significant impact on the phase of our circadian rhythms. However, environmental sensing can add significant user burden (e.g., wearing an extra device for tracking light exposure). Future work will do well to come up with wearable devices for tracking the relevant environmental factors while keeping the overhead (both computational and user burden) low.

3.4.2 Implications and Opportunities

The ability to passively and continuously monitor sleep and circadian stability lays down the foundation of Circadian Computing. Indeed, a low-cost, reliable, and unobtrusive way of inferring sleep and circadian disruptions will expand the current scope of personalized and context-aware technologies. The findings presented here, thus, can have impact across a wide range of application domains. In the following section, I will discuss the implications of these findings and potential opportunities.

Intervention for sleep and circadian stability

Sleep and circadian instability are becoming increasingly prevalent across the world. Sleep pathologies are reaching an epidemic level. Around 70 million

people in the USA alone suffer from chronic sleep disorder [164]. As a result, there has been an increasing interest from both academic and industrial researchers in measuring, assessing, and improving various aspects of sleep. However, the current sleep recommendations still do not take individual circadian patterns into consideration and often are too generic (e.g., “sleep for 7–8 hours”). Our circadian processes are different across individuals and as such, the appropriate timing and duration of sleep can be highly individualized. It is, therefore, important to go beyond the one-size-fits-all models and consider the relationship between sleep and underlying circadian systems while making these recommendations.

Moreover, interventions that only target sleep disturbances run the risk of merely focusing on the symptoms of misaligned biological clocks. Instead, a more holistic approach that takes one’s chronotype, social jet lag, and sleep debt into consideration will be more effective in addressing sleep and circadian disruptions. Towards this goal, the findings presented in this chapter could play a significant role. It can help sleep researchers identify more effective and actionable suggestions (e.g., appropriately timed exposure to sunlight). Further, these findings can help the participants to be more aware of their idiosyncratic patterns and make informed decisions about their daily activities and sleep.

Also, the use of modern devices has been associated with sleep and circadian disruptions [41]. Circadian Computing can play a significant role to minimize such disruptions by making these devices adaptable to individuals’ biological rhythms. For example, a circadian-aware device might dim the brightness or adjust the blue-enriched light emission at appropriate times of the day. Such adaptable and personalized technologies can be applied beyond laptops, tablets

and phones. Indeed, circadian-aware home and office environments can also provide adaptable and individualized light exposure to ensure better sleep hygiene and circadian stability.

Opportunities for Sensing and Measurement

In this chapter, I developed a passive sensing method that leverages phone usage patterns for inferring social jet lag, and circadian disruptions. Given that external variables — particularly light — can influence our biological rhythms, there is an opportunity to extend the proposed method to consider environmental factors as well. In particular, monitoring of light exposure could be useful for granular tracking of circadian phase. Such multi-modal framework that combines both behavioral and environmental data could be more accurate and provide better insights into the cause of disruptions. It can also be more adaptable to environmental changes (e.g., moving to a new timezone).

However, integrating environmental sensors into this framework poses some challenges when it comes to scalability and user burden. For example, light exposure monitoring requires specialized devices that can assess radiant flux density at the cornea [184]. To be able to model the influence of light exposure on circadian phase, such monitoring needs to be unobtrusive and passive. There has been some recent progress in developing light tracking wearables that are low-cost and easy to wear. For example, Figueiro et al. [68] have developed Daysimeter — which is only 2 cm in diameter — for monitoring light exposure. Future work will do well to combine these sensors with behavioral features for not only assessing disruptions but also identifying the causes for such disruptions.

3.5 Summary

In this chapter, I focused on passive monitoring of circadian disruptions. In particular, I focused on sleep given that it both reflects and modulates the underlying circadian system. While a number of lab and survey based methods have been proposed to monitor sleep and circadian disruptions, they often do not scale. As a result, it is impractical to deploy these tools over a large population for a longitudinal period of time. In this study, I leveraged the widespread use of phones in our daily lives for assessing sleep duration, chronotype, social jet lag, and sleep inertia. Based on data from 9 participants over 97 days, I found that this phone use based method can accurately infer sleep onset and duration as well as correctly identifying sleep debt and social jet lag. The accuracy of this sleep algorithm has been further validated by a large study from Cuttione et al. [48].

I also discussed the implications of these findings in the broader context of Circadian Computing. Being able to passively monitor circadian disruptions is a key building block that can be used for a wide range of applications. In particular, these findings can enable the design, development and deployment of circadian aware technologies for more adaptive and personalized support across different domains including effective interventions for sleep and circadian stability.

sOn : Ordered $N \times 1$ timestamp of screen-on events
sOff : Ordered $N \times 1$ timestamp of screen-off events
 θ : Threshold duration for phone usage
 δ : Individual corrective term
output: Calculated sleep duration, onset, and midpoint

```

1  $n \leftarrow 0$ 
2  $t \leftarrow 0$ 
3 for  $i \leftarrow 0$  to  $N$  do
4   /* Filter out non-interactive sessions based on
      time based threshold. */
5    $d_i \leftarrow \text{sOff}_i - \text{sOn}_i$ 
6   if  $d_i > \theta$  then
7      $\text{fOn}_n \leftarrow \text{sOn}_i$ 
8      $\text{fOff}_n \leftarrow \text{sOff}_i$ 
9      $n \leftarrow n + 1$ 
10  end
11 end
12 for  $i \leftarrow 0$  to  $n$  do
13   /* We are interested in sleep during night only.
      So we'll discard any non-usage patterns that
      does not start between 10PM to 7AM (next day).
      */
14   if  $\text{fOff}_i$  is between 10PM to 7AM (next day) then
15      $\text{nonUsage}_i \leftarrow \text{fOn}_{i+1} - \text{fOff}_i$ 
16      $\text{nonUsageOnset}_i \leftarrow \text{fOff}_i$ 
17      $t \leftarrow t + 1$ 
18   end
19 end
20 /* The longest duration of non-usage coincides with
      sleep. */
21  $\text{sleep}' \leftarrow \max_i(\text{nonUsage}_i)$ 
22 /* Sleep onset is marked by the beginning of the
      longest duration of non-usage block */
23  $\text{sleepOnset} \leftarrow \text{nonUsageOnset}[\text{argmax}_i(\text{nonUsage}_i)]$ 
24 /* Finally, using the individual corrective term to
      adjust sleep duration. */
25  $\text{sleep} \leftarrow \text{sleep}' + \delta$ 
26 /* Calculate sleep midpoint from onset and duration.
      */
27  $\text{sleepMidpoint} \leftarrow \text{sleepOnset} + \frac{\text{sleep}}{2}$ 

```

Algorithm 1: Computing sleep duration, onset, and midpoint from phone usage.

CHAPTER 4

PREDICTING RHYTHMS OF ALERTNESS

In the previous chapter, I have described a passive sensing method for monitoring sleep and circadian disruptions. In this chapter, I will focus on modeling a specific circadian process — alertness — in the real-world setup. In following sections, I will describe why assessing alertness is important and how the existing methods are inadequate for tracking it in the wild. I will also describe how phone usage patterns can be leveraged to infer alertness states. I will conclude this chapter by identifying the implications of these findings and potential opportunities.

4.1 Introduction

Our cognitive abilities are not constant over time. In particular, our alertness ebbs and flows over a 24 hour period reflecting circadian rhythm. Alertness can also depend on a number of external variables — including sleep, exercise, eating, and use of stimulants (e.g., caffeine or alcohol). It plays a key role in our day-to-day activities since it is a core subcomponent of our sensory, motor, and cognitive processing [181]. Indeed, *fatigue* — a state of diminished alertness — has been associated with accidents and serious occupational errors.

For example, around 10 – 30% of fatal road accidents happen due to driver fatigue resulting in 1,550 deaths, 71,000 injuries, and \$12.5 billion in monetary losses per year in the US [74]. In workplace, alertness related issues can also have severe negative impacts. The estimated cost from productivity loss due to

fatigue is \$136 billions per year [190]. The role of fatigue leading to serious errors in occupational setups has been well-established [237]. For example, physician are 22% more likely to make surgical errors due to alertness related issues [141]. Overall, shift workers have 61% higher injury or accident rate resulting from fatigue [57].

Given the key role alertness plays in our daily activities, safety, and well-being, a better understanding of its patterns can have a far reaching positive impact. For this, we need a way to continuously monitor one's alertness in the real world setup. However, the existing methods for assessing alertness are ill-equipped for such continuous monitoring over a long period of time.

The current methods for alertness assessment can be categorized into three groups: i) self-assessment surveys, ii) physiological tests, and ii) reaction time tests. Researchers have developed and used a number of self-assessment surveys for assessing alertness. For example, Epworth Sleepiness Scale (ESS) [108], Karolinska Sleepiness Scale (KSS) [6], and Stanford Sleepiness Scale (SSS) [98] have been widely used for this purpose. However, these self-assessment surveys might not track subtle and granular changes in alertness. They are also limited due to added user burden, low adherence, and recall errors. Moreover, there are conflicting findings regarding the reliability of self-assessment when it comes to alertness states [79, 222].

Physiological signals can also be used for assessing alertness states. A number of recent studies have used brainwave [110], eye movement [107], and eye blinking [140] information to infer alertness levels. However, these methods require special hardwares including EEG [110], electrooculography (EOG) [11], eye tracker or a camera [140]. This dependence on special and costly hardwares

means that these methods are constrained to lab environments only. Due to cost and logistical reasons, it is very difficult to deploy these methods at scale over a large population.

In recent decades, reaction time tests have been used for objective measurement of alertness. In particular, Psychomotor vigilance test (PVT) [133] has been widely used by clinical and sleep researchers for assessing one's alertness. PVT can take 3–10 minute. During the test, participant are presented with visual stimuli at random time intervals and asked to respond to it as soon as they see the stimulus. While the earlier version of PVT used specialized hardware, Matthew et al. [112] recently implemented PVT on a phone to make it more accessible and easy to deploy. However, PVT can be very burdensome to users given that it requires 3–10 minutes of undivided attention for each assessment. As such, it is unsuitable for continuous measurement over a long period of time.

In this chapter, I aim to address these issues by developing a method for continuous and unobtrusive assessment of alertness states. My contributions for this work are:

- I co-developed a phone based data collection framework for monitoring patterns of alertness in the real-world setup. This framework enabled collecting alertness data in the wild using PVT tests. The framework also collected diverse stream of self-reported and passively sensed data relevant to one's alertness states. Compared to previous studies, the phone based system moves beyond the controlled lab environment and assesses in-situ alertness of the participants as they go about their daily lives.
- I helped to deploy the system over 20 participants for 40 days. Based on these data, I found that alertness can oscillate approximately 30% depend-

ing on time and body clock types. I also found that Daylight Savings Time (DST), sleep, and stimulant intake can influence alertness. During the data analysis process, I particularly focused on replicating and extending previous lab based studies. Such replication is not only important for scientific progress but also helps to validate the proposed method for collecting alertness data.

- I also focused on developing an unobtrusive and passive method for assessing alertness that can be deployed among distributed and large populations. For this, I trained models for assessing alertness states using phone usage and other behavioral variables as features. I showed that these models achieve high accuracy in estimating both response time and alertness states.

4.2 Methods

In the following section, I will describe the study population, the phone based framework, and the data collection procedure.

4.2.1 Participants

For this study, my collaborators and I focused on college students. As I mentioned in the previous chapter, the college student population often suffers from serious circadian disruptions [217]. Given that alertness is a circadian process, data from this population provides a good opportunity to investigate the impact of such disruptions on cognitive performance. Also, this age group is the

largest and fastest growing users of phone [39]. It, thus, seems appropriate to recruit participants from the college student population to evaluate the proposed phone based framework.

For recruiting, my collaborators and I used public mailing lists, recruitment portals, and snowball sampling. Given that the framework is developed for the Android platform, the inclusion criteria required a participant to be an Android users. The study had 20 participants in total (7 males and 13 females) within 18–29 years age range and it lasted for 40 days.

4.2.2 Data

4.2.3 Momentary Assessments

For collecting data, I developed an Android application that can deliver ecological momentary assessment (EMA) to the participants. These EMAs included both subjective and objective assessment of individual's alertness. The phone application delivered the EMA four times a day at the start of 6-hour long windows (morning, afternoon, evening, and late night) through passive notifications. The participants could complete the EMA any time within this time window.

To ensure data quality, the participants were instructed to only begin an EMA session if they have 5 minutes without any significant distractions. If an EMA session was not completed within the time window, it was expired and a new notification was delivered to avoid redundant or temporally mislabeled data. This flexible study design helped to increase data coverage throughout

the day and maintain the data quality since the PVT task requires sustained attention.

For subjective assessment, the participants completed the Chalder Fatigue Scale [40] and Fatigue Visual Analogue Scales (VAS) [151] surveys on the phone. During an EMA session, the participants also provided information about their activities that might influence their level of alertness. Specifically, the EMA questionnaire asked about consuming caffeine [236], exercising [32], using nicotine [87], napping [32], consuming alcohol [97], eating [189], and loafing (e.g., cyberloafing) [232] during last hour.

Psychomotor Vigilance Test (PVT) Data

For objective assessment of alertness, the participants completed the Psychomotor Vigilance Test (PVT) on their phone. This study used the PVT version developed by Matthew et al. [112]. It shows a visual stimulus (a checkered box) on the screen at random intervals and the participants responded to the stimulus by touching the screen. Response times are measured in milliseconds, and a various statistical summaries of the response times have been used as indicator of alertness levels. Each PVT session in this study lasted for around 3 minutes. PVT was well-suited for this study given that it is immune to any learning effects over time.

Following previous work [112, 228], my collaborators and I decided to operationalize alertness as relative response time (RRT). For calculating RRT from PVT data, I first filtered out false starts — touch events that occurred even before the stimulus is shown to the participant. In this dataset, false start accounted for

2.85% of all touch events. Removing false starts resulted in only data points entered *after* showing the stimulus. Note that each PVT session contains multiple reaction time tests. I calculated the median response time ($MRT_{s,p}$) for each session s per participant p . I then identified and removed outlier sessions for each participant. For this, I considered any session with $MRT_{s,p}$ value falling outside ($\text{mean} \pm 2.5 \times SD$) as an outlier. In total, 6.2% of all sessions were removed as outliers. Next, I established the individual baseline for a participant p by taking the mean $MRT_{s,p}$. Then, the relative response time ($RRT_{s,p}$) of a given session s from a participant p is calculated as the percentage deviation from p 's individual baseline. That is, given a PVT session s for a participant p with $MRT_{s,p}$ as the median reaction time for that session, the corresponding relative response time is calculated as:

$$RRT_{s,p} = \left(1 - \frac{MRT_{s,p}}{MMRT_p}\right) * 100,$$

where $MMRT_p = \frac{1}{N} \sum_{i=1}^N MRT_{i,p}$, is the mean MRT averaged across all N sessions from a participant p . Positive value of RRT indicates increased alertness, and negative values indicate decreased alertness¹.

Given that PVT takes around 3-minutes and can be burdensome on the participant, it was sequenced at the end of the EMA session to prevent any PVT-induced fatigue on the subjective assessments. During the onboard process, my collaborators and I demonstrated how to perform the PVT to the participants with an emphasis on accuracy and speed when reacting to visual stimuli.

¹I have made the code for computing and analyzing RRT available at <https://github.com/saeed-abdullah/alertness-ubicomp-2016>

4.2.4 Sleep and Chronotype Information

Given that alertness is a circadian process, I also investigated the relationship of alertness with sleep, chronotype, and circadian stability. For this, participants completed a sleep journal every day answering questions about bed time, minutes to fall asleep, number of wakeups during the night, wake time, and total sleep duration. The phone application sent a notification at 10:30AM every day to remind the participants to complete their sleep journals.

I followed MCTQ to calculate social jet lag and chronotype of each individual using the collected sleep information. As mentioned earlier, MCTQ uses “corrected” mid-sleep point (MSF_{SC}) as chronotype. To calculate MSF_{SC} , it takes the differences in sleep pattern over work days and free days into consideration (i.e., the sleep debts over workdays and oversleep during free days to compensate for the accumulated sleep debt). For calculating MSF_{SC} , I used weekdays as work days and weekends as free days. Given the age range, the participants are mostly late types compared to the broader population. This is consistent with prior work on a similar age group [3]. Taking the overall distribution of this population, for this study I considered anyone with $MSF_{SC} < 5:00$ AM as “early” types.

I also used MSF_{SC} to calculate *internal time*. While external time (ExT) represents time elapsed since midnight (i.e., 00:00), “internal time” represents time elapsed since the *biological midnight* (i.e., MSF_{SC}). That is, internal time is calculated as:

$$InT = ExT - MSF_{SC}.$$

InT represents a corrected measure of time by factoring an individual's chronotype into consideration. Internal time has been used in previous studies to identify the relationship between alertness and circadian rhythm [228].

4.2.5 Phone Usage Data

The installed application also collected information about phone usage. In particular, it monitored screen on-off events and patterns of application use. These data were stored in the phone and periodically uploaded to a secure server.

Compensation was based on the duration of participation (\$5 for each week), sleep journal completion (\$0.50 for each entry), and the number of completed EMA assessments (\$0.20 for each entry). The Cornell University Institutional Review Board (IRB) approved all procedures.

4.3 Findings

During the study, the participants completed an average of 2.52 (SD: 0.79) subjective and 2.05 (SD: 0.87) objective (PVT) alertness assessments per day. They also completed 2.46 (SD: 0.8) surveys about stimulant intakes on average per day. The average EMA compliance rate across all participants was 63.7%. Excluding late night sessions — when participants are most likely to be asleep — the average compliance rate rose to 79.9%. The sleep journal had similar compliance rate — 72.8% across all participants.

In the following section, I will describe the findings from these data sets.

First, I will focus on replicating and extending previous lab based findings. While a number of recent studies have looked into the rhythms of alertness, this is the first study that collected alertness data in the real world setup for continuous period of time. As such, this dataset provides a unique opportunity to see whether the previous findings can be replicated outside of the controlled environment of a lab. Based on these findings, then I will describe a novel method for assessing alertness levels leveraging phone usage and behavioral features.

4.3.1 Replicating and Extending Extant Lab-Based Findings

Influence of Chronotype and Time-of-day on Alertness

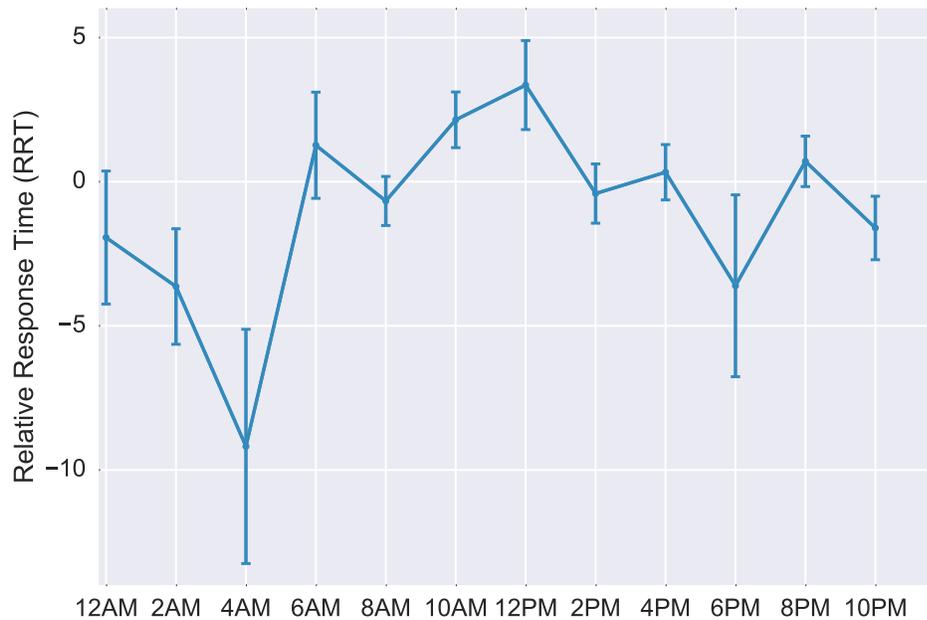


Figure 4.1: Relative response time (RRT) with standard error of mean (SEM) over the day. Positive and negative values of RRT indicate alertness higher and lower than individual baseline, respectively.

To compare the pattern of alertness over time, I aggregated the RRT values from each participant over 2-hour bins. Similar methods have been used by previous lab based studies to compare alertness trends over time [152, 228]. Figure 4.1 shows the patterns of RRT change in this dataset. We can see that there is a short-term and local dip around 2:00PM. This is consistent with a well-known biological phenomenon called “mid-day dip” which coincides with a dip in core temperature and overall cognitive performance [36, 152]. There are also noticeable dips around 4:00AM and 6:00PM. This trend is also consistent with previous lab based findings [223, 228]. However, in this dataset, these patterns are time-shifted (i.e., these dips happen later in the day compared to previous lab based studies). It might be due to the fact that the participants in this study have significantly later chronotypes than those from past studies (e.g., the average MSF_{SC} in this study is $05 : 56 \pm 0.94$ hr compared to $05 : 19 \pm 1.75$ hr in Vetter et al. [228]). The RRT improves during the late evening and night further reflecting the later chronotypes of this study population.

I also investigated how the alertness pattern changes across early ($MSF_{SC} < 5:00$ AM) and late ($MSF_{SC} \geq 5:00$ AM) types. For this, I calculated the difference in median² RRT between these two groups over a given period of time and normalize it by average daily change in RRT. As shown in Figure 4.2, the alertness pattern between early and late types is very different. Specifically, the early types are more alert during the morning (i.e., 17% higher RRT compared to late types). However, their performance worsens during late evening and night (i.e., 15% lower RRT compared to late types). The ANOVA test shows that the response time between early and late types across these time periods are significantly different with $F(5, 964) = 5.36, p < 0.001$.

²Using mean RRT for this comparison results in similar findings.

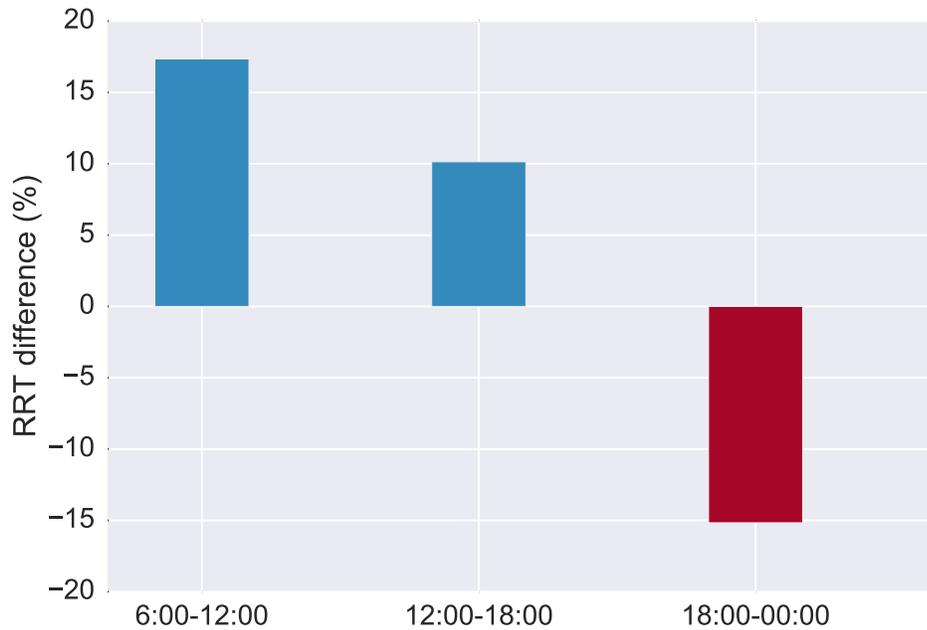


Figure 4.2: RRT of early chronotypes compared to late chronotypes across the day. Blue and red colors indicate higher RRT for early and late types, respectively. In the morning, early chronotypes display much higher alertness than late types, while the opposite is observed later in the day. Response time difference between early and late types across the day is also statistically significant: $F(5, 964) = 5.36, p < 0.001$.

Now, these comparisons are done over external time (*ExT*). As mentioned earlier, previous studies have also compared alertness pattern over internal time (*InT*) to better understand the relationship between alertness and chronotype. The change of RRT over *InT* in this dataset is shown in Figure 4.3. Similar to prior studies [228], we see that there is a local peak in RRT around 2:00 PM (i.e., 14 hours since the biological midnight, MSF_{SC}). However, the overall trend of alertness change is different compared to the findings from [228] (e.g., the dip in 12:00 PM). I think these differences might result from the fact that this study population have significantly later chronotypes. Future study will do well to further investigate the pattern of alertness change of late types across internal

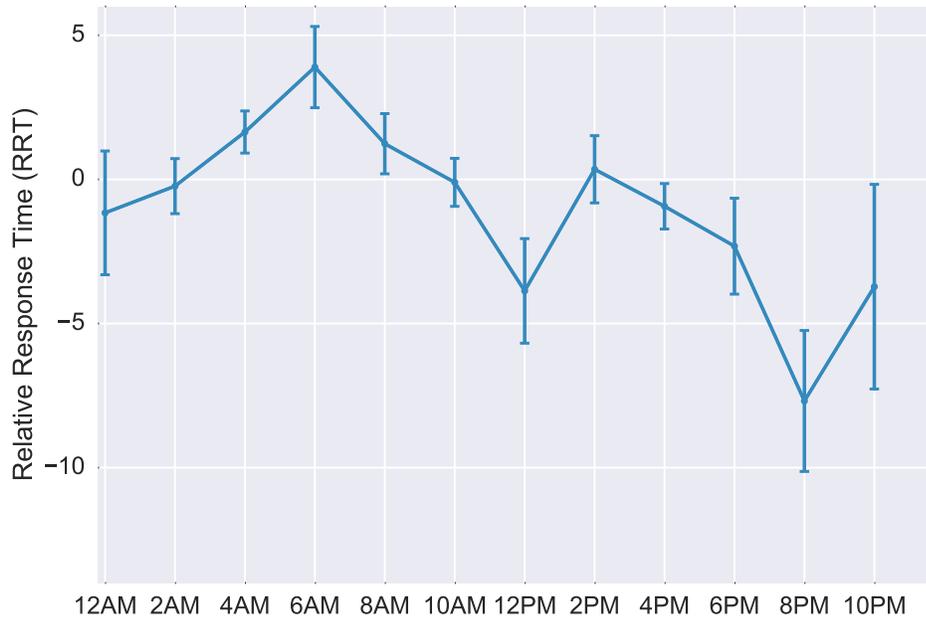


Figure 4.3: RRT with standard error of mean (SEM) with internal time on X-axis

time.

Daylight Savings Time (DST) Impacts Alertness

Daylight Savings Time (DST) is known to cause circadian disruptions. This twice in a year change of external time can result in effects similar to jet lag [111]. More importantly, DST is observed in more than 70 countries and as such, 1.6 billion people worldwide gets impacted by this shift of external time every year. While debates about DST have mostly been centered around economy, a number of recent studies have investigated how the resultant circadian disruptions negatively impact our health and well-being. For example, Lahti et al. [124] found that DST results in fragmented rest and wake activities; Gaski et al. [81] reported that DST can have a detrimental effect on academic performance; Barnes et al. [19] found that DST time changes can increase the risk for work-

place injuries. However, none of these studies investigated how the circadian disruptions resulting from DST might affect alertness patterns as computed by objective tests like PVT. In the following section, I will describe the impact of DST on alertness based on the data from this study.



Figure 4.4: RRT change for early and late chronotypes before and after Spring DST. While the DST transition negatively affects both early and late types, late types suffer more. The difference in RRT before and after DST for both early ($t = -2.37, p = 0.02$) and late ($t = -2.52, p = 0.01$) types is statistically significant.

I will focus on Spring DST here given that the clock forwarding can result in significant circadian disruptions particularly affecting late types for days [111]. Also, the study spanned 23 days before and 17 days after the Spring DST, providing adequate data to compare the alertness trends before and after the DST. For this, I calculated the difference in median³ RRT in morning sessions before and after DST; and then I normalized it by average daily change in RRT. The resultant difference in morning RRT is shown in Figure 4.4. In this dataset,

³Using mean RRT also resulted in similar findings.

RRT in the morning sessions drops after DST for both early (-4.7%) and late (-11.8%) types; however, the late types suffer significantly more than the early types. These differences in RRT are also statistically significant for both early ($t = -2.37, p = 0.02$) and late ($t = -2.52, p = 0.01$) types.

Impact of Sleep

Sleep affects a wide range of cognitive functionalities [85]. Based on the data from this study, I also investigated how sleep might affect one's alertness levels. In particular, I was interested to see how sleep deprivation on previous night might affect alertness on the next day. Given that different individuals have different sleep patterns and requirements, it might not be appropriate to compare alertness and sleep duration directly. Instead, following prior work [228], I used sleep duration *relative* to one's sleep need for this analysis. Sleep need is calculated as the average sleep duration across free days (SD_F) and workdays (SD_W): $SN = (SD_W * 5 + SD_F * 2) / 7$. Given the sleep duration SD , the relative sleep need is then calculated as $SD_r = SD / SN$. Similar to chronotype calculation, I considered weekdays as workdays and weekends as free days for calculating sleep need.

Based on data from shift-workers, Vetter et al. [228] found that reduced sleep duration negatively impact alertness in the morning. However, in this dataset, there was no significant difference in morning RRT between adequate ($SD_r \geq 1.0$) and inadequate sleep ($SD_r < 1.0$) with $t = -1.45, p = 0.14$. This might be due to the fact the participants here are late types compared to the previous study and as such, they are more likely to exhibit low alertness in the morning regardless of their sleep last night. However, for days after Spring DST, ad-

equate sleep at night does seem to improve alertness in the morning. That is, after Spring DST, the median RRT after nights with inadequate sleep ($SD_r < 1.0$) falls by 17.49% with $t = -2.11, p = 0.03$. In other words, sleep deprivation after Spring DST might have a more negative impact on one's alertness given the already disrupted circadian system.

Stimulant Use Impacts Alertness

Positive Stimulants	Negative Stimulants
Caffeine Consumption, Exercising, Napping, Nicotine Intake	Alcohol Consumption, Food Intake, Loafing

Table 4.1: Groupings of stimulants based on how they impact alertness

Based on the data from this study, I also investigated how stimulant intake in our day to day life might impact alertness levels. As mentioned earlier, during the daily EMAs, the participants reported activities in past hour that might influence alertness. I categorized these activities into positive (stimulants that might improve alertness) and negative (stimulants that might diminish alertness) groups as shown in Table 4.1. In this dataset, there was a short-term impact of positive and negative stimulants on alertness. That is, the median RRT increased by 5.08% across all participants after the reported use of positive stimulants, while negative stimulants resulted in -1.37% drop. This is also statistically significant with $t = 2.21, p = 0.03$.

Comparing Self-assessment to Objective Assessment of Alertness

As I have mentioned earlier, objective assessment of alertness using PVT or physiological signals poses a number of logistical challenges. The subjective as-

assessments, on the other hand, are significantly easier to deploy. However, there have been conflicting findings regarding the validity of subjective assessments [175]. Using data from this study, I investigated the reliability of self-assessment surveys in comparison to PVT assessment.

Self-Assessed Variable	RRT difference between self-assessed alertness states
Energy	$t = 2.06, p = 0.04$
Concentration	$t = 2.76, p = 0.005$
Tiredness	$t = -2.1, p = 0.03$

Table 4.2: RRT differs significantly between self-assessed high and low alertness, indicating fatigued individuals are usually aware of their reduced capability.

During each EMA session, the participants rated their energy levels, concentration and tiredness using a 5-point Likert scale. These ratings are highly correlated with each other indicating internal validity of the responses from the participants. I categorized these subjective assessment values into high and low alertness groups using median as a threshold (e.g., all energy ratings below the median falls into the low alertness group). I then compared the RRT values of these two groups [58]. As Table 4.2 shows, RRT differs significantly between subjective high and low alertness scores. In other words, self-assessments are consistent with objective assessment in this dataset and the participants were aware of their reduced cognitive ability. This finding is similar to a number of previous studies [16, 58]. However, it should be noted that the comparison in this study is done in a somewhat coarse level (i.e., “are you tired?” compared to “*how* tired are you?”). Future work should extend these findings to a more granular subjective assessments (e.g., determining if a participant can correctly recognize when one is too tired to drive safely).

4.3.2 Passive assessment of alertness

In this section, I will focus on developing a novel method for continuously and passively assessing alertness that can be deployed at scale over a long period of time. I will first describe how phone usage patterns can indicate different levels of alertness and then I will leverage these features for predicting alertness levels.

Phone Usage Reflects Alertness States

There has been a number of recent studies that investigated how technology use changes reflecting different levels of attention and boredom [135, 178]. However, none of these studies has focused on alertness, which is a critical component of our cognitive performance. In this section, I will focus on identifying how phone use patterns might reflect different levels of alertness. Given phone use is deeply embedded in our day-to-day life, it can provide unique insights in our behavior and cognitive performance.

Based on previous research, I hypothesized that frequent and prolonged duration of phone use would indicate a decreased ability to focus and low levels of alertness. I defined two metrics to identify such behaviors: i) *burstiness*: the total number of phone use sessions in an hour, where each phone use session is defined by subsequent screen on and off events, ii) *total duration* of phone use over a given period of time. To identify how these metrics differ across alertness levels, I categorized PVT sessions into high alertness ($RRT \geq 0$) and low alertness ($RRT < 0$) groups. The distribution of these groups is balanced in this dataset with 56.98% and 43.02% PVT sessions marked as high and low sessions, respectively.

In this dataset, the mean burstiness increases during high alertness sessions. That is, the mean burstiness during high alertness sessions is 9.07 (± 0.45) compared to 7.9 (± 0.30) during low alertness sessions. This difference is also statistically significant with $t = 2.14$, $p = 0.03$. In other words, the participants initiated more phone sessions when they were more alert. However, the duration of phone use per burst shows an opposite trend. During sessions with high alertness, the mean duration of phone use is 116.37 (± 5.5) seconds per burst compared to 123.47 (± 6.62) seconds per burst during sessions with low alertness. In other words, during the period of high alertness, the participants checked their phones more frequently but those phone sessions lasted for a short period of time. However, during period of low alertness, they engaged in prolonged use of phones. To further validate these findings, I investigated the patterns of short phone usage sessions. For this, I considered any phone use session lasting less than 30 seconds as a short session. As expected, the mean frequency of short phone use sessions was 20% higher during periods with high alertness with $t = 1.97$, $p < 0.05$.

These findings show that phone use patterns can be a useful indicator to differentiate between high and low alertness states. In the following sections, I will leverage this data for passive assessment of alertness.

Predicting Alertness States

For modeling and predicting alertness states, I used both behavioral and contextual features. Specifically, I used the following features: local time, internal time, sleep duration, relative sleep need (SD_r), stimulant intake (positive vs negative), subjective assessment scores (i.e., energy, concentration, and tiredness),

and phone usage trends (burstiness, duration, mean time between consecutive sessions, and number of short sessions with duration less than 30 seconds). These features reflect findings from this dataset as well as extant literature about behavioral and biological cues that might impact alertness.

For modeling, I first focused on response time. Sleep and cognitive research have used response time from PVT in a number of different ways. The ability to model and predict response time, thus, is useful to a wider range of community. Predicting response time is a regression task. For this, I used linear regression model with Stochastic Gradient Descent (SGD) and the Huber loss function [102]. I decided to use SGD given its first convergence speed and scalability enabling learning with large amount data. Huber loss function also ensured that these models are robust against outliers.

Given that SGD can be sensitive to feature scaling, I normalized all features to have zero mean and unit variance. Following prior work [29], I randomly shuffled the training data after each epoch to increase convergence speed and performance. For selecting model parameters (e.g., hyper-parameter α), I used 10% of randomly selected training data. The best performing model for this dataset had L1 norm as the regularization term with hyper-parameter $\alpha = 10^{-7}$ and learning rate $\gamma_t = (\alpha \cdot t)^{-1}$. Based on 10-fold cross-validation data, I found that the generalized model predicted response time with a root-mean-square-error (RMSE) of 83.81 milliseconds. I also trained personalized models using data from each participant. As expected, use of personalized models further improved the performance — the average RMSE from 10-fold cross validation dropped to 80.64 milliseconds over all participants.

The performance of these models are very encouraging. To put the accuracy

of these models into context: a response time higher than 500 milliseconds often is considered as a *lapse* — a standard measure of impaired cognitive ability in sleep and cognitive performance research [20, 133]. Given the much higher granularity achieved by these models (80.64 milliseconds on average), I think they can be successfully deployed instead of actually using PVT.

Beyond response time, I also focused on modeling relative response time (RRT) that I have used in this study. Given that one of the major goal of this study is to be able to assess alertness unobtrusively, I decided to only use passively sensed features. That is, for this modeling task, I used local time, internal time, sleep duration, relative sleep need, phone usage burstiness, mean duration of phone usage sessions, average time between successive phone usage sessions, and frequency of short (less than 30 seconds) phone usage sessions as features. While the participants in this study used sleep journals, sleep and chronotype information can be reliably inferred from phone usage patterns as I have described earlier [3].

Following the success of previous modeling task, I decided to use linear regression model with Stochastic Gradient Descent (SGD) and the Huber loss function for predicting RRT as well. I also standardized features to have zero mean and standard variance to avoid any feature scaling issues. For choosing model parameters, I used randomly selected 10% of training data. The best performing model for this task used L1 norm as the regularization term with $\alpha = 10^{-8}$ and learning rate $\gamma_t = \gamma_0 \cdot t^{-\frac{1}{4}}$ where the initial rate, $\gamma_0 = 0.01$. Based on 10-fold cross validation data, the generalized model predicted RRT with RMSE 11.39% across all participants. I also trained individualized models that further improved accuracy with RMSE 10.87% across all participants. In other words,

these models can detect around 11% deviation from individual baseline. Considering the average daily peak-to-peak change in RRT is three times higher than (29.4%), these models are indeed reasonably accurate and reaffirms the feasibility of replacing PVT.

Rank	Response Time	RRT
1	Energy rating	Internal time
2	Internal time	Avg. time between phone usage sessions
3	Stimulant intake	Short session frequency
4	Avg. time between phone usage sessions	Phone usage duration
5	Concentration rating	Relative sleep need

Table 4.3: Top ranking feature groups for modeling response time and RRT based on relative feature elimination (RFE)

To get better insights into these models, I also evaluated the importance of each feature. To perform feature ranking, I used recursive feature elimination (RFE) [90]. It starts with the complete feature set and removes the least contributing feature at each step. That is, at successive step of RFE, I trained a model using the full dataset and the feature with least absolute weight in the model is removed. This procedure continues when there is only one feature left and it is considered as the most important feature. Table 4.3 shows the top ranking features for predicting response time and RRT. These features show consistency with both previous literature and the prior findings from this dataset. For example, extant literature identified chronotype, internal time, relative sleep need, and stimulant use to be important factors related to alertness [58, 205, 228]. Moreover, these top ranking features also indicate that behavioral features like phone usage patterns can be very useful when it comes to model in-situ alertness.

4.4 Discussion

Alertness is a complex circadian process that is known to vary significantly across the time of a day. A number of biological (e.g., chronotype) and behavioral (e.g., sleep duration and stimulant intakes) factors can influence alertness. Being able to understand and predict such changes in alertness can have significant impact across a wide range of domains given that alertness is a crucial aspect of our cognitive performance. While there have been a number of recent studies focusing on a better understanding of alertness, most of these studies are conducted within the controlled environment of a lab. In this chapter, I focused to move beyond the laboratory environment by collecting ecologically valid data in the wild over an extended period of time.

Specifically, I helped to deploy a phone based method for assessing alertness levels of 20 participants over 40 days. Based on these data, I first focused on replicating and extending previous lab based studies. Replication is important to advance scientific knowledge and it also helps validate the proposed phone based method. The replicated findings from this study are mostly consistent with previous lab based studies. Specifically, I found that alertness reflects a circadian pattern — it drops during the early morning, peaks during noon and late evening. Alertness also shows “mid-day” dip phenomenon with a significant drop around 2:00 PM in this dataset. The alertness pattern depends on chronotype as well — early types are more alert in the morning while late types perform better during evening. Given that the time transition due to Daylight Savings Time (DST) is known to cause circadian disruptions, I also investigated how alertness pattern changed after Spring DST. In this dataset, DST results in significant drop in morning alertness, with late types being affected more nega-

tively. I also found that sleep deprivation after Spring DST had a more negative impact on alertness. During the study, the participants also tracked stimulant intakes that might impact alertness. Based on this data, I found that stimulants can have both positive (e.g., caffeine) and negative (e.g., alcohol) impact on alertness.

Beyond replicating and extending previous findings, one of the main goals of the study was to develop a novel method for passively and unobtrusively assess alertness. For this, I developed and trained models using contextual, behavioral and phone usage features. These models achieve high accuracy for predicting both response time and RRT. Compared to PVT tests, these models are significantly low-cost and induce no user burden. As such, these models can be deployed over a long period of time across a large population potentially enabling unique insights into alertness patterns.

In the following section, I will describe the limitations of the study both in terms of method and findings. I will also suggest ways to address these issues in future work. Further, I will contextualize these findings from the broader perspective of Circadian Computing.

4.4.1 Limitations

The findings of the study are based on a small set of participants. Moreover, these participants are all college students and have mostly late chronotypes. It will be useful to extend these findings to a broader and more diverse population. In particular, future work should validate these findings across a wider age range given that older population also show a time-of-day effect in cogni-

tive performance [205]. Also, given that the phone use patterns might be significantly different across different age ranges and the alertness prediction models here make extensive use of these patterns, it is necessary to confirm the accuracy of these models based on data from a more diverse population.

During the analysis, I used a coarse grouping of chronotypes (i.e., early and late types). While such grouping makes sense for this small population with mostly late chronotypes, it would be useful to extend these findings over a continuous distribution of chronotypes. For example, I compared the alertness of early and late types across the day. Future work will do well to investigate if this pattern is consistent across the distribution of chronotypes.

In this study, I only considered phone use for assessing alertness. However, as our digital lives now span multiple devices simultaneously, it will be useful to broaden the proposed method to consider overall technology use. Such a framework could consider computer use, tablets, and TV watching habits to get more insights into behaviors corresponding to high and low alertness states.

Also, given that these models use phone usage as features, they might not be appropriate in some occupational contexts. For example, during a long-haul drive, a driver will not use one's phone (hopefully). In such cases, the accuracy of the model might be low. However, these models can still infer a generalized trend of alertness (e.g., when a dip is expected for a given person) using overall phone use patterns as features.

4.4.2 Implications and Opportunities

Alertness is a crucial process when it comes to cognitive performance. As a result, being able to predict the changes in alertness can have significant impact across a number of domains. In this section, I will detail the implications of these findings and future opportunities. In particular, I will discuss how these findings can lay the foundation for a new class of technologies.

Sensing and Design Implications

As mentioned before, the phone usage based method might not be applicable in some occupational setups. In such cases, it might be useful to have a multi-modal framework that can combine different data streams to infer alertness levels. In particular, the models presented here can be extended and complemented by additional physiological sensing. For example, during a long-haul drive, instead of solely relying on phone usage features, we can use eye movement and pupil information as features. Such multi-modal frameworks would increase the reliability and overall data coverage.

However, the existing physiological sensors are quite costly and obtrusive to be used in the wild. There is thus a need to develop wearables and hardware that will allow collecting physiological data without inducing user burden. In particular, future work should investigate if phone cameras can be used to reliably track pupil size and gaze information for inferring alertness. Such a system can do opportunistic sensing without any user involvement. For example, it can take a picture each time the user unlocks the phone using front-camera and then extract pupil information from these images for inferring levels of alertness.

These passive sensing of alertness can enhance the current scope of personalized and context-aware computing. In particular, I believe that the findings presented here broadens the scope of Circadian Computing. Being able to predict alertness enables future technologies to be more adaptive to our innate biological rhythms and play to our biological strengths (and weaknesses). In the following section, I will describe how these findings can enable circadian-aware technologies across a number of different application domains.

Application in Education

Given the relationship between alertness, memory, and learning, a passive sensing framework for alertness can be invaluable in educational setup. A number of recent studies have found that learning schedules for high school and college students often run against their natural rhythms of alertness [113]. In other words, these students are being taught when they are least able to focus and concentrate. The method and findings described in this study can address these issues and provide students an optimized scheduled for individualized learning.

For example, this assessment method enables a large scale and ecologically valid data about alertness trends of students. Such collection of data could be used for advocating scheduling reform in the education system. At the class-levels, this data can be used for better scheduling that accounts for both early and late types. It can also help in collaborative projects by allowing selection of group members based on their alertness profiles.

Beyond the boundaries of institutes, these findings can be used for empow-

ering individuals. That is, we can extend these methods to develop tools that will help individuals to be more aware of their idiosyncratic patterns. For example, it can help students to make more informed decisions about study and class schedules that would better align with their own personal alertness trends.

Application in Scheduling

Beyond students, passive sensing of one's alertness pattern will enable novel ways of scheduling tasks and collaborative meetings. I think there is an exciting opportunity to design and develop a novel calendar system that better reflects the dynamic and oscillating nature of alertness and cognitive performance. For example, a calendar system can intelligently suggest meeting slots by taking the alertness patterns of participants into consideration. Also, given the rise of distributed workplace, such systems can better facilitate remote team management. For example, it can pair team members across time zones but with matching alertness trends (e.g., a late chronotype from the East Coast with a very early type from the West Coast). Similarly, such a framework could also suggest appropriate timeslots for a given task by considering the task difficulty and one's alertness patterns into consideration. For example, it will recommend less cognitively demanding tasks when a dip in alertness is expected (e.g., doing laundry during when the alertness level is low).

Application in Accident Prevention

Being able to infer in-situ alertness in real-time could be particularly useful in preventing occupational accidents. As I have mentioned earlier, fatigue causes

a significant number of road and industrial accidents. A framework that can monitor user's alertness in these occupational setups can help to significantly reduce the risk of such accidents. In particular, knowing one's alertness trends can be used to preemptively avoid risky situations. For example, if we know that a user tends to have low alertness levels during 4:00–6:00 PM, then this information can be used to recommend avoiding driving during this time period. Moreover, having access to real-time data about one's alertness will enable actionable and potentially effective interventions and countermeasures. Overall, I believe, these findings presented here potentially enable novel intervention systems for preventing occupational accidents resulting from fatigue.

4.5 Summary

Alertness is a circadian process that varies considerably over the time of a day. Given the critical role alertness plays in our cognitive performance, there has been significant research effort to identify factors that might influence it. However, most of these studies are done within the controlled environment of a lab. As a result, there is a need to extend and validate these findings using alertness data collected in the wild.

Towards this goal, in this chapter, I have presented a phone based method for collecting in-situ alertness data. Based on a study with 20 participants over 40 days, I have replicated and extended a number of previous lab based findings. In particular, I have found that alertness oscillates considerably over the course of a day and this pattern varies across early and late types. I have also showed that Daylight Savings Time (DST) has a negative impact on alertness with a

more pronounced effect on late types; and stimulant intake can have a short-term effect (both positive and negative) on alertness.

Moreover, the existing methods for assessing alertness often are not feasible to deploy at a scale over a long period of time. For example, PVT requires 3–10 minutes of undivided attention and thus results in significant user burden. To address these issues, I also focused on developing a passive and unobtrusive framework that can continuously monitor alertness. For this, I developed both generalized and personalized models that can accurately predict alertness levels by leveraging phone usage patterns along with other behavioral and contextual features.

Given the critical role of alertness in almost every aspects of our lives, these finding potentially have a far-reaching positive impact. In particular, passive and continuous sensing of alertness potentially enables a new class of circadian-aware technologies that can reduce occupational accidents, optimize daily schedules in accordance with our biological rhythms, and improve learning outcomes. These findings can significantly extend the current scope of Circadian Computing.

CHAPTER 5

CIRCADIAN DISRUPTION AND BIPOLAR DISORDER

Circadian disruption has been associated with a wide range of mental health issues including alcohol and substance abuses [192, 96], anxiety disorder [171], schizophrenia [241], and bipolar disorder [94]. While past research attributed such disruptions to the pathology of mental illnesses, recent studies hypothesize that circadian systems might be more directly involved in disease etiology [143, 126]. As such, Circadian Computing can provide a critical support in mental health care by developing appropriate sensing and intervention tools. In this chapter, I will describe the design, development, and deployment of such a circadian-aware tool that focuses on bipolar disorder (BD).

5.1 Introduction

Bipolar disorder is a serious mental illness. It can result in recurring episodes of depression and mania. Episodes of depression are characterized by low mood and reduced energy symptoms. The manic episodes, on the other hand, are marked by elated mood along with irritability and a reduced need of sleep. Patients can also be in hypomania — periods with symptoms similar to manic phase but less severe [177].

Bipolar disorder is one of the leading causes of disability worldwide [168]. It affects around 2% of the world population [144]. It manifests in poor functioning and psychotic symptoms [109], which can lead to low quality of life [201] and stigma [226]. BD is associated with high suicide risks [163]. It results in significant societal cost as well. In 2009, the direct and indirect cost associated

with bipolar disorder was \$151 billion in United States alone [56].

The relationship between bipolar disorder and circadian disruptions is well-established. A decreased need for sleep is considered to be a strong marker of a manic state [179]. A number of studies have also reported associations between bipolar disorder and “clock genes” (genes influencing our circadian rhythms) [65, 22]. Ehlers et al. [60] hypothesized that depression can result from disruptions in social zeitgebers — personal relationships and social demands that serve to stabilize our innate biological rhythms. In other words, changes in daily routine and social relationships can disrupt our endogenous circadian rhythms leading to mood symptoms and, in vulnerable individuals, manic or depressive episodes.

There is no cure for bipolar disorder. Clinical interventions, therefore, aims to reduce risk of relapse through longitudinal and effective management of symptoms. Given the relationship between circadian disruptions and bipolar disorder, recent work has focused on maintaining circadian stability of patients [82]. For example, Interpersonal Social Rhythm Therapy (IPSRT) [78] is a psycho-social theory that specifically focuses on maintaining social rhythms of patients for better circadian stability. That is, IPSRT aims to establish regularity in daily and social routines of patients with bipolar disorder.

To keep track of stability in daily and social routines, IPSRT uses Social Rhythm Metric (SRM). It is a 5-item scale as shown in Figure 5.1. SRM allows the patients to track the following daily events: getting out of bed, making first contact with another person, start working, having dinner, and going to bed. Beyond monitoring the timing of these events, it also tracks the “social” level of these events (using the “people” column). Moreover, it also monitors the daily

mood and energy information using a scale ranging from -5 (very low) to +5 (very high).

SRM II-5

Directions: Date (week of): Nov 18 – 24, 2015

- Write the ideal target time you would like to do these daily activities.
- Record the time you actually did the activity each day.
- Record the people involved in the activity: 0 = Alone; 1 = Others present; 2 = Others actively involved; 3 = Others very stimulating

Activity	Target Time	Sunday		Monday		Tuesday		Wednesday		Thursday		Friday		Saturday	
		Time	People	Time	People	Time	People	Time	People	Time	People	Time	People	Time	People
Out of Bed	7:00 am	9:30 am	0	8:30 am	0	7:30 am	0	7:30 am	0	7:15 am	0	7:40 am	0	8:30 am	0
First contact with other person	8:00 am	9:30 am	1	9:30 am	1	8:30 am	1	8:40 am	1	8:15 am	1	8:40 am	1	9:15 am	1
Start work/school/volunteer/family care	9:30 am	10:30 am	0	10:15 am	2	9:30 am	1	9:50 am	2	9:15 am	0	10:40 am	1	11:30 am	0
Dinner	9:00 pm	11:30 pm	2	9:30 pm	0	9:50 pm	1	9:00 pm	0	9:15 pm	0	10:20 pm	1	9:30 pm	1
To Bed	11:30 pm	12:30 am	0	11:30 pm	0	11:50 pm	0	11:40 pm	0	12:15 am	0	12:40 am	0	12:30 am	0
Rate MOOD each day from -5 to +5 -5 = Very depressed +5 = very elated			+ 1		- 2		- 1		0		- 1		+2		+2

Figure 5.1: Sample paper-based Social Rhythm Metric (SRM) form that is used as part of Interpersonal Social Rhythm Therapy (IPSRT).

The use of SRM through IPSRT has been shown to be an effective intervention tool. Based on data from 175 patients over 2 years, Frank et al. [76] found that IPSRT results in symptomatic improvements and reduces the risk of relapse in patients [76]. Miklowitz et al. [146] also found that the use of IPSRT improves overall functioning and life satisfaction in bipolar disorder. Furthermore, IPSRT has been shown to improve occupational functioning as well [77].

While the SRM is effective in assessment and stabilization of social rhythms, its paper-and-pencil based format has some serious limitations as a clinical tool. For patients, longitudinal self-tracking using this tool is difficult. Even well-intentioned patients can forget to complete it. Moreover, given the nature of bipolar disorder, in certain stages of illness retrospective and momentary recall can be particularly challenging for patients. Also, the current tool is not good

for summarizing the data. It also does not provide any real-time feedback to patients or clinicians. As a result, they can miss the time-sensitive warning signs, which might lead to relapse onset.

A number of recent studies have focused on addressing these issues with self-assessment. For example, Agnes et al. [89] used phone sensors to differentiate between manic and depressive phases; Frost et al. [80] also used phones for collecting behavioral data that can provide better insights into disease trajectory. However, none of these studies focus on understanding and maintaining circadian stability. Given that circadian stability is central to the well-being of individuals with bipolar disorder, being able to passively and unobtrusively track it can have significant positive impacts.

To address these issues, in this chapter, I will focus on developing *MoodRhythm* — a phone based technology to help patients with BD maintain circadian stability. My contributions for this work are:

- I helped to develop the MoodRhythm application that allows patients to complete SRM scores in their phones. It also uses the phone sensors for continuous and unobtrusive monitoring of behavioral and contextual features that might indicate circadian disruptions. Furthermore, I developed the back-end framework for storing data in a reliable and secure manner.
- My collaborators and I deployed MoodRhythm among 7 patients for 4 weeks. Based on this data, I modeled and trained machine learning algorithms that can accurately assess SRM score and stability of the patients. This passive and unobtrusive assessment method can be used for longitudinal monitoring of symptoms and identifying early warning signs, which could enable preemptive care in mental illness.

5.2 Method

5.2.1 Participants

The study participants were recruited through the Depression and Manic-Depression Prevention Program at Western Psychiatric Institute and Clinic. For recruitment, my collaborators sent information letters to potential participants stating project goals, expected duties, and time commitments. Interested patients could also contact the research staff to obtain more information regarding the study. The inclusion criteria required patients to be already participating in a treatment program. This ensured that the participants had a confirmed diagnosis of bipolar disorder. Participants were excluded if they had active suicidal ideation, which might require inpatient or intensive outpatient management. Participants unable or unwilling to comply with study procedures were also excluded. The Institutional Review Board (IRB) at the University of Pittsburgh approved all the procedures.

In total, nine patients participated in the study. However, one participant did not comply with the study protocol and there was data uploading issues from another participant. I filtered out data from these two participants before doing any analysis. The demography of the rest of the patients is shown in Table 5.1.

The study lasted for four weeks. My collaborators provided each participant an Android phone (Nexus 5) with MoodRhythm installed in it. They explained how to use the application to the patients during the onboarding interview. The participants completed an initial and a post-study questionnaire, along with an

Participant	Age	Gender
1	25–34	Female
2	55–64	Female
3	45–54	Male
4	25–34	Female
5	35–44	Male
6	25–34	Female
7	25–34	Female

Table 5.1: Demographic of the patients

interview at the end of the study. Given the extent of personal data collection, each patient was compensated with \$50 for a week of participation, \$25 for each completed questionnaire, and \$50 for the final interview. The compensation was not contingent on adherence to the study protocol.

5.2.2 MoodRhythm

Participants used the app MoodRhythm to track their social rhythms. It was developed through participatory design involving both patients and clinicians [139]. The app allowed patients to track 5 SRM items as shown in Figure 5.2. The patients could also add customized items for tracking. Following IPSRT, a participant could set daily target times for tracking these items. The application provides reminders and visual feedback to notify users about an upcoming activity. Specifically, if a participant completes an activity (e.g., having dinner) within 45 minutes of target time, then the bar to the left turns green. When this time window is about to elapse for an incomplete task, the bar turns yellow to remind user to complete it. If the user misses the target, then the bar turns red. In this way, MoodRhythm provides a glanceable feedback about the consistency

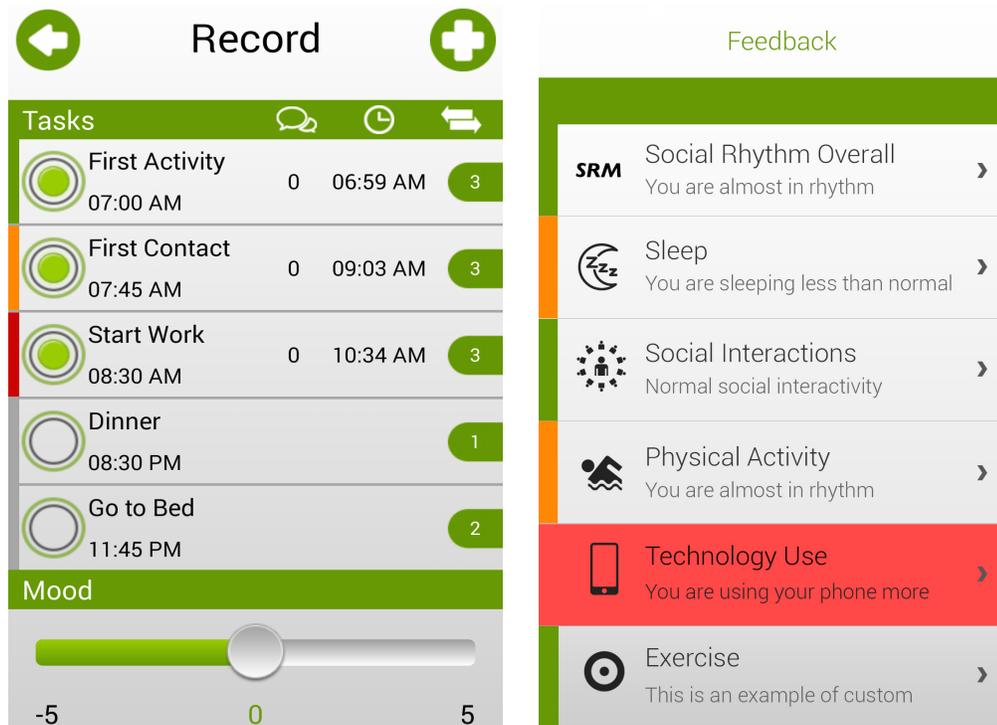


Figure 5.2: MoodRhythm application used by the patients with bipolar disorder.

in completing daily SRM items.

Participants could track their energy and mood levels as well using a scale from -5 (very low) to +5 (very high). MoodRhythm also had a daily journaling feature called “Notes”. The patients could use it to record additional information including medication intake or factors that affected mood and energy on a given day. It also allowed the patients to access SRM item completion data in past days. This provided an instantaneous feedback about overall adherence in the past. Furthermore, MoodRhythm had a visualization module for detailed feedback based on weekly SRM completion and sensor data. For example, it provided a line chart showing the trend in SRM score on a weekly basis. It also had some gamification elements to increase adherence. Specifically, it rewarded

participants by giving them badges after completing a milestone (see Figure 5.2).

As I mentioned above, one of the main goals of this study was to come up with an automated and unobtrusive way to track circadian and social rhythms stability. For this, MoodRhythm used phone sensors to keep track of behavioral and contextual features. In particular, it collected accelerometer, light, location, activity, and voice data. Given that audio data collection and processing can incur significant resource use, the app used an opportunistic sampling in which the phone's microphone was activated every 2 minutes. If any human speech was detected, then the microphone would be kept activated and the app would continue to collect audio data. Otherwise, the app stops collecting and processing audio data till next activation. This opportunistic sampling helped significantly in retaining phone charge — the application could do 16 hour of continuous sensing after a full recharge.

MoodRhythm used the collected audio data to identify presence of human conversation on the fly. Given that audio data can be privacy-sensitive, it did not store raw audio data. Instead, it processed the audio data to extract and store features like spectral content and regularity in real time. These features are useful for detecting presence of human conversation but inadequate to reconstruct any speech content [242]. Further, to filter out false positives (e.g., conversation from a TV program), it used energy intensity and distribution likelihood [182]. Based on this processed data, it is possible to infer the duration and frequency of conversations over a given period of time, along with speaking rates and variations in pitch. These features have been previously used to identify social isolation in older adults [182].

MoodRhythm collected information about activity (e.g., walking vs being sedentary). It kept track of location of the participants, which can be useful to identify daily behavioral patterns (e.g., total distance traveled). It also collected data about phone usage and communication patterns (including SMS and call logs). Data was stored in the phone and periodically uploaded to a secure server.

5.3 Findings

During the study, on average a participant recorded 36.5 (SD: 11.17) energy ratings, 46.12 (SD: 12.71) mood ratings, and 144.43 (SD: 43.1) activities including both SRM and customized items. Also, the average distance traveled per day by a participant was 8.34km (SD: 13.34). On average, the ratio of duration being sedentary to being active per day was 2.09. Also, the audio data contained on average 3.5 (SD: 3.67) hours of human conversation per day.

Features	Correlation
Cluster frequency	0.31***
Distance traveled	0.23***
Conversation frequency	0.25**
Non-sedentary duration	0.39***

Table 5.2: Correlation between sensor stream and trend of self-assessed energy scores computed as rolling average over 7 days (** : $p < .01$, *** : $p < .001$)

For inferring SRM scores, I processed the sensor data over each day to calculate the following daily features: number of location clusters visited, total distance traveled, frequency of human conversations, and duration of being non-sedentary. For calculating the number of location clusters visited by a par-

ticipant, I used the DBSCAN (Density-Based Spatial Clustering of Applications with Noise) algorithm [64]. It requires two parameters for identifying clusters: the maximum distance between two points in a given cluster and the minimum number of samples required to form a cluster. For this study, I used 0.5km as the maximum distance and minimum of 3 points for a cluster.

I selected these features because they are useful indicators for social and physical functioning. A number of recent studies focusing on assessing mental health status have also found these features to be important. For example, speech and conversation features have been used for identifying phases of bipolar disorder [89, 155, 225]. Similarly, previous studies have showed that physical activity [170], location [35], and mobility [88] information can be quite useful in this context. To further assess the usefulness of these features, I compared them with self-assessed energy and mood rating trends from the patients. To identify trends, I calculated the rolling average using a 7-day window. As shown in Table 5.2, these features are strongly correlated with self-assessed energy scores. The correlation with mood trend was similarly positive but it was not statistically significant. For example, the correlation between mood trend and conversation frequency was $r = 0.16, p = 0.06$; with non-sedentary duration, it was $r = 0.15, p = 0.08$.

Next, I focused on predicting SRM scores based on these features. SRM score is a continuous value that ranges from 0 to 7, where higher value indicates better stability. In the paper-and-pencil format, SRM is calculated using data from non-overlapping weeks. However, given the much higher granularity of the sensed data in this study, I calculated SRM scores using a rolling window of 7 days. For modeling, I used Support Vector Regression [211]. Based on 10-fold

cross validation, the generalized model can predict SRM score with an average root-mean-square-error (RMSE) of 1.40. I also trained individualized models using data from each participant. These individualized models further improved the performance with an average RMSE of 0.92. These low RMSE values indicate that these models can accurately predict the SRM scores.

Given these encouraging results, I further investigated whether it is possible to infer stability directly from these features. From a large study with 1,249 healthy participants, Tienoven et al. [224] reported that the mean SRM score for a healthy population is close to 3.5 — halfway between very unstable (SRM score 0) and perfectly stable (SRM score 7). For labeling stability status in this dataset, I used the same threshold of “normal social rhythm”. That is, I labeled SRM scores < 3.5 as unstable states and scores ≥ 3.5 as stable states. This is a binary classification problem with two different classes (stable and unstable). For modeling, I used SVM (Support Vector Machine). For training, I used the same set of features as in the regression task. Based on 10-fold cross validation, this model achieves high accuracy with a precision score of 0.85 and recall of 0.86.

Feature	Ranking	Weight
Distance traveled	1 st	1.56×10^{-2}
Cluster frequency	2 nd	3.27×10^{-3}
Non-sedentary duration	3 rd	-3.79×10^{-4}
Conversation frequency	4 th	7.69×10^{-5}

Table 5.3: Ranking of feature importance for the stable and unstable status classification using recursive feature elimination (RFE). It also shows weights assigned to features in a support vector machine using a linear kernel.

I also did feature ranking to better understand the impact of each feature

on the model. Specifically, I used Recursive Feature Elimination (RFE) [90]. As I have mentioned before (Chapter 4), RFE calculates feature ranking in subsequent steps. At each step, a model is trained on the complete dataset and the feature least contributing to the model is discarded. This continues till there is only one feature left which is considered as the most important one. The ranked features and associate weights quantifying the impact on the model is shown in Table 5.3. In this dataset, the location based features (i.e., cluster and total distance traveled) are the most important ones for predicting stability in bipolar disorder. This is consistent with findings from a number of previous studies that have also identified location based features to be useful indicators of mental health states [35].

Further, I also calculated class probability estimates which provides more granular information about model performance compared to just prediction outputs. Probability estimates associated with model prediction are easily interpretable and can be particularly helpful in conveying uncertainty of a statistical model. As such, these probability scores can help clinicians in making informed decisions, particularly for borderline cases. For example, clinicians can choose to collect more data for prediction outputs with high uncertainty to avoid potential cost associated with misclassification.

For calculating the probability distribution associated with model outputs, I used Platt scaling [180]. This method fits a logistic regression model using the classifier outputs for calculating associated probability. After performing Platt scaling, I found that 75.89% of correctly classified labels have a probability ≥ 0.7 . In other words, the model has high confidence for the majority of the predictions. This finding shows that the model is quite robust for this dataset.

5.4 Discussion

For patients with bipolar disorder, monitoring and maintaining circadian stability is crucial for reducing the risk of relapse. In this chapter, I have described MoodRhythm — a phone application that can passively monitor circadian stability and social rhythms of patients. Based on data from 7 patients with bipolar disorder over 4 weeks, I found that the models using passively sensed features can accurately distinguish between stable and unstable states. The confidence scores associated with the prediction outputs further show that the model is quite robust for this dataset. The ability to monitor and predict circadian disruptions in mental health patients can significantly broaden the scope of circadian computing. Towards this, I will describe the implications of these findings from a broader perspective of mental health care and circadian computing. I will also point out some of the limitations of this study and how future work can address those issues.

5.4.1 Implications and Opportunities

Given that there is no cure for bipolar disorder, it requires lifelong management and constant vigilance against relapse onset. As mentioned above, interventions targeting circadian and social rhythm stability can be an effective tool in this case. However, the existing clinical tools are inadequate for tracking and maintaining circadian stability and social rhythms over a long period of time. For example, while using the paper-and-pencil based SRM, patients can forget to complete entries. Also, there is no way to recover lost records. Its reliance on individual's ability to recall events can pose serious challenges in maintaining

data quality and coverage, particularly when a patient is in manic or depressive phase. As a result, the existing tools often do not capture the subtle cues of behavioral and contextual patterns that might indicate relapse onset. In this section, I will describe how the findings from this study can help to overcome these limitations. I will also point out that being able to passively sense circadian stability can help identify early-warning signs, which can lead to preemptive care instead of reactive care.

Granular and Longitudinal Tracking

The methods and models describe in this chapter is a significant step towards more granular tracking while keeping the user burden low. These methods do not require any user involvement. They also do not rely on individual's ability to recall events. Moreover, the increasing sensing capabilities of phones means that the monitoring of behavioral and contextual patterns can be more comprehensive compared to current clinical tools. The first version of SRM had 17 items [75] but was reduced to 5-items given that it is very difficult to manually track so many items on a daily basis. A phone based method, on the other hand, can track a wider array of behaviors and contexts without adding any user burden.

As a result, the methods and models described in this chapter are not limited to just five SRM items. Indeed, these methods can be used to track a large number of behavioral and contextual cues to identify person specific markers of circadian and social rhythms instability. For example, for some patients the change in online behavior can indicate relapse onset, while for others offline behaviors like conversation frequency and distance traveled might provide more insights. A passive and unobtrusive way of sensing makes monitoring of such

diverse arrays of behavioral patterns feasible for a large population of patients.

Moreover, the sensing methods described here also address the limitations of existing tools when it comes to longitudinal tracking. Given that these methods do not rely on user inputs, they can be particularly useful when a patient is very symptomatic. It is quite unlikely that a patient will remember to complete SRM items when suffering from depression or being in a manic phase. However, the behavioral cues during these periods can provide invaluable insights into the idiosyncratic patterns of illness. A passive and automated sensing method, on the other hand, can keep collecting data even when a patient has relapsed. Such data can be very useful for both immediate interventions and long term clinical planning.

Preemptive Care and Intervention

Being able to collect data over long period of time will also help to identify individualized warning signs. That is, if we have data about behavioral trends before and during a relapse, it is possible to identify how relapse onset manifests in changing these trends for a patient. In other words, this will help in identifying personalized early-warning signs that can be used to prevent relapse onset in the future. For example, the gradual decrease in conversation frequency over time might indicate that the onset of depressive state for a patient. Such early detection of relapse signatures would enable preemptive care — allowing clinical help even before the relapse happens. Furthermore, understanding individualized symptom cues will enable patients and clinicians to come up with a more effective personalized treatment plan.

The existing clinical practices and tools are also severely limited when it comes to timely intervention. Indeed, the delay between patients completing SRM entries and clinicians having access to them can be more than ideal. The sensing method and inference models presented in this chapter can make it significantly easier for timely feedback from the clinicians by identifying and sharing markers of disruptions in real-time. Moreover, beyond just being a tracking tool, phones themselves can be an effective intervention delivery method — providing care whenever patients need it and wherever they need it.

Beyond Bipolar Disorder

In addition to bipolar disorder, SRM has been used in a number of other clinical conditions. Campos et al. [34] used SRM to quantify regularity of daily activities in stroke patients. Câmara et al. [33] used it for tracking lifestyle regularities of patients with Parkinson’s disease. Similarly, Schimitt et al. [204] used SRM for identifying the correlation between chronobiological variables and characteristics of juvenile myoclonic epilepsy. SRM has been used for measuring the rhythmicity of daily behaviors among patients with anxiety disorder [208] and unipolar depression [47] as well.

As such, the methods and findings presented in this chapter potentially apply to these clinical conditions as well. Indeed, given that these studies used SRM mostly for tracking stability in daily activities, the predictive model for automatically inferring SRM scores would be useful for these conditions as well. Overall, I believe that having a passive and unobtrusive method for assessing stability would greatly enhance the practicality of SRM and IPSRT as a clinical tool and a research instrument.

5.4.2 Limitations

The findings in this study are based on a small set of participants. The study population was also not balanced in terms of gender and age group. It would be useful to replicate these findings over a larger and more diverse patient population. In particular, future work should investigate whether younger patients have significantly different behavioral markers during stable and unstable states compared to this study population. Also, the study duration was relatively short with 4 weeks of data collection. It will be useful to investigate how these models perform over a long period of time.

For logistical reasons, the patients did not use their own phones during the study. This might have resulted in departure from normal phone use behaviors of a patient. Also, this requirement of carrying another phone can have significant impact in adherence to study protocol for a longitudinal study. For example, patients might forget to carry or charge the study phone. This would result in data gaps regarding user behaviors and contexts. For future studies, I would thus recommend using participants' own phones whenever possible for monitoring behavioral markers.

During the recruitment, our research team only considered euthymic patients. This is due to the fact that adherence to study protocol (and SRM tracking) can be very low during manic and depressive phases. Future work should investigate how behavioral markers change for a relapsing patient. In particular, it will be useful to compare the performance of the models across euthymic, manic, and depressed states.

I only focused on assessing and predicting SRM scores in this study. There

is an opportunity to extend these findings to include social activities as well. Specifically, a patient tends to be less socially active during depressed states and much more socially involved during a manic episode. Future studies should investigate whether the model performance improves significantly by taking social activity features into consideration.

The findings from this study regarding passive assessment of social rhythms and circadian stability addressed a number of limitations of existing clinical tools. However, it is worth considering whether a completely automated system might impact the therapeutic process. Indeed, there is a potential risk that positive therapeutic elements associated with self-tracking — for example, having a sense of involvement in treatment and control over one’s illness — might be lost in a completely passive sensing system. This might be particularly true for automated systems replacing SRM given that it is both a measurement of stability and a tool for helping patients structure their days.

5.5 Summary

In this chapter, I described a passive and unobtrusive method to infer SRM scores — a clinically validated marker of stability in bipolar disorder. Based on data from 7 patients over 4 weeks, I showed that the model using passively sensed data as features can accurately distinguish between stable and unstable states with a precision of 0.85 and recall of 0.86. Further, the confidence score associated with prediction outputs also showed that the model was quite robust for this dataset.

Given that maintaining circadian and social rhythm stability is a key as-

pect of disease management for patients with bipolar disorder, the findings presented here can have considerable impact on clinical care. In particular, being able to passively and unobtrusively monitor stability in patients will potentially result in more accurate and granular data over longer periods of time while significantly lowering user burden. The model presented here can be further extended to detect personalized early warning signs based on collected behavioral and contextual cues. Furthermore, circadian-aware tools based on these findings could open up novel ways to provide interventions.

Overall, these methods and findings extend the scope of Circadian Computing specifically in the context of mental health. Circadian-aware tools as presented here can significantly help not only in illness monitoring but also enabling preemptive care that focuses on circadian stability of patients with mental illness.

CHAPTER 6

DISCUSSION

In this concluding chapter of the dissertation, I aim to be introspective as well as forward-looking. Throughout this dissertation, I have don the hat of a technologist, with a focus on design, development, and deployment. However, I strongly believe that as a technologist it is also my responsibility to explore the other side — the potential unintended consequences that might result from the broad vision presented here. For this, I will start summarizing the previous chapters and then point out how these findings provide the foundation of Circadian Computing. Then, I will take a critical look at the key assumptions behind Circadian Computing to identify limitations and potential sources of tension. I will also offer concrete strategies to balance the trade-offs to maximize the benefits for our health and well-being. In the second half of this chapter, I aim to look forward to identify potential opportunities for Circadian Computing. In particular, I will focus on better modeling of circadian disruptions using novel data sources and sensors, intervention strategies for circadian stability, and potential application of circadian-aware tools across a range of domains.

6.1 Summary of Prior Chapters

In Chapter 1, I set the stage for this dissertation by introducing the concept of Circadian Computing. I described the critical role circadian rhythms plays in our overall health and well-being. However, the existing tools for monitoring circadian stability are severely limited; specifically, when it comes to tracking in the wild. I pointed out that ubiquitous technologies can help address these issues. Indeed, over the years, ubiquitous computing researchers have worked

on developing pervasive technologies for better health, greater productivity and improved well-being. However, they have mostly ignored the role of our innate biological rhythms — one of the most critical components when it comes to our health and cognitive performance. I explained how Circadian Computing aims to bridge this critical gap by developing circadian-aware technologies across a wide range of application domains.

In Chapter 2, I contextualized this dissertation in relevant prior work. I provided the theoretical background explaining how circadian systems maintain periodicity by using environmental cues like light exposure. I also introduced a number of terminologies and vocabularies used in the remainder of this discussion. Furthermore, I detailed the current findings connecting circadian system to our health and well-being, particularly in the context of sleep, metabolism, performance, and mental health issues. Then, I described existing methods to measure circadian disruptions including biological markers (e.g., core body temperature), physiological markers (e.g., actigraphy data), and self-report measurements (e.g., MCTQ). I also identified several limitations of these tools when the goal is in-situ and longitudinal monitoring. I argued that pervasive technologies based on phones and wearables can address some of these limitations and extend the prior findings from the controlled environment of a lab to real-world setups.

In Chapter 3, I described such a system for passive monitoring of circadian disruptions that leverages phone usage patterns. My collaborators and I deployed the system for 97 days over 9 participants. Based on this data, I showed that the proposed method can accurately infer sleep onset and duration, chronotype, as well as markers of circadian disruptions including social jet lag, and

sleep inertia. I also explained how these findings enable granular and longitudinal monitoring of circadian stability and as such, lead to a critical aspect of Circadian Computing.

In Chapter 4, I focused on the relationship between circadian rhythms and cognitive performance. Our cognitive abilities wax and wane over the time of a day. For example, alertness, which is crucial for our physical and cognitive performance, strongly reflects circadian rhythms. In this chapter, I specifically aimed to model the ebb and flow of alertness in the wild. For this, I first helped to develop a data collection framework and deploy it among 40 participants over 20 days. Based on the collected data, I showed that a number of circadian (e.g., chronotype and time of a day) and external environmental (e.g., Daylight Savings Time transition) factors can significantly impact alertness levels. Leveraging these insights, I then developed models that can passively and unobtrusively predict one's alertness level using circadian and environmental data.

Given the role of alertness in almost every aspect of cognitive performance, these findings have potentially far reaching positive impacts. In particular, I noted how these findings could result in circadian-aware tools focusing on preventing occupational errors and transportation accidents. Furthermore, given the relationship between alertness, learning and memory, these findings can fundamentally reshape educational tools.

In Chapter 5, I focused on broadening the scope of Circadian Computing to mental health care. Circadian and sleep disruptions have been associated with a number of mental illnesses. Specifically, circadian stability is crucial for illness management of patients with bipolar disorder. However, existing clinical tools for monitoring such disruptions are often unable to collect granular and

longitudinal data. In addition, their efficacy at delivering interventions for stability is typically low. This can significantly hinder clinical decision making and ultimately result in sub-optimal illness management. In this chapter, I aimed to address these issues. Specifically, I focused on the design, development and evaluation of MoodRhythm — a phone application that uses sensor data to passively infer circadian stability in bipolar disorder. Based on data from 7 patients over 4 weeks, I showed that the proposed data-driven methods can accurately infer key markers of stability in patients with bipolar disorder. These findings reaffirm the potential roles Circadian Computing can play in healthcare. Specifically, circadian-aware tools can reshape the mental healthcare from reactive to preemptive by identifying early-warning signs. These tools can also help novel forms of intervention and relapse prevention.

In the following sections, I will connect the dots and reflect on the body of work presented so far. I will also discuss opportunities for future work.

6.2 Considerations

Throughout this dissertation, I have focused on positive aspects of circadian-aware tools in our daily lives, health, and well-being. However, any broad vision of technology as presented here will have unintended consequences. In this section, I will aim to identify these blind-spots. Specifically, I will go over the key assumptions and promises of Circadian Computing to identify potential limitations and need for trade-offs. I hope this section will encourage a balanced and conscientious approach to the future development in the realm of Circadian Computing.

6.2.1 The Contradictory Roles of Technology

In this dissertation, I have developed new technologies and methodologies with a focus on understanding and ensuring circadian stability. In particular, I have extensively used sensor-rich phones for collecting data. I have also proposed to use these devices as a delivery medium for in-situ interventions. However, it is worth noting that use of these very devices can precipitate sleep and circadian disruptions. For example, Chang et al. [41] reported that evening use of these light emitting devices delays circadian clock and reduces duration of REM sleep.

These issues are further compounded by the fact that modern technologies often aim to keep the participants “hooked” over a long period of time [235]. In other words, users are heavily incentivized to keep using these devices and services. However, such frequent engagements with these devices can result in negative consequences. A number of studies have found that late night phone and media device use result in poor sleep quality and excessive daytime sleepiness [38]. Device use at night can negatively impact alertness [41] and job performance [127] next day as well.

The negative aspects of device use can be particularly problematic for young adults given their wide adoption of technology. In 2015, Hysing et al. [105] conducted a large cross-sectional population based study consisting of 9,846 adolescents and found that frequent device use is significantly associated with sleep onset latency and sleep deficiency. This is consistent with findings from other studies as well [70]. Beyond sleep, over-reliance on these devices has also been associated with mental health issues including stress and depression [218].

In light of this, it is important to ensure that Circadian Computing does not exacerbate these issues. The sensing algorithms I have presented in Chapter 3 and 4 rely on such consistent use of devices. That is, these algorithms work because we are hooked to our phones — often phone checking is the first activity we perform after getting up and the last activity before going to sleep. Future sensing methods would do well to make sure that their accuracy is not tied to potentially problematic use of devices and technologies. One way of achieving that would be to make use of multi-modal sensor data by combining behavioral and environmental features together. For example, instead of just using phone patterns, Circadian Computing can leverage energy usage data to identify long-term sleep patterns and circadian disruption trends over a large geographical area [23].

The other related challenge here is to design interventions that are effective but do not add to issues mentioned previously. In particular, interventions focusing on circadian stability and optimal scheduling will need to be personalized and data driven. In most cases, interventions also need to be in-situ. From a technical perspective, these requirements are well-addressed by phones and wearables. However, leveraging these devices for delivering interventions for circadian stability will require careful design to ensure that they do not encourage sustained device use.

The current notification system used by phones and wearables might not be appropriate for this purpose. Phone notifications can be too distracting and disrupting [214]. Kushlev et al. [123] found that these notifications can lead to hyperactivity symptoms which are associated with Attention Deficit Hyperactivity Disorder (ADHD). As such, for intervention, Circadian Computing

should focus on providing users with actionable suggestions from the *periphery*. For example, instead of using distracting notifications, such a system can passively provide useful information about one's alertness rhythm by changing background color of the phone screen. Such careful design considerations for user engagement along with a bottom-up approach for circadian-aware devices (described below) can result in a more positive and balanced relation with the our personal devices.

Successful adoption of novel technologies often requires careful design considerations that balance different trade-offs. Circadian Computing will be no different. Designers and developers of circadian-aware technologies need to be particularly conscientious about unintended consequences given their potential impact on health and well-being. Towards that, I hope that the above section will serve as a guideline and a conversation starter. Being mindful about the limitations and potential "blind-spots" can only help to make Circadian Computing a reality.

6.3 Future Opportunities

Now, I will look forward and identify future opportunities for Circadian Computing. That is, *where can we go from here? Which application domains will be specifically benefited from incorporating circadian perspectives? How can we improve the accuracy of methods for identifying individualized circadian trends?* In this section, I will answer these questions and detail future directions that I find particularly compelling.

6.3.1 Novel Sensing for Granular Assessment

In this dissertation, I have described methods for assessing sleep and circadian disruptions. However, this is just the beginning. I believe there are some exciting opportunities to develop novel sensing methods for more granular tracking of circadian stability at a population level. Towards this, future work should take advantage of wearable devices. In recent years, we have seen an exciting number of novel wearables aiming to track different aspects of behavior and health metrics. Data from these wearables can certainly help to better monitor one's circadian trends.

In particular, being able to quantify light exposure will be very useful in modeling and predicting one's circadian rhythms. There has been some interesting development in recent years to identify and monitor light exposure. For example, Koo et al. [118] used meteorological satellite images to infer outdoor artificial light exposure at night over a large geographical area. However, this method is limited to aggregated population and thus can not be used to infer individual exposure to light.

A number of wearable light sensors have been developed to track individual light exposure as well. For example, Actiwatch Spectrum [176] is a commercial product that can be worn on the wrist. However, given that circadian system is sensitive to light *through our eyes*, the data from these wrist-worn devices might not be very accurate. For a more accurate assessment, Bierman et al. [24] have developed *Daysimeter* — a small wearable device with better photometric performance [68]. Using light exposure data in combination with sleep information can help not only to identify circadian disruptions but also point out the factors responsible for such disruptions.

Similarly, food intake can also impact our circadian system [234]. Data about meal timing, thus, can help in modeling one’s circadian phase. A number of recent work has focused on tracking food intake, leveraging self-reporting, video monitoring, force sensors, and wearable devices [230]. Future sensing work within Circadian Computing should take advantage of these methodologies.

Beyond these existing wearables, there are also opportunities to develop novel *bio-sensors*. These bio-sensors could directly track physiological markers to infer circadian phase. For example, cortisol — a steroid hormone — distinctively reflects circadian rhythm [207]. Bio-sensors measuring cortisol over time, thus, can provide granular information about one’s circadian phase. Towards this goal, my collaborators and I have recently started working on a salivary cortisol detection method. Bio-sensors that can accurately detect melatonin from saliva could also be used for accurate assessment of circadian phase [188]. Furthermore, recent studies have used gene expressions to predict circadian phase [103, 5]. Future work could potentially extend this work for a minimally-invasive but accurate method for monitoring circadian patterns.

Overall, I believe that there are some exciting opportunities to develop novel sensing methods that would significantly improve our ability to track one’s circadian phase. These methods would help to extend and complement current goals of Circadian Computing.

6.3.2 Fixing a Broken Clock: Sleep and Circadian Interventions

Effective interventions focusing on sleep and circadian stability could have a significant positive impact on health and well-being at a population scale. In

recent years, both academic and industrial research have focused on improving various aspects of individual's sleep. However, it is imperative to consider circadian factors and the effect of *zeitgebers* in this context. In particular, interventions that only focus on sleep disturbances might be merely treating the symptoms of a misaligned circadian system but not the underlying causes. Indeed, Schroeder et al. [206] argued that sleep problems are often caused by poor "circadian hygiene". Circadian Computing, thus, can address these issues by providing a holistic approach, which considers circadian factors, chronotype information, and sleep patterns. In the following section, I will point out how circadian-aware technologies could enable effective environmental and lifestyle interventions to stabilize a disrupted circadian system.

Light Exposure

Light is one of the most important *zeitgebers*. Exposure to light can significantly shift circadian timing [114]. In particular, even low-intensity light at evening and early night can significantly delay circadian phase [248]. Controlling light exposure is thus crucial to maintain circadian hygiene. However, the ever-increasing use of artificial light at night can significantly disrupt our circadian system [49]. The rise of light emitting electronic devices also poses serious challenges for a user to maintain circadian hygiene. For example, Chang et al. [41] reported that use of light emitting electronic devices before bed-time results in delayed circadian phase, increased sleep latency, and impaired morning alertness.

Circadian Computing can help to address these issues. In recent years, we have seen softwares that automatically change light intensity of electronic de-

vices at night (e.g., f.lux [69], Night Shift [9]). There is an opportunity to push this idea further to a circadian friendly smart-home environment. That is, Circadian Computing can integrate the so called light therapy in most aspects of our daily lives. Specifically, I think a dynamic light system in home environment that minimizes negative impacts of light exposure can be particularly useful. For example, instead of having a static light system, a dynamic light system will change its intensity and color to better accommodate one's circadian profile.

Such a system will not only be limited to light composition and intensity at nights. Schroeder et al. [206] pointed out the need for greater variance between day and night light exposure for robust circadian stability. A number of recent studies have also pointed out that using blue light in the morning can help in sleep and circadian stability [158, 83]. A circadian-aware environment can facilitate such appropriate exposure throughout the day (e.g., light intensity in the environment following the Sun). Indeed, in a recent work, Rea et al. [185] argued for adopting circadian considerations when planning lighting conditions in architectural spaces.

These systems do not need to be limited to one's home. Appropriate light exposure can improve productivity, mood, and sleep quality of office workers [220, 67]. Dynamic lighting system has also been shown to improve academic performance of the students [18]. Circadian-aware systems could be particularly useful for hospitality industry. *Imagine your hotel providing a personalized light exposure based on your circadian profile to minimize your jet lag.*

Chrono-nutrition

For a significant fraction of the society, food is constantly available throughout the day. This could result in departure from traditional eating schedules. However, inappropriate meal timing can have serious negative impacts for our health and well-being. Based on genetic studies, Huang et al. [101] argued that circadian system is tightly coupled with metabolic processes. Similarly, Wehrens et al. [234] found that meal timing can influence circadian phase. Specifically, they reported that a 5-hr delay in meal times resulted in delay in circadian phase.

In light of these findings, I think it is feasible to develop personalized meal scheduling that can minimize circadian disruptions. Indeed, Asher et al. [13] argued for *chrono-nutrition* — scheduling food intake in accordance with our biological rhythms. Circadian Computing can provide frameworks for monitoring one's circadian phase and then provide personalized recommendations for meal timing. Such tools could help individuals to maintain circadian stability over a long period of time. These tools would also be particularly useful for travelers crossing time zones and shift-workers. Furthermore, such circadian-aware meal scheduling could help to minimize risks of metabolic diseases, particularly among shift-workers [53].

A number of recent studies have also identified exercise as a potential factor in shifting circadian phase. Barger et al. [17] reported that exercise can cause circadian phase delay in humans. Similar findings have been reported by other studies as well [14, 59]. Moreover, Yamanaka et al. [244] reported that physical exercise can help to quickly stabilize a disrupted circadian system. As such, appropriate exercise duration, intensity, and timing could be used as an effective

intervention strategy for sleep and circadian stability. To facilitate such stabilization, a circadian-aware tool could provide personalized recommendation of type and duration of exercise depending on one's chronotype and level of disruption.

Our modern lifestyle often is incongruent with our innate biological rhythms. This can result in circadian and sleep disruptions with significant negative impacts on our health and well-being. Environmental and lifestyle interventions can be effective for maintaining a robust and stable circadian system. Circadian Computing can provide an integrated data-driven framework to facilitate these interventions. In this framework, the ability to continuously sense one's circadian phase would drive personalized, in-situ, and actionable recommendations for reducing circadian disruptions effectively.

6.3.3 Optimizing Cognitive and Physical Performance

Our physical and cognitive processes often reflect circadian rhythms. As such, being able to monitor these trends and identifying ways to intervene can potentially help to optimize the outcome of these processes. That is, Circadian Computing can help to optimize our cognitive and physical performance *in accordance with* our individualized biological rhythms.

Learning and Education

Learning, memorization, and problem solving reflect circadian rhythms [25, 31]. Recent studies have also found a relationship between circadian phase and aca-

ademic performance [113, 55]. Furthermore, circadian disruptions can adversely affect our memory and learning capabilities [240], potentially resulting in negative academic performance. For example, based on grading data from 523 students, Zerbini et al. [249] found that late chronotypes perform significantly worse than early chronotypes, which the authors attributed to circadian disruptions resulting from early school schedules.

As such, a consideration of biological rhythms of the students could significantly improve the efficacy of educational institutes. This would be particularly useful for young adults (high school and college students) given that they tend to be mostly late types. Indeed, the current school scheduling is often too early for these students, resulting in sub-optimal memorization and learning performances.

Circadian Computing can help to make a change in this context in two ways. First, it can help to monitor and predict the attentional and learning rhythms of the students. Such information can in turn help educators in making decisions regarding the timing of particular class, learning activities, and exams. For example, Zerbini et al. [249] reported that difference in academic performance across chronotypes are much higher for subjects related to science compared to humanistic or linguistic ones. Based on this finding, an educational institute might decide to schedule science education later in the day.

Secondly, Circadian Computing can help students to make informed decisions regarding their own learning and attentional rhythms. For example, being aware of one's individualized rhythm can help a student to choose an optimized class scheduling. Furthermore, circadian-aware tools can also help in active learning by customizing learning tasks depending on one's chronotype

and task difficulty. For example, these tools can help to better schedule tasks with different memory recall requirements [66].

More importantly, these tools and findings are applicable beyond the traditional boundaries of academic institutes. Circadian Computing can reshape online learning (e.g., Massive Open Online Courses, or MOOCs) as well. Imagine a MOOC course not only focusing on the content but also having a personalized scheduling of learning *just for you*.

Scheduling and Activity Management

Given the predictable rise and fall of our cognitive performance, I believe the domains of task scheduling and activity management are open for significant innovations. For example, a circadian-aware calendar could potentially reshape how we organize our daily tasks and responsibilities. That is, *how would you schedule your day if you could predict your alertness states in advance?* By taking one's chronotype and personalized trends of alertness, a circadian-aware calendar could provide recommendations about when to schedule a task with a given difficulty level. For example, it would recommend rote tasks (e.g., doing laundry) in the morning and cognitively demanding tasks in the evening for a late chronotypes. Having circadian perspective embedded in calendars — one of the most widely used digital tools — can also help to raise awareness about individualized circadian rhythms.

These tools can facilitate better collaboration as well. Given the increasingly global and distributed nature of our workforce, being able to identify collaborators with synchronized alertness profiles can be very useful. For example, a

late chronotype from the West Coast would be a good match with an early type from the East Coast for collaboration. Similarly, these tools could also enable collaborative scheduling. That is, instead of just using mutual availability as a criteria, circadian-aware tools could also consider individualized rhythms of alertness for scheduling collaborative tasks.

Accident Prevention

Circadian Computing can reduce the risk of industrial and transportation accidents. Specifically, alertness and fatigue issues have been associated with occupational and driving safety. According to the National Transportation Safety Board (NTSB) 30% of all road accident fatalities in the US involve fatigue and sleepiness [161]. Given that alertness is a circadian process, circadian disruption can significantly increase the risk of accidents [174]. Indeed, transportation accident patterns reflect a circadian cycle with major peaks around 2 AM, 6 AM, and 4 PM [100]. Industrial accidents also show similar patterns [71].

As such, technologies able to track and predict individualized variations in circadian processes like alertness can significantly reduce the risk of these accidents. In critical cases, these tools could not only monitor one's alertness but also provide in-situ recommendations and interventions to sustain it. There has been some recent work in preventing fatigue in aviation and flight operations [32]. Circadian Computing can extend these countermeasures to other domains as well.

Higher Level Cognitive Functionalities

Alertness is a key biological process that underpins a number of higher level cognitive functionings. For example, two phases of creative ability — divergent and convergent thinking — have different attentional requirements. That is, divergent thinking requires *defocused* attention while convergent thinking requires focused attention. Facilitating focused and defocused attention, thus, can improve different stages of creative ability. Towards this goal, circadian interventions can be particularly useful.

Specifically, our circadian system is sensitive to blue light (in 460 – 480 nm wavelength range) [15]. Blue light has shown to improve a number of cognitive processes including alertness and concentration [166, 157]. Indeed, Bevan et al. [21] found blue light to be more effective than caffeine for enhancing alertness. Based on these findings, I have used blue light as an intervention to improve convergent thinking and creative ability [1]. In this preliminary study, participants used a blue light source with a spectral wavelength range of 475 – 480 nm for 20 minutes. Based on data from 21 participants over 2 weeks, I found that blue light exposure significantly improved the convergent thinking ability of the participants.

I believe similar circadian interventions can also be applicable for other high-level cognitive functions. Schmidt et al. [205] reported that processes underpinning human cognition often can be influenced by circadian factors. As such, Circadian Computing can not only help to identify the opportune timing for these cognitive processes (“a time to think”) but also provide interventions to optimize our cognitive abilities.

6.3.4 Clinical Applications

Bringing circadian perspectives to clinical applications is one of the most promising future directions of Circadian Computing. Circadian-aware tools can be applied across a wide range of clinical applications including mental health care, diagnostic testing, and medication delivery.

Mental Health Care

Mental health is an urgent global issue. More than 450 million people around the globe suffer from mental illnesses [167]. However, existing health care systems are inadequate to address these issues. Indeed, the global planned annual spending on mental health is less than \$2 per person; in low-income countries, the amount is less than \$0.25 per person [169]. As a result, we need technology mediated systems that can complement current health care infrastructure. In other words, we need an end-to-end framework that can support mental health patients and clinicians. Such a framework should do early-detection, in-situ intervention, and provide long term support for the patients. The tools and methodologies developed by Circadian Computing can significantly help in developing such frameworks.

Specifically, a growing number of research studies have linked circadian disruptions with mental illnesses including bipolar disorder, schizophrenia, and depression [173, 96]. As such, being able to monitor circadian disruptions and provide appropriate interventions for circadian stability can be crucial in the context of mental health care. In particular, longitudinal monitoring of circadian stability can help to identify departure of a patient from a stable state. This

would help to develop early-warning systems that could identify personalized cues of relapse onset. In other words, circadian-aware tools that use behavioral and contextual data for longitudinal passive sensing can usher the mental health care systems from being *reactive* to *preemptive* — enabling clinical help even before the onset of relapse.

Similarly, Circadian Computing can help to identify effective and personalized interventions that would improve long-term prognosis of a patient. For example, circadian disruption can precipitate manic or depressive phase onset in bipolar disorder. Lifestyle tools focusing on sustaining stability, thus, can significantly help to prevent relapse onset. Similar tools could be crucial for dementia, depression, and PTSD since these illnesses have been associated with circadian disruption as well [241]. Indeed, future work will do well to extend tools like MoodRhythm (described in Chapter 5) that can collect behavioral and contextual data, and then provide data-driven personalized recommendations to patients for maintaining circadian stability. These tools will also allow clinicians to provide in-situ interventions and, thus, closing the gap between relapse onset and treatment delivery.

Chronotherapy

I am particularly excited about chronotherapy — an emerging field in medicine, which focuses on optimizing health outcomes by aligning treatment steps with our biological rhythms [213]. Such alignment of medication dosage and timing can result in significant clinical improvements. For example, chronotherapy based cancer medication has been shown to be 50% more effective [131]. However, current clinical practice does not take our circadian information into

consideration. This is often due to the difficulty in adapting treatment steps to inter-individual differences.

Circadian Computing could make a difference here by addressing these issues to enable chronotherapy in a scalable way. For example, bio-sensors that can track cortisol level from saliva can be used for identifying one's circadian phase at any given time. Such information could potentially be useful to determine optimized timing for medications. Similarly, chronotherapy will also change the way we take our medications. It would be personalized and will depend on the phase of our internal rhythms. As such, personalized assistive system in this regard would be particularly useful for anyone taking medications. Towards this, Circadian Computing could work on assistive systems similar to Siri [8] or Google Now [86], which will be based on one's circadian profile to improve medication outcomes.

Similarly, Circadian Computing can also improve accuracy of a wide range of diagnostic tests. Symptom intensity of a number of medical conditions reflect circadian rhythms. For example, Smolensky et al. [212] reported that symptom intensity of allergic rhinitis (AR) and bronchial asthma show a diurnal pattern. Similarly, gout [93], Biliary colic [191], and peptic ulcer attacks [153] also show circadian rhythms with worsen symptoms during the night. Our cardiovascular system follows circadian rhythm as well [130]. For example, sudden cardiac death [45], stroke [62], acute myocardial infarction [156], and congestive heart failure [245] peak during the morning. Tests using blood pressure and heart rate, thus, should take this circadian variation in consideration. Similarly, the test outcome for glucose tolerating [250], hematology, coagulation, and hormone [219] might depend on the timing of the test.

As such, being able to monitor circadian phase of a patient could not only improve the accuracy of these tests but also enhance the efficacy of medications. For example, Circadian Computing could suggest how one's circadian phase should be taken into consideration while interpreting data from a diagnostic test. Similarly, it could assist clinicians and patients to identify best time for medication intake that would optimize the outcome. Overall, I believe that in next few decades, we will see a dramatic shift in medicine that will focus more on circadian aspects of medication and diagnosis. Circadian Computing could facilitate this shift by bringing research findings to in-practice care.

6.4 Concluding Remarks

In modern society, humans no longer need to live according to the position of the sun. Technologies, availability of lighting, and social conventions are increasingly disrupting the traditional temporal boundaries. However, our innate biological clocks still tick with a 24 hour periodicity — reflecting the temporal cues consistent throughout our evolutionary history. The resulting circadian rhythms shape almost every neurobehavioral processes including sleeping, cognitive and physical performance, metabolism, and mood. Circadian stability is, thus, critical to our overall health and well-being. Consistent circadian disruption can have serious consequences with an increased risk for cancer, obesity, diabetes, and relapse onset in patients with mental illnesses. Moreover, circadian disruption can also severely impair our alertness and cognitive performance contributing to loss of productivity in workplace and occupational accidents.

As such, continuous assessment of circadian disruption is important for both

individual well-being and public health. Recently there has been an increased focus on monitoring and identifying disruptions in circadian rhythms. However, these methods and findings are often limited to controlled lab environments leaving a gap in knowledge and applicability when it comes to real-world setups. As such, there is an explicit and urgent need for the development of technologies that can assess and monitor circadian disruptions in-situ, over long periods of time, and on a global scale. There is also an opportunity for developing intervention tools for maintaining circadian stability with applications ranging from optimizing performance in accordance with our innate biological rhythms to preventing relapse onset in patients with mental illnesses.

My dissertation is a leading step towards this broad and novel vision of circadian-aware technologies. In my PhD work, I have laid the groundwork for *Circadian Computing* — technologies that support and adapt to our innate biological rhythms. Specifically, I have developed and evaluated methods for unobtrusively assessing circadian disruptions. I have also showed that behavioral and contextual data can be used for modeling and predicting alertness — a circadian process integral to our cognitive performance. I have also developed, deployed, and evaluated a data driven tool focusing on identifying circadian anomalies in patients with bipolar disorder.

With this groundwork in place, I believe that there is an exciting opportunity lying ahead for Circadian Computing. I have pointed out how circadian-aware technologies can potentially reshape a number of application domains including education and learning, optimized scheduling, mental health care, and chronotherapy. I hope this dissertation provides a road map for a new perspective in technology development and contributes to the shared effort of

improving our productivity, health, and well-being.

BIBLIOGRAPHY

- [1] Saeed Abdullah, Mary Czerwinski, Gloria Mark, and Paul Johns. Shining (blue) light on creative ability. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 793–804. ACM, 2016.
- [2] Saeed Abdullah, Mark Matthews, Ellen Frank, Gavin Doherty, Geri Gay, and Tanzeem Choudhury. Automatic detection of social rhythms in bipolar disorder. *Journal of the American Medical Informatics Association*, 23(3):538–543, 2016.
- [3] Saeed Abdullah, Mark Matthews, Elizabeth L Murnane, Geri Gay, and Tanzeem Choudhury. Towards circadian computing: early to bed and early to rise makes some of us unhealthy and sleep deprived. In *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing*, pages 673–684. ACM, 2014.
- [4] Saeed Abdullah, Elizabeth L Murnane, Mark Matthews, and Tanzeem Choudhury. Circadian computing: Sensing, modeling, and maintaining biological rhythms. In *Mobile Health*, pages 35–58. Springer, 2017.
- [5] Makoto Akashi, Haruhiko Soma, Takuro Yamamoto, Asuka Tsugitomi, Shiko Yamashita, Takuya Yamamoto, Eisuke Nishida, Akio Yasuda, James K Liao, and Koichi Node. Noninvasive method for assessing the human circadian clock using hair follicle cells. *Proceedings of the National Academy of Sciences*, 107(35):15643–15648, 2010.
- [6] Torbjörn kerstedt and Mats Gillberg. Subjective and objective sleepiness in the active individual. *International Journal of Neuroscience*, 52(1-2):29–37, 1990.
- [7] Sonia Ancoli-Israel, Roger Cole, Cathy Alessi, Mark Chambers, William Moorcroft, and Charles P Pollak. The role of actigraphy in the study of sleep and circadian rhythms. *Sleep*, 26(3):342–392, 2003.
- [8] Apple. Siri. Available at <https://www.apple.com/ios/siri/>.
- [9] Apple. Use night shift on your iphone, ipad, and ipod touch. Available at <https://support.apple.com/en-us/HT207570>.

- [10] Apple. Your heart rate. what it means, and where on apple watch you'll find it. Available at <https://support.apple.com/en-us/HT204666>.
- [11] Pedro Arezes, João Santos Baptista, Monica P Barroso, Paula Carneiro, Patrício Cordeiro, Nelson Costa, Rui B Melo, A Sergio Miguel, and Gonçalo Perestrelo. *Occupational Safety and Hygiene*. CRC Press, 2013.
- [12] Jürgen Aschoff. Circadian rhythms in man. *Science*, 148(3676):1427–1432, 1965.
- [13] Gad Asher and Paolo Sassone-Corsi. Time for food: the intimate interplay between nutrition, metabolism, and the circadian clock. *Cell*, 161(1):84–92, 2015.
- [14] Erin K Baehr, Charmane I Eastman, William Revelle, Susan H Losee Olson, Lisa F Wolfe, and Phyllis C Zee. Circadian phase-shifting effects of nocturnal exercise in older compared with young adults. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 284(6):R1542–R1550, 2003.
- [15] Helena J Bailes and Robert J Lucas. Human melanopsin forms a pigment maximally sensitive to blue light ($\lambda_{\max} \approx 479\text{nm}$) supporting activation of gq/11 and gi/o signalling cascades. *Proceedings of the Royal Society of London B: Biological Sciences*, 280(1759), 2013.
- [16] Joseph V Baranski, Ross A Pigeau, and Robert G Angus. On the ability to self-monitor cognitive performance during sleep deprivation: a calibration study. *Journal of Sleep Research*, 3(1):36–44, 1994.
- [17] Laura K Barger, Kenneth P Wright, Rod J Hughes, and Charles A Czeisler. Daily exercise facilitates phase delays of circadian melatonin rhythm in very dim light. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 286(6):R1077–R1084, 2004.
- [18] Claus Barkmann, Nino Wessolowski, and Michael Schulte-Markwort. Applicability and efficacy of variable light in schools. *Physiology & behavior*, 105(3):621–627, 2012.
- [19] Christopher M Barnes and David T Wagner. Changing to daylight saving time cuts into sleep and increases workplace injuries. *Journal of applied psychology*, 94(5):1305, 2009.

- [20] Mathias Basner, Daniel Mollicone, and David F Dinges. Validity and sensitivity of a brief psychomotor vigilance test (pvt-b) to total and partial sleep deprivation. *Acta astronautica*, 69(11):949–959, 2011.
- [21] C Martyn Beaven and Johan Ekstr m. A Comparison of Blue Light and Caffeine Effects on Cognitive Function and Alertness in Humans. *PLoS ONE*, 8(10):e76707–7, October 2013.
- [22] Francesco Benedetti, Sara Dallaspezia, Cristina Colombo, Adele Pirovano, Elena Marino, and Enrico Smeraldi. A length polymorphism in the circadian clock gene *per3* influences age at onset of bipolar disorder. *Neuroscience letters*, 445(2):184–187, 2008.
- [23] Jan Beyea. The smart electricity grid and scientific research. *Science*, 328(5981):979–980, 2010.
- [24] Andrew Bierman, Terence R Klein, and Mark S Rea. The daysimeter: a device for measuring optical radiation as a stimulus for the human circadian system. *Measurement Science and Technology*, 16(11):2292, 2005.
- [25] Katharina Blatter and Christian Cajochen. Circadian rhythms in cognitive performance: methodological constraints, protocols, theoretical underpinnings. *Physiology & behavior*, 90(2):196–208, 2007.
- [26] Diane B Boivin. Circadian rhythms and clock genes in psychotic disorders. *The Israel journal of psychiatry and related sciences*, 47(1):27, 2010.
- [27] Anne Bonnefond, Patricia Tassi, ROGE Joceline, and Alain Muzet. A critical review of techniques aiming at enhancing and sustaining worker’s alertness during the night shift. *Industrial health*, 42(1):1–14, 2004.
- [28] Alexander A Borb ly. A two process model of sleep regulation. *Human neurobiology*, 1982.
- [29] L on Bottou. Stochastic gradient descent tricks. *Neural Networks: Tricks of the Trade*, pages 421–436, 2012.
- [30] Malcolm W Brown and John P Aggleton. Recognition memory: what are the roles of the perirhinal cortex and hippocampus? *Nature reviews. Neuroscience*, 2(1):51, 2001.
- [31] Christian Cajochen, Vera Knoblauch, Anna Wirz-Justice, Kurt Kr uchi,

- Peter Graw, and Dieter Wallach. Circadian modulation of sequence learning under high and low sleep pressure conditions. *Behavioural brain research*, 151(1):167–176, 2004.
- [32] John A Caldwell, Melissa M Mallis, J Lynn Caldwell, Michel A Paul, James C Miller, and David F Neri. Fatigue countermeasures in aviation. *Aviation, Space, and Environmental Medicine*, 80(1):29–59, 2009.
- [33] Samir Câmara Magalhães, Carlos Vitorino Souza, Tiago Rocha Dias, Pedro Felipe Carvalhedeo de Bruin, and Veralice Meireles Sales de Bruin. Lifestyle regularity measured by the social rhythm metric in parkinson’s disease. *Chronobiology international*, 22(5):917–924, 2005.
- [34] Tania Fernandes Campos, Aline Braga Galvão Silveira, and Marina Tostes Miranda Barroso. Regularity of daily activities in stroke. *Chronobiology international*, 25(4):611–624, 2008.
- [35] Luca Canzian and Mirco Musolesi. Trajectories of depression: unobtrusive monitoring of depressive states by means of smartphone mobility traces analysis. In *Proceedings of the 2015 ACM international joint conference on pervasive and ubiquitous computing*, pages 1293–1304. ACM, 2015.
- [36] Julie Carrier and Timothy H Monk. Circadian rhythms of performance: new trends. *Chronobiology international*, 17(6):719–732, 2000.
- [37] Mary A Carskadon, Christine Acebo, and Oskar G Jenni. Regulation of adolescent sleep: implications for behavior. *Annals of the New York Academy of Sciences*, 1021(1):276–291, 2004.
- [38] Ben Carter, Philippa Rees, Lauren Hale, Darsharna Bhattacharjee, and Mandar S Paradkar. Association between portable screen-based media device access or use and sleep outcomes: a systematic review and meta-analysis. *JAMA pediatrics*, 170(12):1202–1208, 2016.
- [39] Pew Research Center. U.s. smartphone use in 2015, 2015. Available at <http://www.pewinternet.org/2015/04/01/us-smartphone-use-in-2015/>.
- [40] Trudie Chalder, G Berelowitz, Teresa Pawlikowska, Louise Watts, S Wesley, D Wright, and EP Wallace. Development of a fatigue scale. *Journal of psychosomatic research*, 37(2):147–153, 1993.

- [41] Anne-Marie Chang, Daniel Aeschbach, Jeanne F Duffy, and Charles A Czeisler. Evening use of light-emitting ereaders negatively affects sleep, circadian timing, and next-morning alertness. *Proceedings of the National Academy of Sciences*, 112(4):1232–1237, 2015.
- [42] Zhenyu Chen, Mu Lin, Fanglin Chen, Nicholas D Lane, Giuseppe Cardone, Rui Wang, Tianxing Li, Yiqiang Chen, Tanzeem Choudhury, and Andrew T Campbell. Unobtrusive sleep monitoring using smartphones. In *Pervasive Computing Technologies for Healthcare (PervasiveHealth), 2013 7th International Conference on*, pages 145–152. IEEE, 2013.
- [43] Kwangwook Cho. Chronic ‘jet lag’ produces temporal lobe atrophy and spatial cognitive deficits. *Nature neuroscience*, 4(6):567, 2001.
- [44] Bruno Claustrat, Jocelyne Brun, and Guy Chazot. The basic physiology and pathophysiology of melatonin. *Sleep medicine reviews*, 9(1):11–24, 2005.
- [45] Mylan C Cohen, Kati M Rohtla, Cynthia E Lavery, James E Muller, and Murray A Mittleman. Meta-analysis of the morning excess of acute myocardial infarction and sudden cardiac death. *The American journal of cardiology*, 79(11):1512–1516, 1997.
- [46] Stefan Cohrs. Sleep disturbances in patients with schizophrenia. *CNS drugs*, 22(11):939–962, 2008.
- [47] Emmanuelle Corruble, Ellen Frank, Florence Gressier, Philippe Courtet, Frank Bayle, Pierre-Michel Llorca, Guillaume Vaiva, and Philip Gorwood. Morningness–eveningness and treatment response in major depressive disorder. *Chronobiology international*, 31(2):283–289, 2014.
- [48] Andrea Cuttone, Per Bækgaard, Vedran Sekara, Hkan Jonsson, Jakob Eg Larsen, and Sune Lehmann. Sensible sleep: A bayesian model for learning sleep patterns from smartphone events. *PloS one*, 12(1):e0169901, 2017.
- [49] Charles A Czeisler. Perspective: casting light on sleep deficiency. *Nature*, 497(7450):S13–S13, 2013.
- [50] Charles A Czeisler, Jeanne F Duffy, Theresa L Shanahan, Emery N Brown, Jude F Mitchell, David W Rimmer, Joseph M Ronda, Edward J Silva, James S Allan, Jonathan S Emens, et al. Stability, precision, and near-24-hour period of the human circadian pacemaker. *Science*, 284(5423):2177–2181, 1999.

- [51] Charles A Czeisler, Richard E Kronauer, James S Allan, Jeanne F Duffy, Megan E Jewett, Emery N Brown, and Joseph M Ronda. Bright light induction of strong (type 0) resetting of the human circadian pacemaker. *Science*, pages 1328–1333, 1989.
- [52] Jean de Mairan. Observation botanique. *Histoire de l'Academie Royale des Sciences*, pages 35–36, 1729.
- [53] Christopher M Depner, Ellen R Stothard, and Kenneth P Wright Jr. Metabolic consequences of sleep and circadian disorders. *Current diabetes reports*, 14(7):507, 2014.
- [54] Charna Dibner, Ueli Schibler, and Urs Albrecht. The mammalian circadian timing system: organization and coordination of central and peripheral clocks. *Annual review of physiology*, 72:517–549, 2010.
- [55] Angela K Dills and Rey Hernandez-Julian. Course scheduling and academic performance. *Economics of Education Review*, 27(6):646–654, 2008.
- [56] Steven C Dilsaver. An estimate of the minimum economic burden of bipolar i and ii disorders in the united states: 2009. *Journal of affective disorders*, 129(1):79–83, 2011.
- [57] David F Dinges. An overview of sleepiness and accidents. *Journal of sleep research*, 4(s2):4–14, 1995.
- [58] Jill Dorrian, Nicole Lamond, and Drew Dawson. The ability to self-monitor performance when fatigued. *Journal of Sleep Research*, 9(2):137–144, 2000.
- [59] Charmane I Eastman, Erin K Hoese, Shawn D Youngstedt, and Liwen Liu. Phase-shifting human circadian rhythms with exercise during the night shift. *Physiology & Behavior*, 58(6):1287–1291, 1995.
- [60] Cindy L Ehlers, Ellen Frank, and David J Kupfer. Social zeitgebers and biological rhythms: a unified approach to understanding the etiology of depression. *Archives of general psychiatry*, 45(10):948–952, 1988.
- [61] A Roger Ekirch. Sleep we have lost: pre-industrial slumber in the british isles. *American Historical Review*, 106(2):343–386, 2001.

- [62] William J Elliott. Circadian variation in the timing of stroke onset. *Stroke*, 29(5):992–996, 1998.
- [63] Ericsson. Ericsson mobility report. 2017. Available at <https://www.ericsson.com/assets/local/mobility-report/documents/2017/ericsson-mobility-report-june-2017.pdf>.
- [64] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *KDD*, volume 96, pages 226–231, 1996.
- [65] B Etain, S Jamain, V Milhiet, M Lajnef, C Boudebese, A Dumaine, F Mathieu, A Gombert, K Ledudal, S Gard, et al. Association between circadian genes, bipolar disorders and chronotypes. *Chronobiology international*, 31(7):807–814, 2014.
- [66] Marco Fabbri, Chiara Mencarelli, Ana Adan, and Vincenzo Natale. Time-of-day and circadian typology on memory retrieval. *Biological rhythm research*, 44(1):125–142, 2013.
- [67] Mariana G Figueiro, Bryan Steverson, Judith Heerwagen, Kevin Kampschroer, Claudia M Hunter, Kassandra Gonzales, Barbara Plitnick, and Mark S Rea. The impact of daytime light exposures on sleep and mood in office workers. *Sleep Health*, 3(3):204–215, 2017.
- [68] MG Figueiro, R Hamner, A Bierman, and MS Rea. Comparisons of three practical field devices used to measure personal light exposures and activity levels. *Lighting Research & Technology*, 45(4):421–434, 2013.
- [69] f.lux Software LLC. f.lux. Available at <https://justgetflux.com>.
- [70] Aaron D Fobian, Kristin Avis, and David C Schwebel. The impact of media use on adolescent sleep efficiency. *Journal of developmental and behavioral pediatrics: JDBP*, 37(1):9, 2016.
- [71] Simon Folkard, David A Lombardi, and Mick B Spencer. Estimating the circadian rhythm in the risk of occupational injuries and accidents. *Chronobiology international*, 23(6):1181–1192, 2006.
- [72] Centers for Disease Control and Prevention (CDC). Insufficient sleep is a public health problem, 2015. Available at <https://www.cdc.gov/features/dssleep/>.

- [73] Russell G Foster and Till Roenneberg. Human responses to the geophysical daily, annual and lunar cycles. *Current biology*, 18(17):R784–R794, 2008.
- [74] National Sleep Foundation. Drowsy driving. Available at: <http://drowsydriving.org/about/facts-and-stats/>.
- [75] Ellen Frank, Stefanie Hlastala, Angela Ritenour, Patricia Houck, Xin Ming Tu, Timothy H Monk, Alan G Mallinger, and David J Kupfer. Inducing lifestyle regularity in recovering bipolar disorder patients: results from the maintenance therapies in bipolar disorder protocol. *Biological Psychiatry*, 41(12):1165–1173, 1997.
- [76] Ellen Frank, David J Kupfer, Michael E Thase, Alan G Mallinger, Holly A Swartz, Andrea M Fagiolini, Victoria Grochocinski, Patricia Houck, John Scott, Wesley Thompson, et al. Two-year outcomes for interpersonal and social rhythm therapy in individuals with bipolar i disorder. *Archives of general psychiatry*, 62(9):996–1004, 2005.
- [77] Ellen Frank, Isabella Soreca, Holly A Swartz, Andrea M Fagiolini, Alan G Mallinger, Michael E Thase, Victoria J Grochocinski, Patricia R Houck, and David J Kupfer. The role of interpersonal and social rhythm therapy in improving occupational functioning in patients with bipolar i disorder. *American Journal of Psychiatry*, 165(12):1559–1565, 2008.
- [78] Ellen Frank, Holly A Swartz, and David J Kupfer. Interpersonal and social rhythm therapy: managing the chaos of bipolar disorder. *Biological psychiatry*, 48(6):593–604, 2000.
- [79] Danielle J Frey, Pietro Badia, and Kenneth P Wright. Inter-and intra-individual variability in performance near the circadian nadir during sleep deprivation. *Journal of sleep research*, 13(4):305–315, 2004.
- [80] Mads Frost, Afsaneh Doryab, Maria Faurholt-Jepsen, Lars Vedel Kessing, and Jakob E Bardram. Supporting disease insight through data analysis: refinements of the monarca self-assessment system. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pages 133–142. ACM, 2013.
- [81] John F Gaski and Jeff Sagarin. Detrimental effects of daylight-saving time on sat scores. *Journal of Neuroscience, Psychology, & Economics*, 4(1):44, 2011.
- [82] John R Geddes and David J Miklowitz. Treatment of bipolar disorder. *The Lancet*, 381(9878):1672–1682, 2013.

- [83] Moniek Geerdink, Thijs J Walbeek, Domien GM Beersma, Vanja Hommes, and Marijke CM Gordijn. Short blue light pulses (30 min) in the morning support a sleep-advancing protocol in a home setting. *Journal of biological rhythms*, 31(5):483–497, 2016.
- [84] Jason R Gerstner and Jerry CP Yin. Circadian rhythms and memory formation. *Nature reviews. Neuroscience*, 11(8):577, 2010.
- [85] Namni Goel, Mathias Basner, Hengyi Rao, and David F Dinges. Circadian rhythms, sleep deprivation, and human performance. *Progress in molecular biology and translational science*, 119:155, 2013.
- [86] Google. Google Now. Available at <https://www.google.com/landing/now/>.
- [87] William Griesar, Daniel Zajdel, and Barry Oken. Nicotine effects on alertness and spatial attention in non-smokers. *Nicotine & tobacco research*, 4(2):185–194, 2002.
- [88] Agnes Gruenerbl, Venet Osmani, Gernot Bahle, Jose C Carrasco, Stefan Oehler, Oscar Mayora, Christian Haring, and Paul Lukowicz. Using smart phone mobility traces for the diagnosis of depressive and manic episodes in bipolar patients. In *Proceedings of the 5th Augmented Human International Conference*, page 38. ACM, 2014.
- [89] Agnes Grünerbl, Amir Muaremi, Venet Osmani, Gernot Bahle, Stefan Oehler, Gerhard Tröster, Oscar Mayora, Christian Haring, and Paul Lukowicz. Smartphone-based recognition of states and state changes in bipolar disorder patients. *IEEE Journal of Biomedical and Health Informatics*, 19(1):140–148, 2015.
- [90] Isabelle Guyon, Jason Weston, Stephen Barnhill, and Vladimir Vapnik. Gene selection for cancer classification using support vector machines. *Machine learning*, 46(1):389–422, 2002.
- [91] Marco Hafner, Martin Stepanek, Jirka Taylor, Wendy M Troxel, and Christian van Stolk. Why sleep matters—the economic costs of insufficient sleep. 2016.
- [92] Tian Hao, Guoliang Xing, and Gang Zhou. isleep: Unobtrusive sleep quality monitoring using smartphones. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems, SenSys '13*, pages 4:1–4:14, New York, NY, USA, 2013. ACM.

- [93] Mark D Harris, LORI B Siegel, and JEFFREY A Alloway. Gout and hyperuricemia. *American family physician*, 59(4):925–934, 1999.
- [94] Allison G Harvey. Sleep and circadian rhythms in bipolar disorder: seeking synchrony, harmony, and regulation. *American journal of psychiatry*, 165(7):820–829, 2008.
- [95] Allison G Harvey, Lisa S Talbot, and Anda Gershon. Sleep disturbance in bipolar disorder across the lifespan. *Clinical Psychology: Science and Practice*, 16(2):256–277, 2009.
- [96] Brant P Hasler, Leisha J Smith, Jennifer C Cousins, and Richard R Bootzin. Circadian rhythms, sleep, and substance abuse. *Sleep medicine reviews*, 16(1):67–81, 2012.
- [97] Ian Hindmarch, JS Kerr, and N Sherwood. The effects of alcohol and other drugs on psychomotor performance and cognitive function. *Alcohol and Alcoholism*, 26(1):71–79, 1991.
- [98] E Hoddes, V Zarcone, H Smythe, R Phillips, and WC Dement. Quantification of sleepiness: a new approach. *Psychophysiology*, 10(4):431–436, 1973.
- [99] Jim A Horne and Olov Östberg. A self-assessment questionnaire to determine morningness-eveningness in human circadian rhythms. *International journal of chronobiology*, 4(2):97–110, 1976.
- [100] Jim A Horne and Louise A Reyner. Sleep related vehicle accidents. *Bmj*, 310(6979):565–567, 1995.
- [101] Wenyu Huang, Kathryn Moynihan Ramsey, Biliانا Marcheа, and Joseph Bass. Circadian rhythms, sleep, and metabolism. *The Journal of clinical investigation*, 121(6):2133, 2011.
- [102] Peter J Huber et al. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964.
- [103] Jacob J Hughey. Machine learning identifies a compact gene set for monitoring the circadian clock in human blood. *Genome medicine*, 9(1):19, 2017.
- [104] Jacob J Hughey, Trevor Hastie, and Atul J Butte. Zeitzeiger: supervised

- learning for high-dimensional data from an oscillatory system. *Nucleic acids research*, 44(8):e80–e80, 2016.
- [105] Mari Hysing, Stle Pallesen, Kjell Morten Stormark, Reidar Jakobsen, Astri J Lundervold, and Børge Sivertsen. Sleep and use of electronic devices in adolescence: results from a large population-based study. *BMJ open*, 5(1):e006748, 2015.
- [106] Aarti Jagannath, Stuart N Peirson, and Russell G Foster. Sleep and circadian rhythm disruption in neuropsychiatric illness. *Current opinion in neurobiology*, 23(5):888–894, 2013.
- [107] Qiang Ji and Xiaojie Yang. Real-time eye, gaze, and face pose tracking for monitoring driver vigilance. *Real-Time Imaging*, 8(5):357–377, 2002.
- [108] Murray W Johns. A new method for measuring daytime sleepiness: the epworth sleepiness scale. *Sleep*, 14(6):540–545, 1991.
- [109] Lewis L Judd, Hagop S Akiskal, Pamela J Schettler, Jean Endicott, Andrew C Leon, David A Solomon, William Coryell, Jack D Maser, and Martin B Keller. Psychosocial disability in the course of bipolar i and ii disorders: a prospective, comparative, longitudinal study. *Archives of general psychiatry*, 62(12):1322–1330, 2005.
- [110] Tzyy-Ping Jung, Scott Makeig, Magnus Stensmo, and Terrence J Sejnowski. Estimating alertness from the eeg power spectrum. *IEEE transactions on biomedical engineering*, 44(1):60–69, 1997.
- [111] Thomas Kantermann, Myriam Juda, Martha Merrow, and Till Roenneberg. The human circadian clock’s seasonal adjustment is disrupted by daylight saving time. *Current Biology*, 17(22):1996–2000, 2007.
- [112] Matthew Kay, Kyle Rector, Sunny Consolvo, Ben Greenstein, Jacob O Wobbrock, Nathaniel F Watson, and Julie A Kientz. Pvt-touch: adapting a reaction time test for touchscreen devices. In *2013 7th International Conference on Pervasive Computing Technologies for Healthcare and Workshops*, pages 248–251. IEEE, 2013.
- [113] Paul Kelley, Steven W Lockley, Russell G Foster, and Jonathan Kelley. Synchronizing education to adolescent biology: ‘let teens sleep, start school later’. *Learning, Media and Technology*, 40(2):210–226, 2015.

- [114] Sat Bir S Khalsa, Megan E Jewett, Christian Cajochen, and Charles A Czeisler. A phase response curve to single bright light pulses in human subjects. *The Journal of physiology*, 549(3):945–952, 2003.
- [115] Nathaniel Kleitman. *Sleep and wakefulness*. University of Chicago Press, 1963.
- [116] Anders Knutsson, BjornG Jonsson, Torbjorn Akerstedt, and Kristina Orth-Gomer. Increased risk of ischaemic heart disease in shift workers. *The Lancet*, 328(8498):89–92, 1986.
- [117] Urban Knutsson, Jovanna Dahlgren, Claude Marcus, Sten Rosberg, Mikael BronnegaŁrd, Pontus Stierna, and Kerstin Albertsson-Wikland. Circadian cortisol rhythms in healthy boys and girls: relationship with age, growth, body composition, and pubertal development. *The Journal of Clinical Endocrinology & Metabolism*, 82(2):536–540, 1997.
- [118] Yong Seo Koo, Jin-Young Song, Eun-Yeon Joo, Heon-Jeong Lee, Eunil Lee, Sang-kun Lee, and Ki-Young Jung. Outdoor artificial light at night, obesity, and sleep health: Cross-sectional analysis in the koges study. *Chronobiology international*, 33(3):301–314, 2016.
- [119] Kurt Kräuchi, Christian Cajochen, Esther Werth, and Anna Wirz-Justice. Alteration of internal circadian phase relationships after morning versus evening carbohydrate-rich meals in humans. *Journal of biological rhythms*, 17(4):364–376, 2002.
- [120] Leon Kreitzman and Russell Foster. *The rhythms of life: The biological clocks that control the daily lives of every living thing*. Profile books, 2011.
- [121] Ondrej Krejcar, Jakub Jirka, and Dalibor Janckulik. Use of mobile phones as intelligent sensors for sound input analysis and sleep state detection. *Sensors*, 11(6):6037–6055, 2011.
- [122] Tatsuhiko Kubo, Kotaro Ozasa, Kazuya Mikami, Kenji Wakai, Yoshihisa Fujino, Yoshiyuki Watanabe, Tsuneharu Miki, Masahiro Nakao, Kyohei Hayashi, Koji Suzuki, et al. Prospective cohort study of the risk of prostate cancer among rotating-shift workers: findings from the japan collaborative cohort study. *American Journal of Epidemiology*, 164(6):549–555, 2006.
- [123] Kostadin Kushlev, Jason Proulx, and Elizabeth W Dunn. Silence your phones: Smartphone notifications increase inattention and hyperactivity

- symptoms. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 1011–1020. ACM, 2016.
- [124] Tuuli A Lahti, Sami Leppämäki, Jouko Lönnqvist, and Timo Partonen. Transitions into and out of daylight saving time compromise sleep and the rest-activity cycles. *BMC physiology*, 8(1):3, 2008.
- [125] Nicole Lamond and Drew Dawson. Quantifying the performance impairment associated with fatigue. *Journal of sleep research*, 8(4):255–262, 1999.
- [126] Elaine Waddington Lamont, Daniel Legault-Coutu, Nicolas Cermakian, and Diane B Boivin. The role of circadian clock genes in mental disorders. *Dialogues in clinical neuroscience*, 9(3):333, 2007.
- [127] Klodiana Lanaj, Russell E Johnson, and Christopher M Barnes. Beginning the workday yet already depleted? consequences of late-night smartphone use and sleep. *Organizational Behavior and Human Decision Processes*, 124(1):11–23, 2014.
- [128] Tanja Lange, Stoyan Dimitrov, and Jan Born. Effects of sleep and circadian rhythm on the human immune system. *Annals of the New York Academy of Sciences*, 1193(1):48–59, 2010.
- [129] Uichin Lee, Joonwon Lee, Minsam Ko, Changhun Lee, Yuhwan Kim, Subin Yang, Koji Yatani, Gahgene Gweon, Kyong-Mee Chung, and June-hwa Song. Hooked on smartphones: an exploratory study on smartphone overuse among college students. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pages 2327–2336. ACM, 2014.
- [130] Björn Lemmer. Cardiovascular chronobiology and chronopharmacology. *Biologic Rhythms in Clinical and Laboratory Medicine*. Berlin, Germany: Springer-Verlag, pages 418–427, 1992.
- [131] Francis Lévi, Christian Focan, Abdoulaye Karaboué, Virginie de la Valette, Danièle Focan-Henrard, Benoît Baron, Françoise Kreutz, and Sylvie Giacchetti. Implications of circadian clocks for the rhythmic delivery of cancer therapeutics. *Advanced drug delivery reviews*, 59(9):1015–1035, 2007.
- [132] KJ Lewis, RG Foster, and IR Jones. Is sleep disruption a trigger for postpartum psychosis? *The British Journal of Psychiatry*, 208(5):409–411, 2016.

- [133] Sylvia Loh, Nicole Lamond, Jill Dorrian, Gregory Roach, and Drew Dawson. The validity of psychomotor vigilance tasks of less than 10-minute duration. *Behavior Research Methods, Instruments, & Computers*, 36(2):339–346, 2004.
- [134] Hader A Mansour, Michael E Talkowski, Joel Wood, Kodavali V Chowdari, Lora McClain, Konasale Prasad, Debra Montrose, Andrea Fagiolini, Edward S Friedman, Michael H Allen, et al. Association study of 21 circadian genes with bipolar i disorder, schizoaffective disorder, and schizophrenia. *Bipolar disorders*, 11(7):701–710, 2009.
- [135] Aleksandar Matic, Martin Pielot, and Nuria Oliver. Boredom-computer interaction: Boredom proneness and the use of smartphone. In *Proc. of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 837–841. ACM, 2015.
- [136] Ritsuko Matsumura, Koichi Node, and Makoto Akashi. Estimation methods for human circadian phase by use of peripheral tissues. *Hypertension Research*, 39(9):623–627, 2016.
- [137] Mark Matthews, Saeed Abdullah, Geri Gay, and Tanzeem Choudhury. Tracking mental well-being: Balancing rich sensing and patient needs. *Computer*, 47(4):36–43, 2014.
- [138] Mark Matthews, Saeed Abdullah, Elizabeth Murnane, Stephen Voids, Tanzeem Choudhury, Geri Gay, and Ellen Frank. Development and evaluation of a smartphone-based measure of social rhythms for bipolar disorder. *Assessment*, 23(4):472–483, 2016.
- [139] Mark Matthews, Stephen Voids, Saeed Abdullah, Gavin Doherty, Tanzeem Choudhury, Sangha Im, and Geri Gay. In situ design for mental illness: considering the pathology of bipolar disorder in mhealth design. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services*, pages 86–97. ACM, 2015.
- [140] Ralph Oyini Mbouna, Seong G Kong, and Myung-Geun Chun. Visual analysis of eye state and head pose for driver alertness monitoring. *IEEE transactions on intelligent transportation systems*, 14(3):1462–1469, 2013.
- [141] Frank McCormick, John Kadzielski, Christopher P Landrigan, Brady Evans, James H Herndon, and Harry E Rubash. Surgeon fatigue: a prospective analysis of the incidence, risk, and intervals of predicted

- fatigue-related impairment in residents. *Archives of Surgery*, 147(5):430–435, 2012.
- [142] Sarah P Megdal, Candyce H Kroenke, Francine Laden, Eero Pukkala, and Eva S Schernhammer. Night work and breast cancer risk: a systematic review and meta-analysis. *European Journal of Cancer*, 41(13):2023–2032, 2005.
- [143] Jerome S Menet and Michael Rosbash. When brain clocks lose track of time: cause or consequence of neuropsychiatric disorders. *Current opinion in neurobiology*, 21(6):849–857, 2011.
- [144] Kathleen R Merikangas, Robert Jin, Jian-Ping He, Ronald C Kessler, Sing Lee, Nancy A Sampson, Maria Carmen Viana, Laura Helena Andrade, Chiyi Hu, Elie G Karam, et al. Prevalence and correlates of bipolar spectrum disorder in the world mental health survey initiative. *Archives of general psychiatry*, 68(3):241–251, 2011.
- [145] Jerrold S Meyer and Melinda A Novak. Minireview: hair cortisol: a novel biomarker of hypothalamic-pituitary-adrenocortical activity. *Endocrinology*, 153(9):4120–4127, 2012.
- [146] David J Miklowitz, Michael W Otto, Ellen Frank, Noreen A Reilly-Harrington, Jane N Kogan, Gary S Sachs, Michael E Thase, Joseph R Calabrese, Lauren B Marangell, Michael J Ostacher, et al. Intensive psychosocial intervention enhances functioning in patients with bipolar depression: results from a 9-month randomized controlled trial. *American Journal of Psychiatry*, 164(9):1340–1347, 2007.
- [147] Gregory E Miller, Edith Chen, and Eric S Zhou. If it goes up, must it come down? chronic stress and the hypothalamic-pituitary-adrenocortical axis in humans. *Psychological Bulletin*, 133(1):25–45, 2007.
- [148] Jun-Ki Min, Afsaneh Doryab, Jason Wiese, Shahriyar Amini, John Zimmerman, and Jason I Hong. Toss ‘n’ turn: smartphone as sleep and sleep quality detector. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 477–486. ACM, 2014.
- [149] Dana K Mirick and Scott Davis. Melatonin as a biomarker of circadian dysregulation. *Cancer Epidemiology and Prevention Biomarkers*, 17(12):3306–3313, 2008.

- [150] Jennifer A Mohawk, Carla B Green, and Joseph S Takahashi. Central and peripheral circadian clocks in mammals. *Annual review of neuroscience*, 35:445–462, 2012.
- [151] Timothy H Monk. A visual analogue scale technique to measure global vigor and affect. *Psychiatry research*, 27(1):89–99, 1989.
- [152] Timothy H Monk, Daniel J Buysse, Charles F Reynolds, and David J Kupfer. Circadian determinants of the postlunch dip in performance. *Chronobiology international*, 13(2):123–133, 1996.
- [153] John G Moore and Franz Halberg. Circadian rhythm of gastric acid secretion in active duodenal ulcer: chronobiological statistical characteristics and comparison of acid secretory and plasma gastrin patterns with healthy subjects and post-vagotomy and pyloroplasty patients. *Chronobiology international*, 4(1):101–110, 1987.
- [154] Daniel S Moran and Liran Mendal. Core temperature measurement. *Sports Medicine*, 32(14):879–885, 2002.
- [155] Amir Muaremi, Franz Gravenhorst, Agnes Grünerbl, Bert Arnrich, and Gerhard Tröster. Assessing bipolar episodes using speech cues derived from phone calls. In *International Symposium on Pervasive Computing Paradigms for Mental Health*, pages 103–114. Springer, 2014.
- [156] James E Muller, Peter H Stone, Zoltan G Turi, John D Rutherford, Charles A Czeisler, Corette Parker, W Kenneth Poole, Eugene Passamani, Robert Roberts, Thomas Robertson, et al. Circadian variation in the frequency of onset of acute myocardial infarction. *New England Journal of Medicine*, 313(21):1315–1322, 1985.
- [157] Mirjam Münch, Friedrich Linhart, Apiparn Borisuit, Susanne M Jaeggi, and Jean-Louis Scartezzini. Effects of prior light exposure on early evening performance, subjective sleepiness, and hormonal secretion. *Behavioral Neuroscience*, 126(1):196–203, 2012.
- [158] Mirjam Münch, Claudia Nowozin, Johannes Regente, Frederik Bes, Jan De Zeeuw, Sven Hädel, Amely Wahnschaffe, and Dieter Kunz. Blue-enriched morning light as a countermeasure to light at the wrong time: Effects on cognition, sleepiness, sleep, and circadian phase. *Neuropsychobiology*, 74(4):207–218, 2016.

- [159] Elizabeth L Murnane, Saeed Abdullah, Mark Matthews, Tanzeem Choudhury, and Geri Gay. Social (media) jet lag: how usage of social technology can modulate and reflect circadian rhythms. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 843–854. ACM, 2015.
- [160] Elizabeth L Murnane, Saeed Abdullah, Mark Matthews, Matthew Kay, Julie A Kientz, Tanzeem Choudhury, Geri Gay, and Dan Cosley. Mobile manifestations of alertness: connecting biological rhythms with patterns of smartphone app use. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services*, pages 465–477. ACM, 2016.
- [161] National Transportation Safety Board. Evaluation of u.s. department of transportation efforts in the 1990s to address operator fatigue. safety report ntsb/sr-99/01. 1999.
- [162] Reto Niedermann, Eva Wyss, Simon Annaheim, Agnes Psikuta, Sarah Davey, and René Michel Rossi. Prediction of human core body temperature using non-invasive measurement methods. *International journal of biometeorology*, 58(1):7–15, 2014.
- [163] Danielle M Novick, Holly A Swartz, and Ellen Frank. Suicide attempts in bipolar i and bipolar ii disorder: a review and meta-analysis of the evidence. *Bipolar disorders*, 12(1):1–9, 2010.
- [164] Institute of Medicine. *Sleep Disorders and Sleep Deprivation: An Unmet Public Health Problem*. The National Academies Press, Washington, DC, 2006.
- [165] Department of Motor Vehicles (New York State). Penalties for alcohol or drug-related violations. Available at <https://dmv.ny.gov/tickets/penalties-alcohol-or-drug-related-violations>.
- [166] Yosuke Okamoto and Seiji Nakagawa. Effects of daytime light exposure on cognitive brain activity as measured by the ERP P300. *Physiology & Behavior*, 138:313–318, January 2015.
- [167] World Health Organization. Investing in mental health. 2003. Available at http://www.who.int/mental_health/media/investing_mnh.pdf.
- [168] World Health Organization. *The global burden of disease: 2004 update*. 2008.

Available at http://www.who.int/healthinfo/global_burden_disease/2004_report_update/en/.

- [169] World Health Organization et al. Mental health action plan 2013-2020. 2013.
- [170] Venet Osmani, Alban Maxhuni, Agnes Grünerbl, Paul Lukowicz, Christian Haring, and Oscar Mayora. Monitoring activity of patients with bipolar disorder using smart phones. In *Proceedings of International Conference on Advances in Mobile Computing & Multimedia*, page 85. ACM, 2013.
- [171] George N Papadimitriou and Paul Linkowski. Sleep disturbance in anxiety disorders. *International Review of Psychiatry*, 17(4):229–236, 2005.
- [172] Dhanashree A Paranjpe and Vijay Sharma. Evolution of temporal order in living organisms. *Journal of Circadian Rhythms*, 3(1):7, 2005.
- [173] Stuart N Peirson and Russell G Foster. Sleep and circadian rhythm disruption in psychosis. *Circadian Medicine*, pages 271–282, 2015.
- [174] Pierre Philip and Torbjorn kerstedt. Transport and industrial safety, how are they affected by sleepiness and sleep restriction? *Sleep medicine reviews*, 10(5):347–356, 2006.
- [175] Pierre Philip, Patricia Sagaspe, Jacques Taillard, Nicholas Moore, Christian Guilleminault, Montserrat Sanchez-Ortuno, Torbjörn Akerstedt, and Bernard Bioulac. Fatigue, sleep restriction, and performance in automobile drivers: a controlled study in a natural environment. *SLEEP*, 26(3):277–284, 2003.
- [176] Philips. Actiwatch spectrum. Available at <http://www.usa.philips.com/healthcare/product/HC1046964/actiwatch-spectrum-activity-monitor>.
- [177] Mary L Phillips and David J Kupfer. Bipolar disorder diagnosis: challenges and future directions. *The Lancet*, 381(9878):1663–1671, 2013.
- [178] Martin Pielot, Tilman Dingler, Jose San Pedro, and Nuria Oliver. When attention is not scarce - detecting boredom from mobile phone usage. In *Proc. of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 825–836. ACM, 2015.

- [179] David T Plante and John W Winkelman. Sleep disturbance in bipolar disorder: therapeutic implications. *American Journal of Psychiatry*, 165(7):830–843, 2008.
- [180] John Platt et al. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers*, 10(3):61–74, 1999.
- [181] Michael I Posner. Attention: the mechanisms of consciousness. *Proceedings of the National Academy of Sciences*, 91(16):7398–7403, 1994.
- [182] Mashfiqui Rabbi, Shahid Ali, Tanzeem Choudhury, and Ethan Berke. Passive and in-situ assessment of mental and physical well-being using mobile sensors. In *Proceedings of the 13th international conference on Ubiquitous computing*, pages 385–394. ACM, 2011.
- [183] Shantha MW Rajaratnam and Josephine Arendt. Health in a 24-h society. *The Lancet*, 358(9286):999–1005, 2001.
- [184] Mark S Rea, Mariana G Figueiro, Andrew Bierman, and John D Bullough. Circadian light. *Journal of circadian rhythms*, 8(1):2, 2010.
- [185] MS Rea and MG Figueiro. Light as a circadian stimulus for architectural lighting. *Lighting Research & Technology*, page 1477153516682368, 2016.
- [186] Roberto Refinetti and Michael Menaker. The circadian rhythm of body temperature. *Physiology & behavior*, 51(3):613–637, 1992.
- [187] Russel J Reiter. Melatonin: the chemical expression of darkness. *Molecular and cellular endocrinology*, 79(1):C153–C158, 1991.
- [188] Russel J Reiter. The melatonin rhythm: both a clock and a calendar. *Cellular and Molecular Life Sciences*, 49(8):654–664, 1993.
- [189] LA Reyner, SJ Wells, V Mortlock, and JA Horne. ‘post-lunch’ sleepiness during prolonged, monotonous driving — effects of meal size. *Physiology & behavior*, 105(4):1088–1091, 2012.
- [190] Judith A Ricci, Elsbeth Chee, Amy L Lorandean, and Jan Berger. Fatigue in the us workforce: prevalence and implications for lost productive work time. *Journal of occupational and environmental medicine*, 49(1):1–10, 2007.

- [191] Basil Rigas, James Torosis, Carl J McDougall, Kirt J Vener, and Howard M Spiro. The circadian rhythm of biliary colic. *Journal of clinical gastroenterology*, 12(4):409–414, 1990.
- [192] Timothy Roehrs and Thomas Roth. Sleep, sleepiness, sleep disorders and alcohol use and abuse. *Sleep medicine reviews*, 5(4):287–297, 2001.
- [193] Till Roenneberg. *Internal time: Chronotypes, social jet lag, and why you're so tired*. Harvard University Press, 2012.
- [194] Till Roenneberg. Chronobiology: the human sleep project. *Nature*, 498(7455):427–428, 2013.
- [195] Till Roenneberg. Having trouble typing? what on earth is chronotype? *Journal of biological rhythms*, 30(6):487–491, 2015.
- [196] Till Roenneberg, Karla V Allebrandt, Martha Merrow, and Céline Vetter. Social jetlag and obesity. *Current Biology*, 22(10):939–943, 2012.
- [197] Till Roenneberg, Tim Kuehnle, Myriam Juda, Thomas Kantermann, Karla Allebrandt, Marijke Gordijn, and Martha Merrow. Epidemiology of the human circadian clock. *Sleep medicine reviews*, 11(6):429–438, 2007.
- [198] Till Roenneberg, Tim Kuehnle, Peter P Pramstaller, Jan Ricken, Miriam Havel, Angelika Guth, and Martha Merrow. A marker for the end of adolescence. *Current Biology*, 14(24):R1038–R1039, 2004.
- [199] Till Roenneberg, Anna Wirz-Justice, and Martha Merrow. Life between clocks: daily temporal patterns of human chronotypes. *Journal of biological rhythms*, 18(1):80–90, 2003.
- [200] Melanie Rüger and Frank AJL Scheer. Effects of circadian disruption on the cardiometabolic system. *Reviews in Endocrine and Metabolic Disorders*, 10(4):245–260, 2009.
- [201] Samuli I Saarni, Satu Viertiö, Jonna Perälä, Seppo Koskinen, Jouko Lönnqvist, and Jaana Suvisaari. Quality of life of people with schizophrenia, bipolar disorder and other psychotic disorders. *The British Journal of Psychiatry*, 197(5):386–394, 2010.
- [202] Robert L Sack, Dennis Auckley, R Robert Auger, Mary A Carskadon, Kenneth P Wright Jr, Michael V Vitiello, and Irina V Zhdanova. Circadian

rhythm sleep disorders: part i, basic principles, shift work and jet lag disorders. *Sleep*, 30(11):1460–1483, 2007.

- [203] Sohrab Saeb, Thaddeus R Cybulski, Stephen M Schueller, Konrad P Kording, and David C Mohr. Scalable passive sleep monitoring using mobile phones: Opportunities and obstacles. *Journal of medical Internet research*, 19(4), 2017.
- [204] RL Schimitt, JA Bragatti, R Levandovsky, MP Hidalgo, and MM Bianchin. Social rhythm and other chronobiological findings in juvenile myoclonic epilepsy. *Biological Rhythm Research*, 46(3):371–377, 2015.
- [205] Christina Schmidt, Fabienne Collette, Christian Cajochen, and Philippe Peigneux. A time to think: circadian rhythms in human cognition. *Cognitive neuropsychology*, 24(7):755–789, 2007.
- [206] Analyne M Schroeder and Christopher S Colwell. How to fix a broken clock. *Trends in Pharmacological Sciences*, 34(11):605–619, November 2013.
- [207] Brahim Selmaoui and Yvan Touitou. Reproducibility of the circadian rhythms of serum cortisol and melatonin in healthy subjects: a study of three different 24-h cycles over six weeks. *Life sciences*, 73(26):3339–3349, 2003.
- [208] M Katherine Shear, Jeff Randall, Timothy H Monk, Angela Ritenour, Xin Tu Frank, Charles Reynolds, David J Kupfer, et al. Social rhythm in anxiety disorder patients. *Anxiety*, 1(2):90–95, 1994.
- [209] Shu-qun Shi, Tasneem S Ansari, Owen P McGuinness, David H Wasserman, and Carl Hirschie Johnson. Circadian disruption leads to insulin resistance and obesity. *Current Biology*, 23(5):372–381, 2013.
- [210] Benjamin L Smarr, Kimberly J Jennings, Joseph R Driscoll, and Lance J Kriegsfeld. A time to remember: The role of circadian clocks in learning and memory. *Behavioral neuroscience*, 128(3):283, 2014.
- [211] Alex J Smola and Bernhard Schölkopf. A tutorial on support vector regression. *Statistics and computing*, 14(3):199–222, 2004.
- [212] Michael H Smolensky, Bjoern Lemmer, and Alain E Reinberg. Chronobiology and chronotherapy of allergic rhinitis and bronchial asthma. *Advanced drug delivery reviews*, 59(9):852–882, 2007.

- [213] Michael H Smolensky and Nicholas A Peppas. Chronobiology, drug delivery, and chronotherapeutics. *Advanced Drug Delivery Reviews*, 59(9-10):828–851, August 2007.
- [214] Cary Stothart, Ainsley Mitchum, and Courtney Yehnert. The attentional cost of receiving a cell phone notification. *Journal of experimental psychology: human perception and performance*, 41(4):893, 2015.
- [215] Kurt Straif, Robert Baan, Yann Grosse, Béatrice Secretan, Fatiha El Ghissassi, Véronique Bouvard, Andrea Altieri, Lamia Benbrahim-Tallaa, Vincent Coglianò, WHO International Agency for Research on Cancer Monograph Working Group, et al. Carcinogenicity of shift-work, painting, and fire-fighting, 2007.
- [216] Jacques Taillard, Patrick Sanchez, Patrick Lemoine, and Jacques Mouret. Heart rate circadian rhythm as a biological marker of desynchronization in major depression: A methodological and preliminary report. *Chronobiology international*, 7(4):305–316, 1990.
- [217] Daniel J Taylor and Adam D Bramoweth. Patterns and consequences of inadequate sleep in college students: substance use and motor vehicle accidents. *Journal of Adolescent Health*, 46(6):610–612, 2010.
- [218] Sara Thomée, Annika Härenstam, and Mats Hagberg. Mobile phone use and stress, sleep disturbances, and symptoms of depression among young adults—a prospective cohort study. *BMC public health*, 11(1):66, 2011.
- [219] Yvan Touitou and Erhard Haus. *Biologic rhythms in clinical and laboratory medicine*. Springer Science & Business Media, 2012.
- [220] Wout JM Van Bommel. Non-visual biological effect of lighting and the practical meaning for lighting for work. *Applied ergonomics*, 37(4):461–466, 2006.
- [221] Hans PA Van Dongen, Greg Maislin, Janet M Mullington, and David F Dinges. The cumulative cost of additional wakefulness: dose-response effects on neurobehavioral functions and sleep physiology from chronic sleep restriction and total sleep deprivation. *Sleep*, 26(2):117–126, 2003.
- [222] HPA Van Dongen, Maurice D Baynard, Greg Maislin, and David F Dinges. Systematic interindividual differences in neurobehavioral impairment from sleep loss: evidence of trait-like differential vulnerability. *Sleep*, 27(3):423–433, 2004.

- [223] HPA Van Dongen and David F Dinges. Circadian rhythms in sleepiness, alertness, and performance. In *Principles and practice of sleep medicine*, pages 435–443. 4, 2005.
- [224] Theun Pieter van Tienoven, Joeri Minnen, Sarah Daniels, Djiwo Weenas, Anke Raaijmakers, and Ignace Glorieux. Calculating the social rhythm metric (srm) and examining its use in interpersonal social rhythm therapy (ipsrt) in a healthy population study. *Behavioral Sciences*, 4(3):265–277, 2014.
- [225] Nicola Vanello, Andrea Guidi, Claudio Gentili, Sandra Werner, Gilles Bertschy, Gaetano Valenza, Antonio Lanata, and Enzo Pasquale Scilingo. Speech analysis for mood state characterization in bipolar patients. In *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*, pages 2104–2107. IEEE, 2012.
- [226] GH Vázquez, F Kapczinski, PV Magalhaes, R Córdoba, C Lopez Jaramillo, AR Rosa, M Sanchez de Carmona, M Tohen, The Ibero-American Network on Bipolar, et al. Stigma and functioning in patients with bipolar disorder. *Journal of affective Disorders*, 130(1):323–327, 2011.
- [227] Céline Vetter. *Clocks in action: Exploring the impact of internal time in real life*. Südwestdeutscher Verlag für Hochschulschriften, 2011.
- [228] Céline Vetter, Myriam Juda, and Till Roenneberg. The influence of internal time, time awake, and sleep duration on cognitive performance in shiftworkers. *Chronobiology International*, 29(8):1127–1138, 2012.
- [229] Jacqueline M Vink, Jacqueline M Vink, Alexia S Groot, Gerard A Kerkhof, and Dorret I Boomsma. Genetic analysis of morningness and eveningness. *Chronobiology international*, 18(5):809–822, 2001.
- [230] Tri Vu, Feng Lin, Nabil Alshurafa, and Wenyao Xu. Wearable food intake monitoring technologies: A comprehensive review. *Computers*, 6(1):4, 2017.
- [231] Manav V Vyas, Amit X Garg, Arthur V Iansavichus, John Costella, Allan Donner, Lars E Laugsand, Imre Janszky, Marko Mrkobrada, Grace Parra, and Daniel G Hackam. Shift work and vascular events: systematic review and meta-analysis. *Bmj*, 345:e4800, 2012.
- [232] David T Wagner, Christopher M Barnes, Vivien KG Lim, and D Lance Ferris. Lost sleep and cyberloafing: Evidence from the laboratory and

- a daylight saving time quasi-experiment. *Journal of Applied Psychology*, 97(5):1068, 2012.
- [233] XS Wang, MEG Armstrong, BJ Cairns, TJ Key, and RC Travis. Shift work and chronic disease: the epidemiological evidence. *Occupational medicine*, 61(2):78–89, 2011.
- [234] Sophie M T Wehrens, Skevoulla Christou, Cheryl Isherwood, Benita Middleton, Michelle A Gibbs, Simon N Archer, Debra J Skene, and Jonathan D Johnston. Meal Timing Regulates the Human Circadian System. *Current biology : CB*, 27(12):1768–1775, June 2017.
- [235] Jacob Weisberg. We are hopelessly hooked. *The New York Review of Books*, 63(3):6–9, 2016.
- [236] Nancy Wesensten, William Killgore, and Thomas Balkin. Performance and alertness effects of caffeine, dextroamphetamine, and modafinil during sleep deprivation. *Journal of sleep research*, 14(3):255–266, 2005.
- [237] Ann Williamson, David A Lombardi, Simon Folkard, Jane Stutts, Theodore K Courtney, and Jennie L Connor. The link between fatigue and safety. *Accident Analysis & Prevention*, 43(2):498–515, 2011.
- [238] Marc Wittmann, Jenny Dinich, Martha Mero, and Till Roenneberg. Social jetlag: misalignment of biological and social time. *Chronobiology international*, 23(1-2):497–509, 2006.
- [239] Gretchen Wolff and Karyn A Esser. Scheduled exercise phase shifts the circadian clock in skeletal muscle. *Medicine and science in sports and exercise*, 44(9):1663, 2012.
- [240] Kenneth P Wright Jr, Joseph T Hull, Rod J Hughes, Joseph M Ronda, and Charles A Czeisler. Sleep and wakefulness out of phase with internal biological time impairs learning in humans. *Journal of Cognitive Neuroscience*, 18(4):508–521, 2006.
- [241] Katharina Wulff, Derk-Jan Dijk, Benita Middleton, Russell G Foster, and Eileen M Joyce. Sleep and circadian rhythm disruption in schizophrenia. *The British Journal of Psychiatry*, 200(4):308–316, 2012.
- [242] Danny Wyatt, Tanzeem Choudhury, Jeff Bilmes, and James A Kitts. Inferring colocation and conversation networks from privacy-sensitive audio

with implications for computational social science. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(1):7, 2011.

- [243] Xiaojiang Xu, Anthony J Karis, Mark J Buller, and William R Santee. Relationship between core temperature, skin temperature, and heat flux during exercise in heat. *European journal of applied physiology*, 113(9):2381–2389, 2013.
- [244] Yujiro Yamanaka, Satoko Hashimoto, Yusuke Tanahashi, Shin-ya Nishide, Sato Honma, and Ken-ichi Honma. Physical exercise accelerates reentrainment of human sleep-wake cycle but not of plasma melatonin rhythm to 8-h phase-advanced sleep schedule. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 298(3):R681–R691, 2010.
- [245] Kok-Meng Yee, Stuart D Pringle, and Allan D Struthers. Circadian variation in the effects of aldosterone blockade on heart rate variability and qt dispersion in congestive heart failure. *Journal of the American College of Cardiology*, 37(7):1800–1807, 2001.
- [246] Andrei Zavada, Marijke CM Gordijn, Domien GM Beersma, Serge Daan, and Till Roenneberg. Comparison of the munich chronotype questionnaire with the horne-östberg’s morningness-eveningness score. *Chronobiology international*, 22(2):267–278, 2005.
- [247] Phyllis C Zee, Hrayr Attarian, and Aleksandar Videnovic. Circadian rhythm abnormalities. *Continuum: Lifelong Learning in Neurology*, 19(1 Sleep Disorders):132, 2013.
- [248] Jamie M Zeitzer, Derk-Jan Dijk, Richard E Kronauer, Emery N Brown, and Charles A Czeisler. Sensitivity of the human circadian pacemaker to nocturnal light: melatonin phase resetting and suppression. *The Journal of physiology*, 526(3):695–702, 2000.
- [249] Giulia Zerbini, Vincent Vinne, Lana K M Otto, Thomas Kantermann, Wim P Krijnen, Till Roenneberg, and Martha Merrow. Lower school performance in late chronotypes: underlying factors and mechanisms. *Nature Publishing Group*, 7(1):4385, June 2017.
- [250] PZ Zimmet, JR Wall, R Rome, L Stimmler, and RJ Jarrett. Diurnal variation in glucose tolerance: associated changes in plasma insulin, growth hormone, and non-esterified. *Br Med J*, 1(5906):485–488, 1974.