Enabling Scalable Sleep Monitoring with Mobile Sensing and Machine Learning

Priyanka Mary Mammen
University of Massachusetts Amherst
pmammen@cs.umass.edu

Prashant Shenoy
University of Massachusetts Amherst
shenoy@cs.umass.edu

ABSTRACT

Sleep is a critical component for overall health of a human being. Despite it's importance, a majority of the population is sleep deprived leading to several physical and mental health issues. Traditional methods of sleep monitoring are expensive and not scalable, limiting access to important health information. The advent of the Internet of Things and mobile sensing devices provides an opportunity for more accessible and scalable sleep monitoring. Community-scale sensing has the potential to enable aggregate public health monitoring and informed decision-making for individuals. However, scaling mobile health sensing technologies to the community level presents several challenges including data availability, model generalizability, robustness etc.

This thesis focuses on addressing the challenges of sleep monitoring at the community level by developing non-intrusive, scalable, personalizable, and robust sleep detection techniques. The first technique, called *WiSleep* , a system that utilizes an unsupervised machine learning approach that detects sleep durations from WiFi activity of mobile devices. *WiSleep* leverages the strong correlation between a phone's network activity and sleep periods and uses an ensemble of Bayesian models designed to handle irregular sleep patterns. The second technique, called *SleepLess*, is a semi-supervised machine learning approach that enables personalized sleep estimations for users without labeled data. Finally, I propose an approach combining uncertainity quantification with explainability to handle prediction inaccuracies from sleep prediction models. Overall, this research aims to provide innovative solutions to the challenges of mobile sleep monitoring and contribute to the broader field of mobile health sensing.

1 INTRODUCTION

Sleep is a vital activity that significantly impacts human well-being, productivity and performance [23]. Prior research has shown that 30% of the adult population does not get enough sleep, with many adults sleeping less than 7 hours per day [8, 16]. Both work-related stress and the increasing use of mobile devices throughout the day, particularly in the evenings, have increased sleep disorders [26]. The repercussions of sleep deprivation leading to serious health consequences such as heart disease, stroke, and depression [1, 20] has become a public health burden. The American Academy of Pediatrics confirms sleep deprivation as a public health epidemic, especially among students [10, 20].

Numerous solutions have emerged for sleep monitoring. Polysomnography is a gold standard in medical research [24] practical for short-term monitoring. Contactless methods utilizing doppler radar and RF signals have been proposed [15, 21], but they require specific instrumentation in building infrastructures. With the recent advances in Internet of Things (IoT), it is now easy to deploy low-cost sensing mechanisms for sleep monitoring. IoT-based sensing (known as mobile health sensing) can take two forms: *wearable sensing*, which leverages on-body sensors [18], and *contactless sensing*, which leverages sensors in the environment or passive sensors [30] such as smartphones [2]. Smartphones are ubiquitous and can offer affordable sleep sensing through microphones, cameras, phone activity and screen usage [5, 11, 13, 19, 22, 34].

Sleep is an intimate experience; hence many sleep monitoring technologies are highly personalized for individual use. Monitoring data sources specific to sleep are challenging to acquire for public health understanding and benefits [20]. Such information could benefit professional health administrators to keep abreast of a community's needs and well-being. In particular, college students residing in campus dormitories make an insightful study population of irregular sleepers due to overwhelming academic demands. Further, many college campuses, such as in the United States, are known for their social events during the semester. The active party culture exacerbates bad sleeping habits among students [27]. These irregular habits can significantly and negatively impact students' concentration and academic performance [17].

While mobile health (mHealth) technologies have revolutionized our ability for personalized sleep monitoring, the next frontier in mHealth involves scaling the sleep monitoring technologies to monitor a community or population of users. The combination of community scale monitoring, which enables aggregate sleep monitoring of large groups and personalized health monitoring, which enables fine-grain sleep monitoring of individuals, will provide insights into broader health problems. Realizing this vision involves addressing new challenges that require mobile health technologies to be scalable, easily personalizable and robust while addressing the needs of both personalized and large scale sleep monitoring. Three main challenges are (1) *data availability* - outside laboratory settings, collecting clean and noise-free labeled data at a large scale is difficult and expensive. (2) *generalizability* - while we develop sleep prediction models using mobile device data, one prediction model will not be able to give accurate predictions for all users in a population due to the heterogeneity in the device and human behaviors. (3) *robustness* - data collected from mobile devices is often noisy, has outliers, and may experience data drifts. These factors can significantly impact the prediction model's performance, leading to inaccurate predictions, and decreased reliability.

In this thesis, I seek to address the above challenges in developing scalable and personalizable mobile health sensing techniques that can be employed both to monitor the sleep activity of a community (a public health concern) and those of individuals (a personalized health monitoring issue).

2 BACKGROUND

2.1 Smartphone based Sleep Sensing

Smartphones are ubiquitous and have become a necessary part of our everyday life. By leveraging the device's advanced sensing and computational capabilities, researchers and healthcare providers can gain valuable insights into human behavior and health, which can lead to improved outcomes and quality of life for individuals. One of the significant benefits of smartphones is their ability to collect data in real-time and over extended periods. By using the device's built-in accelerometer sensor, researchers can track physical activities such as walking, running, and climbing stairs [3]. However, developing a robust classifier for physical activity tracking is difficult due to the differences across users and smartphone devices.

For sleep sensing using smartphones, researchers use a combination of techniques used for physical activity tracking, mobility patterns and phone activity and ambient condition detection. Some of the efforts utilize an array of phone sensors such as accelerometer [19, 25], light sensor [19, 25], microphone [13, 19, 22, 25], proximity sensor [19], and WiFi network activity rates [31]. The works above have proven rich (unlabeled) data we can use to predict sleep. However, collecting labeled ground truth of users continue to be a big challenge [6].

2.2 Machine Learning for Sleep Monitoring

Machine learning techniques are widely used for sleep monitoring applications Most of the techniques described in the prior section are based on supervised learning approaches, whereby models require large amounts of training data to build accurate prediction models. Several works developed separate models for each user in the study which require at least 2 weeks of labelled data with sleep/wake-up estimation errors less than 40 minutes [31]. On the other hand, unsupervised methods have been explored specifically to do away with requiring training data. For example, Cuttone et. al develop a Bayesian approach to infer bed time and wake-up time from smart-phone screen events. Although convenient, these works have reported bed time and wake-time errors in the range of 1-2 hours and they don't offer any form of personalization.

Our investigation on semi-supervised learning approaches began via Self Training, followed by other standard techniques such as Co-Training, Auto Encoder, Data Generative, and Adversarial Training. These approaches have been adopted by prior work to solve sleep stage classification [14, 29, 33, 36]. In understanding self-training, Zhang et al. reported error accumulation as a potential problem [32]. In contrast, other works have reported lesser error accumulation with co-training [28, 35], including successful detection of everyday human behavior such as walking, running, and climbing stairs [12]. In fact, much work in human activity recognition has utilized adversarial learning [7] and autoencoder [4, 9] to develop a generalizable and robust classification model for everyday human behavior.

3 CONTRIBUTIONS

3.1 Community Scalable Sleep Monitoring

In the first component of my thesis, I present the design and implementation of WiSleep, a novel population-scale sensing system for sleep monitoring of a large cohort of on-campus students. The primary objective of this work is to address the challenge of scalability in sleep monitoring by adopting a network-side approach that leverages WiFi and smartphones.

The key insight of WiSleep is that a strong correlation exists between a phone's network activity and sleep periods. This correlation is analogous to the well-established relationship between a user's screen usage and sleep patterns, demonstrated in prior client-side methods. WiSleep takes advantage of the coarse-grain network activity data available from WiFi Access Points (APs), which can detect and record AP association and disassociation events, to infer when a user is asleep. WiSleep uses an unsupervised learning approach based on an ensemble of Bayesian models that are designed to handle irregular sleep patterns common among college students.

By adopting a network-side approach, WiSleep overcomes many limitations of traditional sleep monitoring methods, which are often cumbersome and intrusive and can only be applied to a small number of participants. WiSleep has the potential to provide valuable insights into sleep patterns and behaviors for large-scale studies, which could ultimately lead to improved sleep health and well-being for individuals and populations.

3.2 Personalized Sleep Monitoring

While public-health monitoring of communities requires general models for health sensing that works across a range of users, it is also desirable that these models be accurate and handle idiosyncrasies exhibited by the users. However, the major challenge here is the collection of labeled data from the users to develop personalized models. As the unlabelled data from smartphones is easy to gather, in the second component, I design *SleepLess*, a semi-supervised machine learning approach to enable personalization in the sleep prediction models using WiFi activity data collected from smartphones. Specifically, it uses a pre-trained model on an existing set of users to produce pseudo-labels for unlabeled data of a new user and achieves personalization by fine-tuning over selectively picking the pseudo-labels. Our user study among 23 users found model yielding around 96% accuracy, between 12-27 minutes of sleep time error and 18-25 minutes of wake time error. Comparison against other approaches that sought to predict with fewer labeled data found, similarly yielding best performance.

3.3 Uncertainty Aware Sleep Monitoring [Proposed Work]

Accurate predictions are essential for any human activity recognition model especially when it is used for clinical decision making or personal recommendation systems. However, due to noisy data or outlier data, sleep prediction models using smartphone data may produce inaccurate predictions. To effectively handle such inaccuracies, we can rely on uncertainty quantification and model explainability techniques. By identifying uncertain predictions and distinguishing between in-distribution and out-of-distribution data, we can selectively pick highly confident results and also work on model refinements and improvements. I propose to compare various uncertainty quantification and model explainability techniques on various sleep datasets and identify/develop a suitable method to handle inaccurate predictions for sleep monitoring.

REFERENCES

[1] Bruce M Altevogt, Harvey R Colten, et al. 2006. *Sleep disorders and sleep deprivation: an unmet public health problem.* National Academies Press.

[2] Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra, and Jorge L Reyes-Ortiz. 2012. Human activity recognition on smartphones using a multiclass hardware-friendly support vector machine. In *International workshop on ambient assisted living*. Springer, 216–223.

[3] Muhammad Arif, Mohsin Bilal, Ahmed Kattan, and S Iqbal Ahamed. 2014. Better physical activity classification using smartphone acceleration sensor. *Journal of medical systems* 38 (2014), 1–10.

[4] Sourav Bhattacharya, Petteri Nurmi, Nils Hammerla, and Thomas Plötz. 2014. Using unlabeled data in a sparse-coding framework for human activity recognition. *Pervasive and Mobile Computing* 15 (2014), 242–262.

[5] A. Cuttone, P. Bakgaard, V. Sekara, H. Jonsson, JE Larsen, and S. Lehmann. 2017. SensibleSleep: A Bayesian Model for Learning Sleep Patterns from Smartphone Events. *PLoS ONE* 12, 1 (2017).

[6] Massimiliano De Zambotti, Nicola Cellini, Aimee Goldstone, Ian M Colrain, and Fiona C Baker. 2019. Wearable sleep technology in clinical and research settings. *Medicine and science in sports and exercise* 51, 7 (2019), 1538.

[7] Abu Zaher Md Faridee, Avijoy Chakma, Archan Misra, and Nirmalya Roy. 2022. STranGAN: Adversarially-learnt Spatial Transformer for scalable human activity recognition. *Smart Health* 23 (2022), 100226.

[8] Centers for Disease Control, Prevention (CDC, et al. 2009. Perceived insufficient rest or sleep among adults-United States, 2008. *MMWR. Morbidity and mortality weekly report* 58, 42 (2009), 1175.

[9] Anupriya Gogna and Angshul Majumdar. 2016. Semi supervised autoencoder. In *International Conference on Neural Information Processing*. Springer, 82–89.

[10] Adolescent Sleep Working Group et al. 2014. School start times for adolescents. *Pediatrics* 134, 3 (2014), 642–649.

[11] Weixi Gu, Longfei Shangguan, Zheng Yang, and Yunhao Liu. 2015. Sleep hunter: Towards fine grained sleep stage tracking with smartphones. *IEEE Transactions on Mobile Computing* 15, 6 (2015), 1514–1527.

[12] Donghai Guan, Weiwei Yuan, Young-Koo Lee, Andrey Gavrilov, and Sungyoung Lee. 2007. Activity recognition based on semi-supervised learning. In *13th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA 2007)*. IEEE, 469–475.

[13] Tian Hao, Guoliang Xing, and Gang Zhou. 2013. iSleep: unobtrusive sleep quality monitoring using smartphones. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*. 1–14.

[14] Bai Haoran and Lu Guanze. 2021. Semi-Supervised End-to-End Automatic Sleep Stage Classification Based on Pseudo-Label. In *2021 IEEE International Conference on Power Electronics, Computer Applications (ICPECA)*. IEEE, 83–87.

[15] Chen-Yu Hsu, Aayush Ahuja, Shichao Yue, Rumen Hristov, Zachary Kabelac, and Dina Katabi. 2017. Zero-Effort In-Home Sleep and Insomnia Monitoring using Radio Signals. In *Proc. ACM Interact. Mob. Wearable Ubiquitous Technology*.

[16] Patrick M Krueger and Elliot M Friedman. 2009. Sleep duration in the United States: a cross-sectional population-based study. *American journal of epidemiology* 169, 9 (2009), 1052–1063.

[17] Ganpat Maheshwari and Faizan Shaukat. 2019. Impact of poor sleep quality on the academic performance of medical students. *Cureus* 11, 4 (2019).

[18] David Metcalf, Sharlin TJ Milliard, Melinda Gomez, and Michael Schwartz. 2016. Wearables and the internet of things for health: Wearable, interconnected devices promise more efficient and comprehensive health care. *IEEE pulse* 7, 5 (2016), 35–39.

[19] Jun-Ki Min, Afsaneh Doryab, Jason Wiese, Shahriyar Amini, John Zimmerman, and Jason I Hong. 2014. Toss'n'turn: smartphone as sleep and sleep quality detector. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 477–486.

[20] Geraldine S Perry, Susheel P Patil, and Letitia R Presley-Cantrell. 2013. Raising awareness of sleep as a healthy behavior. *Preventing chronic disease* 10 (2013).

[21] Tauhidur Rahman, Alexander T Adams, Ruth Vinisha Ravichandran, Mi Zhang, Shwetak N Patel, Julie A Kientz, and Tanzeem Choudhury. 2015. Dopplesleep: A contactless unobtrusive sleep sensing system using short-range doppler radar. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 39–50.

[22] Yanzhi Ren, Chen Wang, Jie Yang, and Yingying Chen. 2015. Fine-grained sleep monitoring: Hearing your breathing with smartphones. In *Computer Communications (INFOCOM), 2015 IEEE Conference on*. IEEE, 1194–1202.

[23] Mark R Rosekind, Kevin B Gregory, Melissa M Mallis, Summer L Brandt, Brian Seal, and Debra Lerner. 2010. The cost of poor sleep: workplace productivity loss and associated costs. *Journal of Occupational and Environmental Medicine* 52, 1 (2010), 91–98.

[24] Warren R Ruehland, Fergal J O'Donoghue, Robert J Pierce, Andrew T Thornton, Parmjit Singh, Janet M Copland, Bronwyn Stevens, and Peter D Rochford. 2011. The 2007 AASM recommendations for EEG electrode placement in polysomnography: impact on sleep and cortical arousal scoring. *Sleep* 34, 1 (2011), 73–81.

[25] Sohrab Saeb, Thaddeus R Cybulski, Stephen M Schueller, Konrad P Kording, and David C Mohr. 2017. Scalable passive sleep monitoring using mobile phones: opportunities and obstacles. *Journal of medical Internet research* 19, 4 (2017), e118.

[26] Sara Thomée, Annika Härenstam, and Mats Hagberg. 2011. Mobile phone use and stress, sleep disturbances, and symptoms of depression among young adults-a prospective cohort study. *BMC public health* 11, 1 (2011), 66.

[27] Karen Vail-Smith, W Michael Felts, and Craig Becker. 2009. Relationship between sleep quality and health risk behaviors in undergraduate college students. *College Student Journal* 43, 3 (2009), 924–930.

[28] Wei Wang and Zhi-Hua Zhou. 2007. Analyzing co-training style algorithms. In *European conference on machine learning*. Springer, 454–465.

[29] I Zeki Yalniz, Hervé Jégou, Kan Chen, Manohar Paluri, and Dhruv Mahajan. 2019. Billion-scale semi-supervised learning for image classification. *arXiv preprint arXiv:1905.00546* (2019).

[30] Kuo-Hui Yeh. 2016. A secure IoT-based healthcare system with body sensor networks. *IEEE Access* 4 (2016), 10288–10299.

[31] Camellia Zakaria, Gizem Yilmaz, Priyanka Mammen, Michael Chee, Prashant Shenoy, and Rajesh Balan. 2022. SleepMore: Sleep Prediction at Scale via Multi-Device WiFi Sensing. *arXiv preprint arXiv:2210.14152* (2022).

[32] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. 2021. Understanding deep learning (still) requires rethinking generalization. *Commun. ACM* 64, 3 (2021), 107–115.

[33] Chuanhao Zhang, Wenwen Yu, Yamei Li, Hongqiang Sun, Yuan Zhang, and Maarten De Vos. 2022. CMS2-Net: Semi-Supervised Sleep Staging for Diverse Obstructive Sleep Apnea Severity. *IEEE Journal of Biomedical and Health Informatics* 26, 7 (2022), 3447–3457.

[34] Chen Zhenyu, Nicholas Lane, Guiseppe Cardone, Mu Lin, Tanzeem Choudhury, and Andrew Campbell. 2013. Unobtrusive Sleep Monitoring Using Smartphones. In *Proceedings of Pervasive Health*.

[35] Zhi-Hua Zhou and Ming Li. 2010. Semi-supervised learning by disagreement. *Knowledge and Information Systems* 24, 3 (2010), 415–439.

[36] Yang Zou, Zhiding Yu, BVK Kumar, and Jinsong Wang. 2018. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European conference on computer vision (ECCV)*. 289–305.