Personalized User Modelling for Context-Aware Lifestyle Recommendations to Improve Sleep

Vaibhav Pandey*
vaibhap1@uci.edu
University of California, Irvine
Irvine, CA, USA

Dhruv Deepak Upadhyay*
ddupadhy@uci.edu
University of California, Irvine
Irvine, CA, USA

Nitish Nag
University of California, Irvine
Irvine, CA, USA
nagn@uci.edu

Ramesh Jain
University of California, Irvine
Irvine, CA, USA
jain@ics.uci.edu

ABSTRACT

Sleep is a significant contributor to leading a healthy lifestyle. Each day, most people go to sleep without any idea about how their night's rest will be or how they can leverage their data to improve it. For an activity that humans spend near a third of their life doing, there is a surprising amount of mystery around it. Despite current research, creating personalized sleep models in real-world settings has been challenging. Existing literature provides several connections between daily activities and sleep quality. Unfortunately, these insights do not generalize well in many individuals. For these reasons, it is essential to create a data-driven personalized sleep model. This research proposes a sleep model that captures causal relationships between daily activities and sleep quality and presents the user with specific feedback recommendations to improve sleep quality. Using N-of-1 experiments on longitudinal user data and event mining, the model generates a probabilistic understanding between lifestyle choices (exercise, eating, circadian rhythm, environmental selection) and their respective impact on sleep quality. Our experimental results identified and quantified relationships while extracting confounding variables through a causal framework. We then utilize the generated model to provide lifestyle recommendations to optimize sleep outcomes in a context-aware health recommendation system.

KEYWORDS

sleep recommendation system, health recommendation system, context aware, event mining, hypothesis verification, N-of-1 experiments, n=1, causal inference, healthy lifestyle, personalized health, precision health.

ACM Reference Format:
Vaibhav Pandey, Dhruv Deepak Upadhyay, Nitish Nag, and Ramesh Jain. 2020. Personalized User Modelling for Context-Aware Lifestyle Recommendations to Improve Sleep. In *Proceedings of the 5th International Workshop on Health Recommender Systems co-located with 14th ACM Conference on Recommender Systems (HealthRecSys'20), Online, Worldwide, September 26, 2020.* , 7 pages.

*Both authors contributed equally to this research.

HealthRecSys '20, September 26, 2020, Online, Worldwide

1 INTRODUCTION

Lifestyle factors have a high impact on our health outcomes. Our eating, sleeping, and movement patterns determine large parts of our short and long-term health [Hills et al. 2015; Nag 2020; Nag et al. 2018; Shochat 2012]. Keeping track of how we behave in different contextual situations, and the impact these behavioral patterns have on our health is difficult for medical professionals and individuals. At the same time, with the increasing prevalence of chronic diseases such as diabetes and hypertension, understanding the effect of lifestyle on different aspects of our health becomes a critical research challenge [Lee et al. 2017; Sarris et al. 2014]. The rising popularity of wearable and IoT devices provide us with an opportunity to address this problem computationally. There is a considerable body of research dedicated to finding and logging life events from multimodal data streams [Gurrin et al. 2014; Oh and Jain 2017; Sellen and Whittaker 2010]. A multitude of consumer devices such as smartwatches and smart home systems measure aspects of our daily life as events and data streams and control our local environment [Casino et al. 2018]. Using the data streams and events captured by these devices, we can find recurring behavioral patterns and associated health outcomes to create an explainable rule-based model of the person [Pandey et al. 2018]. Explainability is a desired quality in health prediction and recommendation systems as it can verify the quality of predictions and builds user engagement and trust in the system.

We can use such a model to provide the right guidance at the right time for health management. Health recommendation systems provide us a way to apply cybernetic principles to manage a person's health [Nag et al. 2017]. Using lifestyle interventions, we can build a navigation system that guides us through our day much in the same way that modern navigation systems inform drivers about the most optimum path towards their destination [Nag and Jain 2019]. We need to design a context-aware personal recommendation system that changes the person's context with every event that happens during the day. The dynamic context allows us to provide an optimal recommendation at every point of the day, and calibrate the recommendations as different events occur.

In this work, we present a sleep recommendation system that considers various lifestyle factors as contextual variables and possible recommendations for optimizing a sleep parameter. We create a rule-based model to understand the effects of different lifestyle factors (such as exercise during the day, and mealtimes) on sleep parameters. These rules are used in a recommendation system framework for providing the most effective interventions at any time throughout the day. These interventions could be lifestyle recommendations presented to the user or a command to one of the IoT or smart devices that control the user's environment (e.g., HVAC systems, Light bulbs, Music or ambient sounds).

2 RELATED WORKS

Our work spans across context-aware and health recommendation systems, sleep specific monitoring and prediction, and causal event mining. We review the current approaches and limitations below.

2.1 Context Aware and Health Recommendation Systems

The field of sleep and health recommendation systems is relatively new. Studies have explored the pitfalls of using the conventional recommendation systems for health and devised alternatives using entity properties and relationships [López-Nores et al. 2012]. Context-awareness is an essential quality for health recommendation systems [Schäfer et al. 2017]. Context-aware recommendation systems (CARS) have been explored in different domains [Villegas et al. 2018]. CARS have traditionally incorporated context information in collaborative filtering models in one of three ways, 1) Contextual Pre-filtering, 2) Contextual Post-filtering, and 3) Contextual Modelling [Adomavicius and Tuzhilin 2015]. There can be different types of contextual information relevant to a recommendation system. These usually fall into one of the following categories: temporal, location, individual (user characteristics), activity (about the activity), and relational (when multiple entities are involved)[Villegas et al. 2018]. We have adopted a contextual modeling approach and incorporated the contextual information in the rule-based model itself. Multiple studies have explored personalized recommendation systems for different aspects of user-health, such as diet [Khan et al. 2019]. These utilize different learning techniques to develop personalized models for individuals but usually lack explainability, which is an important characteristic of health recommendation systems.

2.2 Sleep Monitoring and Prediction Applications

Polysomnography is considered to be the gold standard for understanding sleep quality. The test records various metrics such as brain waves, oxygen levels in the blood, heart rate, breathing, and eye and leg movements [Kushida et al. 2005]. It requires a sleep expert and multiple medical sensors. The study's accuracy does come at the cost of needing too many resources and equipment to be performed every night reliably. Actigraphy is another popular technique used to measure sleep quality. It measures sleep quality using a wearable device (e.g., a watch) by recording movement during a sleep event. Its simplicity comes at the cost of accuracy, as it can only infer sleep quality via movement measurements. For our study, we used a combination of actigraphy and sound to record movements during sleep events. Several sleep applications use audio from the phone mic to record and report sleep quality (e.g., SleepCycle, Sleepscore).

Previous works have attempted to predict sleep quality using smartphone mic data in conjunction with machine learning [Min et al. 2014]. Other studies have attempted to use actigraphy graphs and utilize Deep Learning to predict sleep quality [Sathyanarayana et al. 2016]. There are even studies that attempt to forgo the idea of tracking sleep and use factor graph models based on daily activity to predict sleep quality with 78% accuracy [Bai et al. 2012]. While these models are useful, they do not address the individual variability in sleep and do not provide a way to incorporate different lifestyle aspects.

2.3 Effect of Lifestyle Activities on Sleep Quality

Current literature has made many attempts to identify daily activities that affect a person's sleep quality. Studies have shown that sleep and exercise are related, and higher physical activity levels can lead to better sleep latency [Yang et al. 2012] [Kline 2014][Kelley and Kelley 2017]. A systematic review has also shown that dietary patterns and the types of food eaten throughout the day lead to better sleep quality and duration [St-Onge et al. 2016]. The environmental factors (temperature and humidity) are also vital to our sleep duration and quality [Troynikov et al. 2018]. From these studies and many more, it is clear that choices made throughout the day affect the quality of the next night's sleep.

3 CAUSAL RULE-BASED MODELLING: EVENT MINING

Creating a model of the person's behavior and health is central to building personalized health recommendation systems. In this work, we present an approach to build a rule-based explainable model for predicting sleep outcomes in different contextual situations. We apply event mining [Jalali 2016] principles to perform N-of-1 experiments on a user's data [McDonald et al. 2017] that allows us to find causal relationships between different lifestyle events and biological outcomes. The process is described in figure 1.

Event mining allows us to discover and specify patterns between different events in a person's life. We utilize these patterns to create hypotheses that might describe a person's behavior. A hypothesis needs to specify the intervention event and the associated confounding factors that affect the relationship between the intervention and the outcome. The confounding factors are specified using the temporal delay operator, $\Delta[t_b, t_e]$, that relates the events that occur within the specified time interval $[t_b, t_e]$. The confounding factors are specified as a set of patterns P between the lifestyle events and the outcome. Thus, a hypothesis would be specified as $E_i \underset{P}{\longrightarrow} E_o$, where we want to measure the causal effect of intervention E_i on the outcome event E_o while controlling for the events specified by the set of patterns P. These patterns and hypotheses can be derived from existing knowledge and human intuition, allowing us also to leverage the results of population studies performed in clinical and

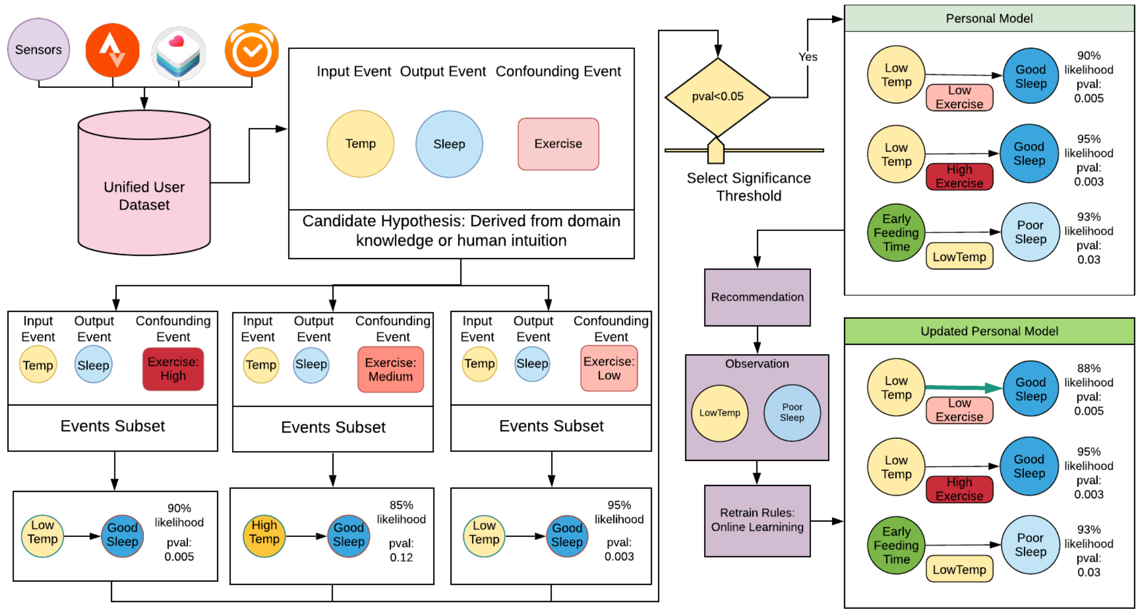

Figure 1: Rule based personal model for predicting sleep outcomes. We divide the occurrences of the outcome events into smaller subsets based on the values of co-occurring contextual factors. This minimizes the variance in the outcome due to the confounding variables within each subset. Subsets that exhibit significantly different distribution for different values of input events are converted to rules and added to the model.

controlled settings.
Combining the event mining operators with causal inference principles allows us to perform N-of-1 experiments on the user's longitudinal data. We utilize the candidate hypothesis specification to create different subsets of data based on the values of co-occurring confounding events. These subsets minimize the variance in the outcome due to confounding factors and mimic a version of the *do-operator* [Pearl 2009]. We can use different statistical techniques, such as linear regression or t-tests to find the effect of the intervention event on the outcome within each subset. This allows us to find the effect of the intervention on the outcome in an unbiased manner, and if we can capture all the confounding variables in the set of patterns, we would obtain the causal effect of the intervention on the outcome. The result of the test is stored as a conditional rule that uses confounding variables and the intervention event as the predicate. The distribution of the outcome events in the subset is used to make a prediction.

A set of these contextual rules can be used to predict health outcomes. We would need to find the most relevant rule by matching the user's current context with the set of rules and utilize it to make the prediction. We describe it in more detail in the next section. A rule-based model, while lacking in complexity, offers the advantage of explainability and online training. Every prediction and recommendation generated from this model can be explained using the associated contextual factors, thus eliminating recommendations based on spurious relationships. This is an essential characteristic of health models and recommendation systems. As the user behavior changes over time, the model needs to adapt to the changing user parameters and trained using the latest observed data. We can easily update the rules by updating the outcome variable's distribution whenever the rule matches the user's current context.

4 MULTI-ITEM HEALTH RECOMMENDATIONS

The rule-based health model allows us to find the health outcomes in different contextual situations. The user's activities during the day (such as exercise, meals, work-related stress) and their local environmental parameters (such as temperature, humidity, and ambient light and sounds) determine these contextual variables. Thus, we can utilize this model in a recommendation system setting to determine the set of parameters (both user behavior and environmental variables) for optimizing a health outcome (e.g., sleep quality).
Every action taken by the user and every environmental exposure changes the user's health state [Nag 2020], which changes the context for future actions and recommendations. We need to retrieve the relevant events from the user's events and data streams that impact their health state [Pandey et al. 2020]. Different contextual parameters are defined as aggregations of these events. For example, Total Screen Time during the day is an important confounding factor for understanding an individual's sleeping habits. It can be determined by aggregating the duration of all the screen activity events (such as working, watching TV, and social media activity) during the day. These aggregations can be performed using events

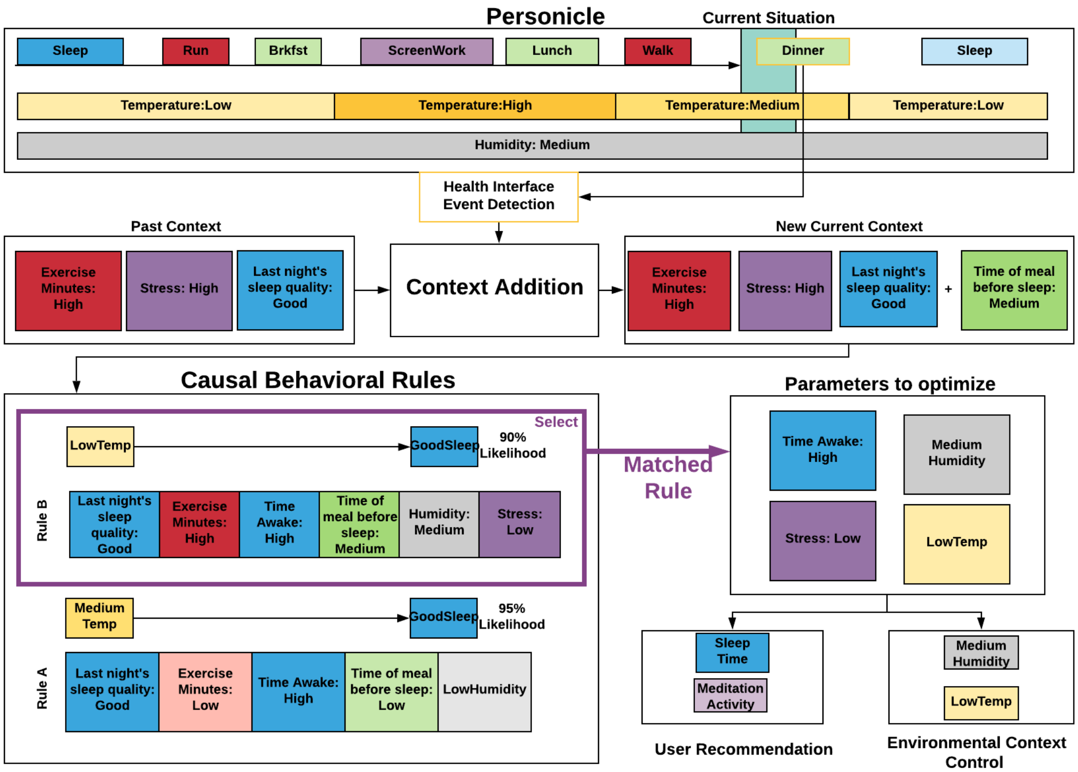

Figure 2: Live context calculation. The system updates user context every time they log an event. We retrieve all the sub-events and parameters relevant for context calculation (e.g., time of meal from dinner event). The retrieved contextual information is added to the existing context, and the updated context is used to generate a new set of recommendations. These recommendations are then sent either to the user or to a device controlling the user's environmental factors.

based triggers encoded as condition-action rules. As new events appear in the person's events log, the retrieved events can be aggregated to change the user's live context parameters. We can use the latest context values to provide a set of recommendations that would optimize the user's health outcomes.

We match the live contextual parameters for the person with the contextual parameters of the various rules present in the model. If the current context matches multiple rules, then we utilize the rule with the highest likelihood of the desired outcome. Once we have identified the rules that match the current context, we can utilize the unmatched contextual parameters and the intervention event to find the set of parameters that can lead to the optimal outcome. We can either present the recommendation to the user (if the recommendation is an action to be taken by the user) or communicate with a smart device that controls the user's environmental context (e.g., smart home devices, HVAC systems, smart bulbs). The recommendation system produces a set of actions that would maximize the likelihood of the optimal outcome; thus, the proposed recommendation system is different from typical recommendation systems as the recommendation consists of multiple items.

Since any event during the day can change the user's context, the recommendations are recomputed anytime an event changes the user's context. This process is depicted in figure 2. Thus, at any point during the day, the recommendation system would provide a list of timestamped actions to be performed by different agents (the user, or an automated device) to optimize the sleep outcomes.

5 EXPERIMENTS AND RESULTS

We ran experiments to create a personal rule-based model for optimizing a person's sleep quality metrics by providing lifestyle and local environmental recommendations. We utilized data collected by one individual for more than two years using readily available consumer applications and wearable and IoT devices. We performed two sets of experiments on the collected dataset to create and evaluate the model. The first set of experiments find the average causal effect of input variables on sleep quality metrics. We used Welch's t-tests and a p-value of 0.05 to determine statistical significance. The second set of experiments tested the prediction accuracy for a static pre-trained model vs. an online training model.

Table 1: Sleep Quality Measure and Event Thresholds

Variable	Classification Ranges	Event Name
Sleep Latency	[0, 15] (15, 30] (30, ∞)	Good Average Poor
Awake Minutes	[0, 20] (20, ∞)	Good Poor
Awakenings >5 mins	[0, 1] (1, ∞)	Good Poor
Sleep Efficiency	[0.85, 1.00] [0, 0.85)	Good Poor

5.1 Data Set

The data set includes exercise and lifestyle parameters for a 31 year old male collected continuously over 2 years via the user's Garmin Fenix 5 smart watch, their smartphone, and an IoT sensor that collected local temperature and humidity values. Sleep Cycle was primarily used to keep track of sleep events. Apple Health Kit was used to help compile sleep quality measures recorded by the Garmin smart watch, daily step counts, and daily floors climbed. The accelerometer measures of the smartwatch and the audio recordings of Sleep Cycle were used to create sleep quality measures. Strava was used to keep track of exercise events. An image based food log recorded feeding times with phone camera metadata, and a SensorPush IoT sensor was used to collect temperature and humidity during sleep events. All of these data sources were then temporally matched to record lifestyle events that took place throughout the day accurately.

We used the thresholds mentioned in Table 1 and Table 2 to convert the data streams to relevant events for the event mining analysis. We used nine lifestyle/environmental events: Previous Night's Sleep Quality Measures, Exercise Minutes in the Day, Interval Between Eating and Sleeping, Minutes Awake Between Sleep Events, Temperature, and Humidity when going to bed. The possible output events are sleep quality measures (Table 1). We used 70% of the data to build the model, and 30% of the data to test the model. The train-test split was created based on temporal order.

5.2 Causal Rules and Effects from N-of-1 Experiments

We perform different N-of-1 tests on the user's data to find the average effects of different lifestyle and environmental events on sleep quality parameters while controlling for other lifestyle parameters as confounding factors. We treat one of the input event's possible values as the baseline and compare the distribution of the outcome variable for other values of the event with the baseline distribution. If changing the input event value causes a significant change in the outcome distribution while controlling for confounding variables, the rule is deemed significant. This gives us the causal effect of different values of an input event on the observed outcome. We repeat this experiment while controlling for different sets of variables and aggregate the causal effects to find the input event's average causal effect. If the difference is not found to be significant, then we merge

Table 2: Lifestyle Factors and Event Thresholds

Variables	Classifications Ranges	Event Name
Exercise Minutes Per Day	[0] (0, 50] (50, 150] (150, ∞)	None Poor Average Good
Exercise Minutes Per Week	[0, 150] (150, 300] (300, ∞)	Poor Average Good
Interval Between Eating and Sleeping	[0] (0, 180] (180, ∞)	Missing Poor Good
Minutes Awake Between Sleep Events	[0, 900] (900, 1020] (1020, ∞)	LT 15 Hours Btwn 15-17 Hours GT 17 Hours
Starting Temperature	[0, 60] (60, 67] (67, ∞)	Cold Comfortable Warm
Starting Humidity	[0, 30] (30, 50] (50, 100]	Low Ideal High

the two distributions and use the combined distribution at the time of contextual matching.

The results of these experiments are in Figure 3. One interesting observation is that an average temperature(60-67 F^o) seems to improve every sleep quality measure except for sleep latency. This is an important observation as it shows that not all quality measures are correlated with each other and that an improvement in one does not necessarily equate to an improvement in all other sleep quality measures. Another interesting insight is that exercise improves sleep latency the most. On average, we can tell that exercising a lot will reduce sleep latency by 10.5 minutes with just a small workout will help reduce sleep latency by an average of 8 minutes.

5.3 Context Matching and Sleep Predictions

We also want to demonstrate the contextual matching of rules and test the accuracy of the model's predictions, as that will determine the efficacy of any recommendations we generate. We train a linear regression model corresponding to every rule in the model, and use the data subset that matches the rule to train the model. This model is then used to predict sleep outcomes for situations matching with the rule.

We used two training strategies for the prediction model; 1) Pre-trained static models, and 2) the warm start online training. We expect that over time the user's sleep behavior would change, and thus an online learning strategy would eventually start outperforming the pre-trained model.

The model's input features are Exercise minutes during the day, Feeding Time, Time Awake, Humidity, and Temperature while going to bed. We match the user's context with the context of the rules, and the most significant rule that matches the context is used to provide the recommendation.

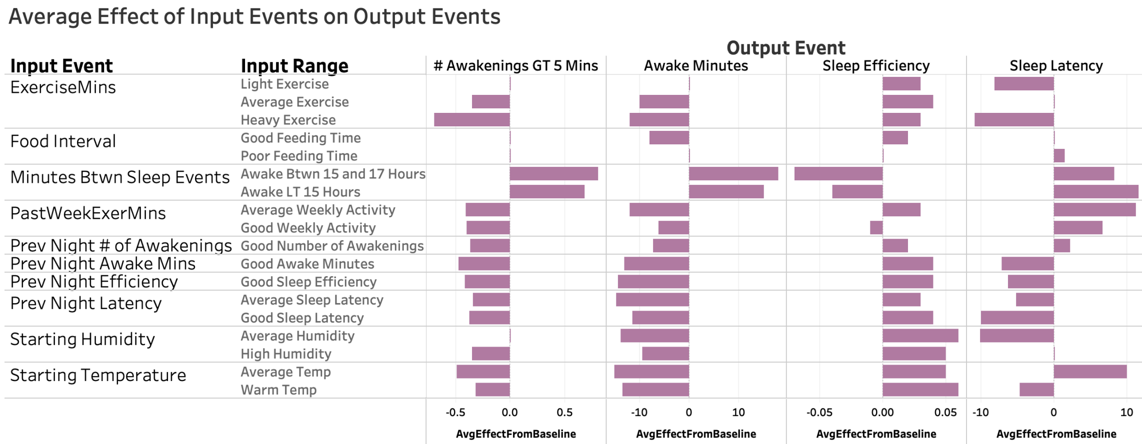

Figure 3: Average Effects that each input event has on the output event when compared to each input event's base category. If a metric is 0 then no significant relations were found.

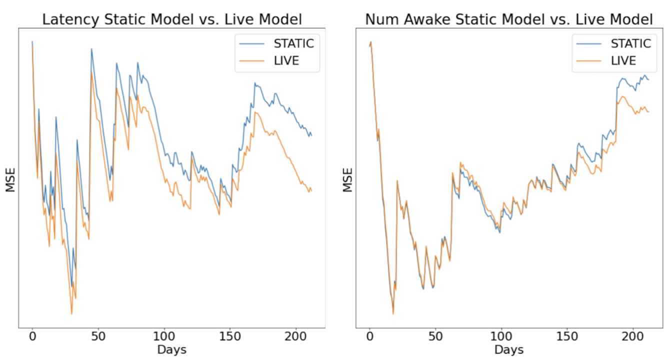

Figure 4: Comparison of pre-trained model vs. online training. Online training allows the model to adapt to user's changing sleep behavior resulting in lower error in predictions.

We create a set of contextual variables at the end of the day for each day in the dataset. These values are then used to find a matching rule. If multiple matches are found, then we used the rule with a higher statistical significance. The linear model associated with the matched rule would then be used to predict the sleep outcome parameter. The key difference between the pre-trained and online models is that the online model would be updated continuously using the data in the test set. This way, the online model has the opportunity to adapt to the user over time. The results of the model predictions are in 4. The results illustrate an improvement in the performance of the online model over the pre-trained model. Eventually, we expect the online model would achieve a much smaller MSE as it adapts to the changing sleep behavior exhibited by the user.

6 DISCUSSION AND FUTURE WORK

We have shown the need for and created a sleep model that utilizes event mining and causal inference principles to provide useful feedback about the relationships between lifestyle events and sleep quality. With enough data, this model can be potent. Coupled with the context-aware health recommendation system, it can give people control over their sleeping habits that have not been previously possible. The insights from the model are easily understandable and should promote user engagement as the recommendations are not coming from a black-box model but are simple relationships between daily habits. The context-aware recommendation approach allows us to provide recommendations at different points during the day. Even if the user fails to follow any recommendations, we can provide them with a new set of recommendations and modify their local environment to best suit their sleep requirements. This helps us move from recommendation-based guidance to navigational guidance.

Although this framework incorporates many useful data sources and provides useful insights into users' sleep behavior, there are many ways to improve. Many other lifestyle factors affect our sleep outcomes but are not included in our study, such as stress and nutritional intake. Our events based framework would allow us to include these events and data streams with minimal additional effort.

We have proposed a recommendation system to optimize one health outcome. However, in a real-world application, the users may want to optimize multiple outcomes simultaneously. This can be an exciting opportunity for the recommendation systems research community, and we hope to stimulate future work expanding to a larger user base and with different applications.

REFERENCES

Gediminas Adomavicius and Alexander Tuzhilin. 2015. Context-Aware Recommender Systems. In *Recommender Systems Handbook*. Springer US, Boston, MA, 191–226. https://doi.org/10.1007/978-1-4899-7637-6{_}6

Yin Bai, Bin Xu, Yuanchao Ma, Guodong Sun, and Yu Zhao. 2012. *Will You Have a Good Sleep Tonight? Sleep Quality Prediction with Mobile Phone.* Technical Report. https://www.andrew.cmu.edu/user/yuzhao1/papers/2012BodyNets1.pdf

Fran Casino, Constantinos Patsakis, Edgar Batista, Octavian Postolache, Antoni Martínez-Ballesté, and Agusti Solanas. 2018. Smart Healthcare in the IoT Era: A Context-Aware Recommendation Example. In *2018 International Symposium in Sensing and Instrumentation in IoT Era, ISSI 2018.* Institute of Electrical and Electronics Engineers Inc. https://doi.org/10.1109/ISSI.2018.8538106

Cathal Gurrin, Alan F. Smeaton, and Aiden R. Doherty. 2014. LifeLogging: Personal big data. , 125 pages. https://doi.org/10.1561/1500000033

Andrew P. Hills, Steven J. Street, and Nuala M. Byrne. 2015. Physical Activity and Health: "What is Old is New Again". In *Advances in Food and Nutrition Research*. Vol. 75. Academic Press Inc., 77–95. https://doi.org/10.1016/bs.afnr.2015.06.001

Laleh Jalali. 2016. Interactive Event-driven Knowledge Discovery from Data Streams. (2016).

George A. Kelley and Kristi Sharpe Kelley. 2017. Exercise and sleep: a systematic review of previous meta-analyses. *Journal of Evidence-Based Medicine* 10, 1 (2 2017), 26–36. https://doi.org/10.1111/jebm.12236

Mansura A Khan, Ellen Rushe, Barry Smyth, and David Coyle. 2019. *Personalized, Health-Aware Recipe Recommendation: An Ensemble Topic Modeling Based Approach.* Technical Report. https://github.com/MAK273/SupportingFileForHealthRecsys2019

Christopher E Kline. 2014. The bidirectional relationship between exercise and sleep: Implications for exercise adherence and sleep improvement. *American journal of lifestyle medicine* 8, 6 (11 2014), 375–379. https://doi.org/10.1177/1559827614544437

Clete A Kushida, Michael R Littner, Timothy Morgenthaler, Cathy A Alessi, ; Dennis Bailey, ; Jack Coleman, Leah Friedman, Max Hirshkowitz, Sheldon Kapen, Milton Kramer, Teofilo Lee-Chiong, ; Daniel, L Loube, Judith Owens, ; Jeffrey, P Pancer, and ; Merrill Wise. 2005. *Practice Parameters for the Indications for PSG-AASM Practice Parameters SLEEP Practice Parameters for the Indications for Polysomnography and Related Procedures: An Update for 2005.* Technical Report 4. https://academic.oup.com/sleep/article-abstract/28/4/499/2696969

Duck chul Lee, Angelique G. Brellenthin, Paul D. Thompson, Xuemei Sui, I. Min Lee, and Carl J. Lavie. 2017. Running as a Key Lifestyle Medicine for Longevity. , 45–55 pages. https://doi.org/10.1016/j.pcad.2017.03.005

Martín López-Nores, Yolanda Blanco-Fernández, José J. Pazos-Arias, and Alberto Gil-Solla. 2012. Property-based collaborative filtering for health-aware recommender systems. *Expert Systems with Applications* 39, 8 (6 2012), 7451–7457. https://doi.org/10.1016/j.eswa.2012.01.112

Suzanne McDonald, Francis Quinn, Rute Vieira, Nicola O'Brien, Martin White, Derek W. Johnston, and Falko F. Sniehotta. 2017. The state of the art and future opportunities for using longitudinal n-of-1 methods in health behaviour research: a systematic literature overview. , 307–323 pages. https://doi.org/10.1080/17437199.2017.1316672

Jun-Ki Min, Afsaneh Doryab, Jason Wiese, Shahriyar Amini, John Zimmerman, and Jason I. Hong. 2014. Toss 'n' turn. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*. ACM Press, New York, New York, USA, 477–486. https://doi.org/10.1145/2556288.2557220

Nitish Nag. 2020. *Health State Estimation.* Ph.D. Dissertation. University of California, Irvine. http://arxiv.org/abs/2003.09312

Nitish Nag and Ramesh Jain. 2019. A Navigational Approach to Health: Actionable Guidance for Improved Quality of Life. *Computer* 52, 4 (4 2019), 12–20. https://doi.org/10.1109/MC.2018.2883280

Nitish Nag, Vaibhav Pandey, Hyungik Oh, and Ramesh Jain. 2017. Cybernetic Health. (5 2017). http://arxiv.org/abs/1705.08514

Nitish Nag, Vaibhav Pandey, Preston J. Putzel, Hari Bhimaraju, Srikanth Krishnan, and Ramesh Jain. 2018. Cross-Modal Health State Estimation. In *MM 2018 - Proceedings of the 2018 ACM Multimedia Conference*. Association for Computing Machinery, Inc, New York, New York, USA, 1993–2002. https://doi.org/10.1145/3240508.3241913

Hyungik Oh and Ramesh Jain. 2017. From multimedia logs to personal chronicles. In *MM 2017 - Proceedings of the 2017 ACM Multimedia Conference*. Association for Computing Machinery, Inc, New York, New York, USA, 881–889. https://doi.org/10.1145/3123266.3123375

Vaibhav Pandey, Nitish Nag, and Ramesh Jain. 2018. Ubiquitous event mining to enhance personal health. In *UbiComp/ISWC 2018 - Adjunct Proceedings of the 2018 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2018 ACM International Symposium on Wearable Computers*. ACM Press, New York, New York, USA, 676–679. https://doi.org/10.1145/3267305.3267684

Vaibhav Pandey, Nitish Nag, and Ramesh Jain. 2020. Continuous Health Interface Event Retrieval. In *Proceedings of the 2020 International Conference on Multimedia Retrieval*. Association for Computing Machinery (ACM), New York, NY, USA, 486–494. https://doi.org/10.1145/3372278.3390705

Judea Pearl. 2009. Causal inference in statistics: An overview. *Statistics Surveys* 3, 0 (2009), 96–146. https://doi.org/10.1214/09-SS057

Jerome Sarris, Adrienne O'Neil, Carolyn E. Coulson, Isaac Schweitzer, and Michael Berk. 2014. Lifestyle medicine for depression. , 13 pages. https://doi.org/10.1186/1471-244X-14-107

Aarti Sathyanarayana, Shafiq Joty, Luis Fernandez-Luque, Ferda Ofli, Jaideep Srivastava, Ahmed Elmagarmid, Teresa Arora, and Shahrad Taheri. 2016. Sleep Quality Prediction From Wearable Data Using Deep Learning. *JMIR mHealth and uHealth* 4, 4 (11 2016), e125. https://doi.org/10.2196/mhealth.6562

Hanna Schäfer, Santiago Hors-Fraile, Raghav Pavan Karumur, André Calero Valdez, Alan Said, Helma Torkamaan, Tom Ulmer, and Christoph Trattner. 2017. Towards Health (Aware) Recommender Systems. In *Proceedings of the 2017 International Conference on Digital Health - DH '17*. ACM Press, New York, New York, USA, 157–161. https://doi.org/10.1145/3079452.3079499

Abigail J. Sellen and Steve Whittaker. 2010. Beyond total capture: A constructive critique of lifelogging. *Commun. ACM* 53, 5 (5 2010), 70–77. https://doi.org/10.1145/1735223.1735243

Tamar Shochat. 2012. Impact of lifestyle and technology developments on sleep. , 19–31 pages. https://doi.org/10.2147/NSS.S18891

Marie-Pierre St-Onge, Anja Mikic, and Cara E Pietrolungo. 2016. Effects of Diet on Sleep Quality. *Advances in Nutrition* 7, 5 (9 2016), 938–949. https://doi.org/10.3945/an.116.012336

Olga Troynikov, Christopher G. Watson, and Nazia Nawaz. 2018. Sleep environments and sleep physiology: A review. , 192–203 pages. https://doi.org/10.1016/j.jtherbio.2018.09.012

Norha M. Villegas, Cristian Sánchez, Javier Díaz-Cely, and Gabriel Tamura. 2018. Characterizing context-aware recommender systems: A systematic literature review. *Knowledge-Based Systems* 140 (1 2018), 173–200. https://doi.org/10.1016/J.KNOSYS.2017.11.003

Pei Yu Yang, Ka Hou Ho, Hsi Chung Chen, and Meng Yueh Chien. 2012. Exercise training improves sleep quality in middle-aged and older adults with sleep problems: A systematic review. *Journal of Physiotherapy* 58, 3 (9 2012), 157–163. https://doi.org/10.1016/S1836-9553(12)70106-6