DISCOVERING OPTIMAL VARIABLE-LENGTH TIME SERIES MOTIFS IN LARGE-SCALE WEARABLE RECORDINGS OF HUMAN BIO-BEHAVIORAL SIGNALS

Tiantian Feng, Shrikanth S. Narayanan

Signal Analysis and Interpretation Lab, University of Southern California, Los Angeles, CA Email: tiantiaf@usc.edu, shri@ee.usc.edu

ABSTRACT

Continuously-worn wearable sensors produce copious amounts of rich bio-behavioral time series recordings. Exploring recurring patterns, often known as motifs, in wearable time series offers critical insights into understanding the nature of human behavior. Challenges in discovering motifs from wearable recordings include noise removal, pattern generalization, and accounting for subtle variations between subsequences in one motif set. In this work, we introduce a time series processing pipeline to summarize an optimal set of variable-length motifs in a real-world wearable recording data-set collected in a hospital workplace setting. We propose the use of the Savitzky-Golay filter for noise removal without significant data distortion. We then combine the previously developed HierarchIcal based Motif Enumeration (HIME) algorithm with a principled optimization approach to obtain the most repetitive patterns in longterm wearable time-series. We also describe challenges in using just a single method to detect motifs in wearable time series in our experiments. We demonstrate our pipeline can effectively identify meaningful variable-length motifs in large-scale heart rate signals collected continuously from over 100 individuals both at and outside their workplace over 10 weeks through two machine learning experiments.

Index Terms— Human behavior patterns, wearable, motif analysis, principle optimization, SAX

1. INTRODUCTION

Driven by the recent advances in microelectronics technology, wearable sensors today can capture vital physiological and bio-behavioral about an individual over a prolonged period in real-world settings. Such rich multimodal time series data offer opportunities for enhanced understanding of individual behaviors including social interactions [1], sleep patterns [2] and emotion variations [3]. However, the underlying patterns of these time series data (such as heart rate, step counts of physical activity over time) in the wild are often not known to the researcher, and a visual inspection is time-consuming and costly, if not manually impossible, and frequently results in a limited understanding of the data.

Motifs are most repetitive similar patterns that frequently appear in time series. Motif discovery on wearable recordings can help identify underlying phenomena and provide useful insight to understand biobehavioral patterns and mechanisms. Since the introduction of the motif discovery problem, many approaches have been introduced [4, 5, 6, 7]. Conversion of the signal into a symbolic representation (named SAX) has been extensively exploited in search-based algorithms. [7] formulate motif discovery as an optimization problem by learning values of motifs which maximize the frequency of

recurring subsequences. However, existing SAX-based motif discovery approach could return many different motifs with close distances which may harm the interpretation of the data. Meanwhile, direct optimization for motif discovery enforces the motif window size to be the same, while the lengths of similar subsequences in real-life biobehavioral recordings could vary and are undetermined.

In this paper, we present a novel processing pipeline for discovering the most 'meaningful' variable-length motifs in large-scale biobehavioral time series from wearable sensors. Given a minimum motif length m_{min} , the pipeline automatically determines a set of motifs longer than m_{min} and the Euclidean distance between every other motif is greater than a threshold distance. The pipeline includes filtering, HierarchIcal based Motif Enumeration [6], and a principled optimization for learning most meaningful motifs. In this work, we particular define motifs which maximize the number of recurrent subsequences in time series as **meaningful**, while pairwise distances between different motifs are above a threshold value. The main contributions of this work are as follows:

1. We propose a novel combination of SAX-based motif matching algorithm and principle optimization for learning meaningful motifs in long-term real-world biobehavioral recordings. We regard the optimization process as a summerization task.

2. We show, on a wearable sensor timeseries data-set of PPG heart rate data we recently collected from over 100 clinical staff in a large hospital (both at and outside work) over a 10 week period, that our pipeline captures useful structures that could be used for human behavior analysis and modeling.

2. BIOBEHAVIORAL TIMESERIES DATA SET

In early 2018, we conducted a set of comprehensive human-subject experiments of the "TILES: Tracking Individual Performance with Sensors" study to examine the physiological, environmental, and behavioral variables affecting job performance and employee wellness. Throughout a ten week period, we collected data through multimodal wearable sensors from over 100 hospital nursing volunteers working at a large critical care hospital located in Los Angeles, CA. From this sample of individuals, the sample for the present paper included 138 individuals working as full-time nursing professions at the hospital. The current sample included 81 females (rest males) and 84 individuals worked the day shift (rest, worked night shift). In particular, we asked participants to wear the Fitbit Charge 2 at all times throughout the 10-week data collection period, both at and outside work. Fitbit Charge 2 [8] is a wristband device that offers information about physical activities, sleep quality, and heart rate through photoplethysmography (PPG). In this paper, we primarily focus on exploring time series motifs from the PPG heart rate (HR). Table summarizes the recording statistics of HR data in our dataset.

Table 1. Average/Std of valid HR recording hours in TILES study.

	All Nurses	Day-shift	Night-shift
Average	1162.63	1182.50	1131.30
Standard Deviation	403.48	401.51	404.60

3. NOTATION

In this section, we introduce some notations and definitions used in this paper.

Definition 3.1. The time series data T of length n is represented as a set of observations $(t_i, t_{i+1}, ..., t_n)$ ordered by time, with $t_i \in \mathbb{R}$. In our paper, t_i represents the PPG HR at time point i. For motif discovery, all time-series are normalized to remove scaling and offset effects, given the primary interest is to find similar matching 'shape'. Typical normalization procedure used in motif discovery is z-normalization.

Definition 3.2. Given a time series T of length n, a time series subsequence $S_{p,q}$ is a contiguous set of samples starting from point p and ending at q with length m = q - p + 1, s.t $1 \le p \le n - m + 1$.

Definition 3.3. A word $w = w_1 w_2 \dots w_L$ is the symbolic representation of a sub-sequence $S_{p,q}$, also notated as $R(S_{p,q})$, with $w_i \in \Sigma$, where Σ is the representation alphabet set with size of a. For example, when a = 3, $\Sigma = \{a, b, c\}$. Sub-sequence S_{q_1,p_1} and S_{q_2,p_2} match if their symbolic representations are the same, namely, $R(S_{q_1,p_1}) = R(S_{q_2,p_2})$

Definition 3.4. Given a subsequence $S_{p,q}$ and its symbolic representation $R(S_{p,q})$, $S_{p,q}$ is an instance of a word w if $R(S_{p,q}) = w$.

Definition 3.5. In this work, we particular define motifs which maximize the number of recurrent subsequences in time series as **meaningful**, while pairwise Euclidean distances between different motifs are above a threshold value. This definition is similar to [7].

4. METHOD

In this section, we introduce our pipeline in generating a set of meaningful varied-length time series motifs from wearable recordings. We first aggregate HR readings every three minutes. The second step uses the Savitzky-Golay filter to smooth the wearable time series data without substantially distorting the signal. HierarchIcal based Motif Enumeration (HIME) is then used to extract a motif set in symbolic representation. Finally, to find the most meaningful motifs from the previous step, we introduce an optimization problem similar to [7], to learn a motif sets with size K. Our method is summarized as a sequence of steps: 1) Data aggregation, 2) Savitzky-Golay filtering, 3) HIME-SAX based algorithm that returns a set of motifs and 4) Top-k motifs learning. These are further described below.

4.1. Data Aggregation

The PPG heart rate samples are made available by the Fitbit Charge 2 sensors at intervals less than one minute, but the time differences between two consecutive samples are irregular. Discovering motifs in time series that are not sampled at fixed rate is meaning-less, and consequently we aggregate the measured heart rate samples to create a time series T_{Agg} with a set of heart rate observations

 $((t_{Agg})_1, ...(t_{Agg})_i, (t_{Agg})_{i+1}, ..., (t_{Agg})_n)$, such that time difference between $(t_{Agg})_i$ and $(t_{Agg})_{i+1}$ is constant. Another reason for the aggregation of heart rate data is that prior literature suggests that PPG heart rate should be averaged at over 1-minute duration to obtain a reliable measure [9]. A large time difference between $(t_{Agg})_i$ and $(t_{Agg})_{i+1}$ could lead to significant loss of resolution, as a result we empirically chose the interval as 3 minutes. We want to highlight that the aggregation step may not necessary for every kind of wearable time series, such as step count.

4.2. Savitzky-Golay filter for signal enhancement

Wearable sensors can easily pick up noises from the environment or body movement and thus filtering becomes an essential prerequisite for further processing of signal. Savitzky-Golay filter [10], often known as S-G filter, has been successfully used to reduce noises in signals while simultaneously maintaining the shape and height of peaks. In our context, we adopted the S-G filter to smooth the aggregated HR since our interest is in identifying time series subsequences where their 'shapes' match. The idea of S-G filter is to perform a least-squares polynomial approximation of a subset of consecutive data points to a polynomial and then a convolution of all the polynomials is obtained. Namely, given the time series T_{Agg} we obtained from 4.1, S_{Agg} is a subsequence of length 2m + 1, and S-G filter tries to fit a polynomial with degree $z (p_i = \sum_{x=0}^{z} a_x i^x)$ to the S_{Agg} centered at i = 0 which minimize the least-squared error below:

$$\varepsilon = \sum_{i=-m}^{m} (p_i - (t_{Agg})_i)^2 \tag{1}$$

Prior study [11] has shown that large window sizes and low polynomial degrees may yield distorted signals. To minimize the deforming of the signal, we experimentally choose a small window size of 5, and cubic order polynomials in the S-G filter. These parameters could be tuned systematically based on criteria and heuristics defined in [11, 12], but we leave this endeavor for future work.

4.3. HierarchIcal based Motif Enumeration (HIME)

HierarchIcal based Motif Enumeration (HIME) [6] is a greedy algorithm to detect variable length-motifs. Similar to most existing SAX motif discovery algorithms, HIME first converts a z-normalized subsequence $S_{p,q}$ with length of m_{start} in time series T (obtained from 4.2), to Piecewise Aggregate Approximation (PAA) [13] segments of size L. We empirically select $m_{start} = 24$ in this pipeline, since HR changes with patterns in range of 30 -90 minutes in many meaningful contexts , like after eating [14], during sleep [15], and exercise. HIME later transverses the similar sequences in neighbor points and update length m of subsequences in the algorithm, and we set the minimum length of the return instance as $m_{min} = 10$ (equivalent to 30 minutes). Secondly, each PAA segment is mapped to a symbolic representation $w_i \in \Sigma$, where the size of alphabet set (Σ) is a. The region of each alphabet character is approximately equal-probable under Gaussian distribution [5].

Numerosity Reduction(NR) [16] is usually the procedure followed in most existing works to remove similar neighboring subsequences off by one time point to avoid unnecessarily long motifs, but it can miss some patterns due to skipped subsequences. Instead, HIME applies induction graphs to represent each subsequence off by one point as nodes, where consecutive subsequences in time series ($S_{p,q}$ and $S_{p+1,q+1}$) are connected by an edge. The induction graph stores the next non-similar subsequence of a current subsequence by attaching a forward edge in between. HIME performs a left to the right passing of all nodes in Induction Graph, and recursively enumerate variable-length subsequences and to detect repeated SAX symbolic representation. After visiting one node, HIME uses a greedy approach to remove short motifs which are fully covered by longer motifs to maintain a small motif set for decreasing computational costs. A detailed description of HIME is in [6].

We employ HIME [6] in our pipeline, as author highlighted that it detects motifs using a single SAX word of length L representing the subsequence $S_{p,q}$ of length m where m could vary. It essentially implies that instances with the same SAX representation have close Eucliedien distances in between by measuring their PAA segments of size L. It is critical to form the optimization problem in the next subsection since input matrix should have fixed number of columns. Besides, as a pre-processing step prior to optimally select top-K motifs, HIME filters out subsequences which rarely occur in time series, and returns a instance set of size J where J is much smaller than length (n) of T. Thus HIME also significantly reduces computations in the optimization problem we formed as a later step.

4.4. Top-K motifs Learning

In the final step, we are trying to learn the most 'meaningful' (see 3.5) motifs from instances extracted by HIME, which can be recognized as a summerization task. We have found that HIME could return over 500 frequently occurring motifs (formed by single SAX word) from the HR time series that we collected over 10 weeks for each participant. One straightforward solution is to choose the K most frequently occurring subsequences, but we could easily select two motifs where their symbolic representation are similar. For example, we could find two motifs 'EEEDDEEE' and 'EEDDDEEE' which occur most frequently in an individual's HR data, but selecting both of them may be erroneous if subtle differences between these similar motifs may likely be caused by sensor noise. On the other hand, a simple clustering solution, like K-means, might regard two HIME-motifs 'EEEDDEEE' and 'EEECCEEE' as one cluster, which assumes they might belong to two different motifs.

[7] proposed to learn $M_{do} \in \mathbb{R}^{K \times m}$ that contains K motifs from $U_{do} \in \mathbb{R}^{(n-m+1) \times m}$, where n is the length of time series T, and the *i*-th row in U_{do} represents z-normalized subsequence $S_{i,i+m-1}$ ('do' stands for direct optimization). However, we argue that the length of possible recurrent patterns in wearable sensor recordings are always unknown, so choosing m is difficult if we directly deploy the method. HIME identifies a set of subsequences with varied length, which are ideal input for this optimization task. In particular, we concatenate PAA segments of each subsequence with length of L to create U ($U \in \mathbb{R}^{J \times L}$). The reason we could do this is that each subsequence is compared as a single SAX word representation of length L in HIME. We define the total frequency given motif matrix M with K motifs, and input matrix U with Jinstances as F(M), and we obtain:

$$F(M) = \sum_{k=1}^{K} \sum_{j=1}^{J} F_{k,j}$$
(2)

$$F_{k,j} = \begin{cases} 1, & \text{if } \sum_{l=1}^{L} (M_{k,l} - U_{j,l})^2 < D\\ 0, & \text{otherwise} \end{cases}$$
(3)

D is the distance threshold, and we determine D by calculating pairwise Euclidean distances in U and choose a small percentile as D. A candidate M should contain K motifs that maximize the function F(M), while K motifs should be different from each other by a threshold distance. Intuitively, we choose 2D to be the minimum distances between each learned motif in M. Finally, optimization task can be formulated as:

$$M^* = \underset{M \in \mathbb{R}^{K \times L}}{\operatorname{argmax}} F(M)$$
(4)
subject to
$$\sum_{l=1}^{L} (M_{k,l} - M_{h,l})^2 > 2D$$
where $k \in \{1, ..., K\}$, and $h \in \{k + 1, ..., K\}$

As [7] suggested, we approximate F(M) using a Gaussian kernel to handle zero derivative of F(M) and discontinuity at the point where $\sum_{l=1}^{L} (M_{k,l} - U_{j,l})^2 = 0$:

$$\hat{F}(M) = \frac{1}{KJ} \sum_{k=1}^{K} \sum_{j=1}^{J} \hat{F}_{k,j}$$
(5)

where:
$$\hat{F}_{k,j} = \exp^{\frac{\alpha}{T} \sum_{l=1}^{L} (M_{k,l} - U_{j,l})^2}$$
 (6)

The parameter α controls the smoothness of approximation function, and we keep $\hat{F}(M)$ in the range of 0 - 1 for optimization reasons. Gradient ascent solution for this optimization is then derived using this approximation. Details of the optimization algorithm can be found in [7].

5. RESULTS

We present here the motif discovery results using our pipeline on the TILES PPG recordings. As described in 4, we first aggregate the PPG HR over a period of 3 minutes. We then filter the aggregated time series T_{Agg} using an S-G filter with cubit fitting and a windows size of 5. T are then processed using HIME, and a set of subsequences of size J are returned. We chose L = 8 (Length of SAX word) and instances of minimum length (m_{min}) as 10 and starting length (m_{start}) of 24 in HIME. We automatically determine the lower bound of a by the proposed method in [6]. We concatenated PAA segments of each instance to create U where we keep the number of PAA segments as L. Finally, we set to learn K (K = 5)motifs which maximize the frequency function defined in 4.4 from U. We select $\alpha = 2$ for the Gaussian approximation in equation 6, and we choose distance threshold at percentile of 1%, 2%, 3%. We apply the pipeline to extract motifs for each participant. Figure 1 shows 2 motif examples for a participant which repetitively occurred during sleep. The heart rate motifs shown in 1(a) could potentially indicate a sleep cycle pattern according to [15], and motifs in 1(b) offer hints to identify wake-up event. We further conducted two experiments below to validate motifs we extracted.

5.1. Work status prediction

In the first experiment, we propose to use motifs to identify contextual information of a participant, particularly working status on a daily basis. We extracted the frequency of each motif in *M* between 7am-7pm and 7pm-7am on a daily basis for day-shift and night-shift nurses, respectively, to predict working status. To compare with motifs we learned, we picked up the five most frequently occurred SAXmotif from HIME, and extracted the same frequency feature. We train the SVM model for each participant and report the averaged 5-fold validation accuracy results for all participants. Table 2 shows the average 5-fold validation accuracy results of work-status prediction, and we observe that motifs extracted from our pipeline offers better accuracy than HIME motifs.



(a) 4 subsequences of a motif that occurred repeatedly during the start of sleep



(b) 4 subsequences of a motif that occurred repeatedly during the end of sleep

Fig. 1. Plots show two HR motifs which appear repeatedly at different sleep periods. We note that there are still slight dissimilarities in shape between same set of instances, which may be caused by sensor misplacement, body movement and environment noise, but we observe that these instances match well with each other qualitatively.

Table 2. Prediction accuracy of work status using motif frequencies from our pipeline and HIME algorithm

	HIME	Motif, D	Motif, D	Motif, D
	Motif	at pct=1%	at pct=2%	at pct=3%
Accuracy	65.20%	65.62%	67.20%	68.54 %

5.2. Sleep quality prediction

In this experiment, we propose to use motif features to classify sleep quality. The Pittsburgh Sleep Quality Index (PSQI) was administered during enrollment sessions and was used to assess sleep quality [17]. The total PSQI score average is 7.00 (SD=1.76) for day-shift nurses and 8.59 (SD=2.87) for night-shift nurses, on a scale ranging from 0 (best quality) to 21 (worst quality). We further binaries the score with the threshold of 7 according to [17] for each participant, where a score above 7 indicates potential sleep disorders.

For motifs we extracted from the pipeline, we pick up the motif that appears most frequently during the sleep, namely $M_{i,:}$. We then extracted the standard deviation and median from $M_{i,:}$ ($M_{i,:} \in \mathbb{R}^{1 \times L}$) and its delta differences. Since length m of subsequences for the same motif could vary, we further extract the standard deviation/median of motif duration occurred during sleep through the 10-week study. Meanwhile, we extract the standard deviation and median of sleep duration, sleep efficiency and sleep REM duration from Fitbit sleep summaries to predict sleep quality for comparison. We discard participants with less than 30 nights of sleep data in this experiment. Similar to the first experiment, we train the SVM model and report the 5-fold validation accuracy results in Table 4. We observe that motif features consistently yield higher prediction accuracy than Fitbit summaries, and returned significantly higher validation accuracy when D is at pct = 2%.

6. DISCUSSION AND FUTURE WORK

Existing methods for motif discovery do not completely address the challenges faced when processing biobehavioral time series obtained from wearable sensors in real world condition i.e., the ability to

Table 3. A list of features from which they were derived in Fitbit sleep summary and most frequently occurred motif during sleep.

Stream	Statistics	Data
		Sleep duration
Fitbit Summary	Median/Std	Sleep efficiency
		REM Sleep duration
Learned Motif		PAA segments
	Median/Std	PAA segments delta difference
		Motif duration

Table 4. Prediction accuracy of binaries sleep quality score using motif feature and Fitbit sleep summary

Feature	Accuracy	F-1
Fitbit Sleep summary	60.86%	61.90%
Motif statistics, D at pct=1%	67.30%	66.35%
Motif statistics, D at pct=2%	75.36 %	77.47 %
Motif statistics, D at pct=3%	66.67%	67.51%

detect variable-length motif while returning the most 'meaningful' motifs (See definition 3.5). In our pipeline, we showed that a combination of HIME and principled optimization could help systematically discover useful variable-length motifs in long terms wearable sensor recordings, such as HR considered in our experiments. We demonstrate foremost that motif features returned from our pipeline yield higher accuracy in predicting contextual conditions like work status than HIME-motifs, and also outperforms Fitbit sleep summary in classifying sleep quality.

We found that there are only few published work focused on detecting motifs in timeseries obtained from commercial wearable sensors [18, 19]. To the best of our knowledge, our work to apply motif discovery on long-term biobehavioral time series from wearables is among the early efforts. Despite potential accuracy issues in heart rate measurement from Fitbit Charge 2 [20], it minimally interferes with our analysis since our scope is to study the variations in the signal where the DC component is removed by z-normalization.

There are several compelling research directions for expanding on this work which we aim to explore in the future. We plan to learn time series motifs using unsupervised learning approaches such as with auto-encoders. Besides, we will also examine the dictionary learning method in exploring motifs in our data-set. We intend to make a comparison on results generated from these approaches and summarize a more generalized solution especially from the viewpoint of the utility of the discovered motifs in predictive modeling such as mental state and work performance [21].

7. ACKNOWLEDGEMENT

The research is based upon work supported by the Office of the Director of National Intelligence(ODNI), Intelligence Advanced Research Projects Activity(IARPA), via IARPA Contract No 2017 - 17042800005. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing official policies or endorsements, either expressed or implied, of ODNI, IARPA, or U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

8. REFERENCES

- Y. Mohammad and T. Nishida, "Using physiological signals to detect natural interactive behavior," *Applied Intelligence*, vol. 33, no. 1, pp. 79–92, Aug. 2010.
- [2] A. Roebuck, V. Monasterio, E. Gederi, M. Osipov, J. Behar, A. Malhotra, T. Penzel, and G. D. Clifford, "A review of signals used in sleep analysis," *Physiological Measurement*, vol. 35, no. 1, pp. R1, 2014.
- [3] K. L. Hui, and R. S. Sherratt, "Coverage of emotion recognition for common wearable biosensors," *Biosensors*, vol. 8, no. 2, 2018.
- [4] T. Rakthanmanon, B. Campana, A. Mueen, G. Batista, B. Westover, Q. Zhu, J. Zakaria, and E. Keogh, "Searching and mining trillions of time series subsequences under dynamic time warping," in *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'12)*, 2012, pp. 262–270.
- [5] J. Lin, E. Keogh, L. Wei, and S. Lonardi, "Experiencing SAX: A novel symbolic representation of time series," *Data Mining* and Knowledge Discovery, vol. 15, no. 2, pp. 107–144, Oct. 2007.
- [6] Y. Gao, and J. Lin, "Efficient discovery of time series motifs with large length range in million scale time series," *International Conference on Data Mining (ICDM' 17)*, pp. 1213– 1222, 2017.
- [7] J. Grabocka, N. Schilling, and L. Schmidt-Thieme, "Latent time-series motifs," *ACM Transactions on Knowledge Discovery from Data*, vol. 11, no. 1, pp. 6:1–6:20, July 2016.
- [8] "Fitbit charge 2," 2018.
- J. Allen, "Photoplethysmography and its application in clinical physiological measurement," *Physiological Measurement*, vol. 28, Jan. 2007.
- [10] A. Savitzky, and M. Golay, "Smoothing and differentiation of data by simplified least squares procedures," *Analytical chemistry*, vol. 36, pp. 1627–1639, July 1964.
- [11] G. Viv-Truyols, and P. Schoenmakers, "Automatic selection of optimal savitzkygolay smoothing," *Analytical chemistry*, vol. 78, pp. 4598–608, Aug. 2006.
- [12] M. Sadeghi, and F. Behnia, "Optimum window length of savitzky-golay filters with arbitrary order," *arXiv Preprint*, Aug. 2018.
- [13] J. Lin, E. Keogh, S. Lonardi, and P. Patel, "Finding motifs in time series," in *Proceedings of the Second Workshop on Temporal Data Mining*, 2002, pp. 53–68.
- [14] B. waaler, M. Eriksen, and K. Toska, "The effect of meal size on postprandial increase in cardiac output," *Acta Physiologica*, vol. 142, no. 1, pp. 33–39.
- [15] L. Joanne, and J. Ashley, "Variations of heart rate during sleep as a function of the sleep cycle," *Electroencephalography and Clinical Neurophysiology*, vol. 35, no. 2, pp. 193–198, 1973.
- [16] Y. Li and J. Lin, and T. Oates, "Visualizing variable-length time series motifs," in *Proceedings of the 12th SIAM International Conference on Data Mining*, Apr. 2012, pp. 895–906.
- [17] D. Buysse, C. Reynolds, T. Monk, S. Berman, and D. Kupfer, "The pittsburgh sleep quality index: A new instrument for psychiatric practice and research," *Psychiatry Research*, vol. 28, no. 2, pp. 193–213, 1989.

- [18] O. Yürüten, J. Zhang, and P. Pu, "Decomposing activities of daily living to discover routine clusters," in *Proceedings of the Twenty-Eighth Conference on Artificial Intelligence (AAAI*" 14), 2014, pp. 1348–1354.
- [19] I. Fox, L. Ang, M. Jaiswal and R. Pop-Busui, and J. Wiens, "Contextual motifs: Increasing the utility of motifs using contextual data," in *KDD*, 2017.
- [20] S. Benedetto, C. Caldato, E. Bazzan, D. Greenwood, V. Pensabene, and P. Actis, "Assessment of the fitbit charge 2 for monitoring heart rate," *PLOS ONE*, vol. 13, no. 2, pp. 1–10, Feb. 2018.
- [21] H. Kao, H. Hosseinmardi, S. Yan, M. Hasan, S. Narayanan, K. Lerman, and E. Ferrara, "Discovering latent psychological structures from self-report assessments of hospital workers," in *Proceedings of the 5th International Conference on Behavioral, Economic, and Socio-Cultural Computing (BESC18)*, 2018, pp. 1–6.