Visual Informatics 5 (2021) 39–53

Contents lists available at ScienceDirect

Visual Informatics

journal homepage: www.elsevier.com/locate/visinf

ARGUS: Interactive visual analysis of disruptions in smartphone-detected Bio-Behavioral Rhythms

Hamid Mansoor *, Walter Gerych, Abdulaziz Alajaji, Luke Buquicchio, Kavin Chandrasekaran, Emmanuel Agu, Elke Rundensteiner

Worcester Polytechnic Institute, USA

A R T I C L E I N F O

Article history:
Received 13 April 2021
Received in revised form 2 June 2021
Accepted 5 July 2021
Available online 14 July 2021

Keywords:
Interactive visual analytics
Circadian rhythms
Smartphone-sensed data

A B S T R A C T

Human Bio-Behavioral Rhythms (HBRs) such as sleep-wake cycles (*Circadian Rhythms*), and the degree of regularity of sleep and physical activity have important health ramifications. Ubiquitous devices such as smartphones can sense HBRs by continuously analyzing data gathered passively by built-in sensors to discover important clues about the degree of regularity and disruptions in behavioral patterns. As human behavior is complex and smartphone data is voluminous with many channels (sensor types), it can be challenging to make meaningful observations, detect unhealthy HBR deviations and most importantly pin-point the causes of disruptions. Prior work has largely utilized computational methods such as machine and deep learning approaches, which while accurate, are often not explainable and present few actionable insights on HBR patterns or causes. To assist analysts in the discovery and understanding of HBR patterns, disruptions and causes, we propose *ARGUS*, an interactive visual analytics framework. As a foundation of ARGUS, we design an intuitive Rhythm Deviation Score (RDS) that analyzes users' smartphone sensor data, extracts underlying twenty-four-hour rhythms and quantifies their degree of irregularity. This score is then visualized using a glyph that makes it easy to recognize disruptions in the regularity of HBRs. *ARGUS* also facilitates deeper HBR insights and understanding of causes by linking multiple visualization panes that are overlaid with objective sensor information such as geo-locations and phone state (screen locked, charging), and user-provided or smartphone-inferred ground truth information. This array of visualization overlays in *ARGUS* enables analysts to gain a more comprehensive picture of HBRs, behavioral patterns and deviations from regularity. The design of ARGUS was guided by a goal and task analysis study involving an expert versed in HBR and smartphone sensing. To demonstrate its utility and generalizability, two different datasets were explored using *ARGUS* and our use cases and designs were strongly validated in evaluation sessions with expert and non-expert users.

1. Introduction

The healthcare system in the USA is under-resourced with patients receiving little care outside of appointments. Patient assessments are often infrequent, typically months apart, and often result in late diagnoses that worsen their prognoses. Emerging research is exploring the possibility of using sensor-rich smartphones that are owned by over 80% of the population[1] in the USA, to passively detect various illnesses, and gather health behavior and other contextual information. Ailments such as depression (Wang et al., 2014; Gerych et al., 2019) and influenza (Madan et al., 2011) can be detected by using sensor data collected from smartphones and applying machine learning models to them. This novel paradigm is called *smartphone health sensing or smartphone ailment phenotyping* (Onnela and Rauch, 2016) and can be applied to assess the important health-related habits of smartphone owners.

Humans are creatures of habit. Human bodies contain many small "biological clocks" that control various biological processes and regulate "Circadian Rhythms" ("Circa" means *about* and "diem" means *a day*) or "Bio-Behavioral Rhythms" (Kreitzman and Foster, 2011; Abdullah et al., 2017; Roenneberg, 2012; Matthews et al., 2016). These rhythms typically reflect twenty-four-hour cycles of human biological processes such as sleep-wake cycles, hormonal changes and blood pressure changes.

* Corresponding author.
E-mail addresses: hmansoor@wpi.edu (H. Mansoor), wgerych@wpi.edu (W. Gerych), asalajaji@wpi.edu (A. Alajaji), ljbuquicchio@wpi.edu (L. Buquicchio), kchandrasekaran@wpi.edu (K. Chandrasekaran), emmanuel@wpi.edu (E. Agu), rundenst@wpi.edu (E. Rundensteiner).

1 https://www.pewresearch.org/internet/fact-sheet/mobile/.

https://doi.org/10.1016/j.visinf.2021.07.001
2468-502X/

Disruptions in Human Bio-Behavioral Rhythms (HBRs) can have significant medical ramifications such as deteriorating mental health, obesity, and heart disease (Vetter, 2020; Roenneberg et al., 2012). Regularity of sleep and physical activity are also particularly important measures of health and their deviations have been linked to serious ailments such as psychiatric disorders (Walker et al., 2020) and are common across different ages and occupations (Ohayon et al., 2002; Gaultney, 2010). Roenneberg et al. (2012) conducted a large-scale study and found that up to 70% of the general population suffers from deviations in their sleep cycles.

Continuous monitoring of HBRs to detect deviations from normalcy can inform timely interventions and improve overall health. However, capturing and mining human behavior data to discover important patterns is challenging, especially in the real world. Smartphones are uniquely suited for capturing such data because they are ubiquitous and are equipped with multiple sensors such as accelerometers, gyroscopes, ambient light and GPS that can capture data to provide important clues about a person's behavioral rhythms. For instance, it is a common pattern for people to not interact with their phones overnight while they sleep. This lack of interaction can be captured by the smartphone and provide a useful measure of how much or little sleep people get (Abdullah et al., 2014). The ability to detect, monitor, and mine such patterns can make the smartphone a *proxy* for human behaviors. Smartphone sensing of human behavior has been used to reliably detect important behavioral changes (Wang et al., 2018) and monitor subjects' health and wellness, including mental health (Rabbi et al., 2011) and academic performance and stress situations (Wang et al., 2014).

While smartphones can gather very rich human behavior data quickly, the difficulty of discovering meaningful patterns and making sense of such data becomes more difficult as the number of participants, their duration of participation and the number of sensor streams analyzed increase. Prior work has largely utilized computational methods such as machine and deep learning approaches (Vaizman et al., 2018a,b), which often are not explainable and present few actionable insights on HBR patterns or causes. Interactive Visual Analysis (IVA) is a powerful approach for analysts to make sense of large multivariate data. In this work, we focus on leveraging IVA to enable human behavior analysts to discover and monitor disruptions in HBRs using autonomously-collected smartphone behavioral data, with the aim to generate actionable insights. We are interested in not only providing a way to identify HBR disruptions but equally important to *explain* those disruptions. We introduce ARGUS, a visual analytics framework to represent multi-stream, heterogeneous data using intuitive visual metaphors that enable analysts to make sense of the data with ease. ARGUS was designed and implemented for use by a diverse range of analysts in fields such as psychology, ubiquitous computing and data science.

Existing visual analytics frameworks for circadian rhythm are effective at detecting deviations in rhythm, but are not interactive and lack the ability to generate explainable insights (Geissmann et al., 2019; Fischer et al., 2016). Existing interactive visual analytics for human behavior disruptions generally focus on specific features of human behavior (Pu et al., 2011; Senaratne et al., 2017), such as mobility, but are not able to integrate multiple channels (sensor types and data streams) of smartphone-collected human behavior. In contrast, our approach leverages *interactive analysis* to provide linked visualization panes that leverage multiple channels to enable behavioral experts to discover deviations in HBRs as inferred from smartphone sensor data *and* explain them. We develop an intuitive Rhythm Deviation score, which then we visualize using a custom glyph which is based on an established visual metaphor called the z-glyph (Cao et al., 2018). Smartphone gathered data is also visualized in several linked panes that highlight how often sensor streams and certain extracted features *co-occur*. For instance, a locked smartphone screen combined with sensed low light and quiet conditions could indicate nights during which the user slept well.

Overall, our contributions are:

1. A novel intuitive HBR "*Rhythmicity Deviation Score (RDS)*" computed from autonomously gathered smartphone data using the Lomb–Scargle periodogram. It captures HBR disruptions that can be visualized for effective, fast, and reliable analysis.
2. The *ARGUS* IVA platform that visualizes our novel RDS using a glyph metaphor, while also linking other behavioral panes that contextualize the rhythmicity score. Together, ARGUS captures an explainable picture of users' HBRs. ARGUS not only visualizes deviations in behavioral rhythms but also provides analysts with the opportunity to uncover easily potential *explanations* of those deviations. Argus integrates IVA support for population-level HBR meta-analyses, easy identification of significant HBR disruptions, cross-channel exploration and contextualization of individual participant HBRs RDSs, and visualization of corresponding raw sensor values.
3. An expert-led goal and task analysis to articulate the challenges faced by analysts in the domain of smartphone health sensing and behavioral rhythms detection and the visualization goals and tasks to mitigate them.
4. A comprehensive evaluation of *ARGUS* using both expert feedback to determine feasible use cases for HBR disruption analysis, in addition to non-expert evaluation to determine the understandability of our visual metaphors and the overall ease of use of our tool.
5. An insightful walk-through along well-designed use cases by specialists in behavior studies involving two real-world datasets illustrates our approach and demonstrates its effectiveness for discovering, visualizing and explaining HBR disruptions.

This paper is an extension of a short paper (Mansoor et al., 2020a) that was presented at Eurovis 2020 (Eurovis, 2020) and was awarded an Honorable Mention in its track. This article has evolved and expanded upon our initial submission in several key ways: We provide clearer and more detailed background theoretical information about bio-behavioral rhythms, their health ramifications and the process for determining our novel rhythm deviation score. In addition, we provide more details about the rationale behind our design choices and the metaphors used, including alternative designs. We also provide a deeper review of the extant literature on the subject of interactive visual analysis of smartphone sensed data. Finally, we report the results from an evaluation with non-experts (in addition to an evaluation with an expert) to determine the efficacy of our visual metaphors and the ease of use of ARGUS.

2. Related work

2.1. Capturing human behavior using smartphones

IVA is a powerful approach for making sense of human behavior data (see Table 1). Gathering and analyzing human behavior data continuously for long periods is challenging *in-the-wild*. Smartphone sensing, wherein data from the smartphone's sensors are continuously gathered and analyzed to infer its owners' behaviors, has recently emerged as an inexpensive and scalable method of human behavioral analysis. Smartphone sensing has been effectively used in inferring depression levels (Saeb et al., 2016), social anxiety levels (Boukhechba et al., 2018b; Rashid et al., 2020), students' Grade Point Averages (Wang et al., 2014), and the smartphone owner's current context/situation (Vaizman et al., 2018a). Abdullah et al. (2014, 2017) even successfully used time periods during which a person's smartphone screen was locked as a reliable proxy for determining when they were asleep.

2.2. Visualizing human behavior patterns

Circadian Rhythm refers to how cyclical/regular a person's sleep-wake cycle and routine is, which has important health ramifications (Vetter, 2020; Walker et al., 2020). As such, monitoring these rhythms and intervening, if necessary, is crucial. Smartphones can prove extremely useful in this endeavor as they are ubiquitous and require low intervention. There has been a lot of research interest in capturing such behavioral patterns using computational models applied to smartphone-sensed data (Yan et al., 2020). Visualizations of human behavioral data can improve the state-of-the-art in this research area as they are useful for discovering, contextualizing, and understanding disruptions in circadian rhythms. Fischer et al. (2016) attempted to visualize "circadian misalignment" in shift workers using data from wrist-worn actimeter sensors combined with their sleep logs. They devised an intuitive method called "Composite Phase Deviation" which enabled them to generate and visualize density plots where the area and shape connote the extent of misalignment and variability in user behavior. They introduced the concept of "islands" and "pancakes" to refer to certain areas of the density plots to find variability in sleep data. Geissmann et al. (2019) created "Rethomics", a framework in the R language to analyze circadian rhythms. Their framework also implemented intuitive data visualizations to present circadian information about animals.

However, unlike our work, these frameworks do not visualize and link multiple smartphone information streams to produce interactive analysis that enhances *sense making and explainability*.

2.3. Interactive Visual Analysis (IVA) to detect anomalous human behavior

Deviation from normal human behavior may be indicative of health problems (Vetter, 2020). Data visualizations can be an effective method of identifying anomalous human behaviors and making deviations from normalcy clearer, which in turn enables timely interventions. Examples include the detection of mal-intents such as spreading unverified rumors (Resnick et al., 0000), committing financial fraud (van den Elzen et al., 2013) or scamming people (Koven et al., 2018). Cao et al. (2015) created TargetVue, an intuitive tool to detect anomalous behaviors on social media and detect bots. They introduced the concept of a *z-glyph*, which is an intuitive visual metaphor, expanded further in a subsequent work by Cao et al. (2018) to highlight deviations from normal behaviors. These prior IVA works are particularly relevant for detecting changes in human behavior and behavioral rhythms (gathered through smartphones) as they show how novel interactive visual solutions enable analysts to quickly generate insights and find useful patterns in digital human data. Our work adds to this field by leveraging IVA for *in-the-wild* human behavior disruption that has direct implications on health. Specifically, we use IVA to find disruptions in natural 24-h behavioral cycles.

2.4. Interactive visual analytics of sensor-based health data

Typical IVA tools for health analysis focus on utilizing data that was gathered explicitly for the purpose of health monitoring such as Electronic Health Records (EHR) records (Abdullah, 2020; Malik et al., 2015; Plaisant et al., 2003; Meyer et al., 2013) or health-sensing gadgets like Fitbits (Heng et al., 2018; Liang et al., 2016), carbon monoxide sensors (Polack Jr et al., 2018) and head-mounted sensors (Garcia Caballero et al., 2019). Polack et al. (2017) presented a position article in which they outlined several opportunities and challenges in using interactive visual analysis to advance the state-of-the-art in passive and opportunistic health monitoring through wearables such as smartwatches and smartphones. Our work expands on that by focusing more explicitly on smartphone-sensed data as smartphones are the most ubiquitously owned device with sensors. Health sensing using smartphones can be challenging since they were not designed explicitly for the task of health monitoring. IVA can bridge this gap by presenting the data gathered from smartphones to provide additional contextualization whether in terms of fixing data quality (Mansoor et al., 2019a,b) or giving the analyst the ability to utilize their human intuition and apply semantic labels to anonymized data (Mansoor et al., 2020b). Other works have utilized intuitive metaphors like the calendar metaphor (Gupta et al., 2017) to present daily activity levels and innovative interaction techniques like alignable daily timelines between daily repeating events (Zhang et al., 2018) such as meals for patients with diabetes. Such works show the utility of applying IVA to enhance the analysis of such data which typically relies on reliable data labeling and clear semantic context, both of which are often not available for in-the-wild collected data.

2.5. Interactive visual analysis of data gathered from personal digital devices

The vast proliferation of mobile phones has created many opportunities to gather rich datasets about human behaviors such as their mobility patterns and social interactions (Calabrese et al., 2015; Cuttone, 2017). IVAs can be useful for mining such data, contextualizing and explaining human behaviors. Pu et al. (2011) leveraged IVA that combined established visualization techniques such as parallel coordinates plots with intuitive, novel techniques such as "Voronoi-diagram-based" data visualizations to analyze the mobility patterns of three users. Senaratne et al. (2017) use an IVA approach to analyze spatial and temporal similarities in human movements using a passively gathered mobile phone dataset. They employed matrix visualizations of the user movements. These prior works illustrate the usefulness of IVA to explore and understand human movement (a very important facet of human life), its variations, patterns and disruptions.

IVA techniques can be further augmented with novel glyphs and visual metaphors that are useful for representing complex mobile phone data. Shen and Ma (2008) created MobiVis, an IVA tool that implemented the "Behavior Ring", a radial metaphor to represent individual and group behaviors compactly. Their tool enabled intuitive visual data mining by semantic filtering to facilitate effective analysis of "social-spacial-temporal" data that phones gather. This approach illustrates the utility of compact visual metaphors and IVA to understand complex phone data.

Unlike our method, these methods do not incorporate the concept of *cycles* and *rhythms* and the disruptions thereof. Specifically, we facilitate interpersonal/intergroup analyses in order to identify users of interest that are then analyzed in greater detail.

3. Goal and task analysis: Interactive visual analytics to monitor bio-behavioral rhythms

Given that this domain contains such diverse and heterogeneous datasets from which many features are typically extracted, Interactive Visual Analytics (IVA) can assist in making sense of

Table 1
Interactive Visual Analytics (IVA) works for human behavior analysis.

IVA topic	Behavior type Visualized	IVA works and aspect visualized	IVA techniques
Anomalous behavior	Fraud	Dis-information (Resnick et al., 0000) Financial Fraud, scams (van den Elzen et al., 2013)	Sankey diagram, timelines Circular chord diagrams, timelines
	Social Media and communication	Bots and emails (Cao et al., 2015) Bots (Cao et al., 2018)	Interactive glyphs, timelines Interactive glyphs
	Suspicious Online Interactions	Misuse of admin privileges (Nguyen et al., 2020, 2018)	Semantic linking
Personal digital device data	Social circle	Friend networks, calls, texts (Shen and Ma, 2008)	Node-link diagrams, timelines
	Population level movement patterns	Urban mobility patterns (Cuttone, 2017; Senaratne et al., 2017) Smartphone-detected mobility (Pu et al., 2011)	Matrix representation and maps Voronoi diagram, maps
	Detected activities	Fitbit activity (Gupta et al., 2018a,b)	Interactive glyphs, timelines
Health analysis	Electronic health record	Patient hospitalizations and medications (Malik et al., 2015; Plaisant et al., 2003) Renal problems (Abdullah, 2020)	Interactive timelines, discrete event overlay Clustering
	Smoking	Relapse (Polack Jr et al., 2018)	Node-link diagrams, interactive timelines
	Geriatric care	Indoor movement (Payandeh and Chiu, 2019) Smart home (Le et al., 2014)	Node-linked graphs, circular histograms Stream graph display, circular chord diagrams
	Sleep	Quality and quantity (Choe et al., 2015) Quality and quantity (Liang et al., 2016)	Histograms, emojis Histograms, bubble charts
	Medication	Medication error (Kakar et al., 2019b) Drug–drug interactions (Kakar et al., 2019a)	Treemaps, interactive timelines Force directed layouts

this data and identifying and monitoring people with deviations in their behavioral patterns.

We conducted goal and task analysis sessions with an expert in bio-behavioral rhythms who was also experienced in analyzing human behavior gathered using ubiquitous sensing devices. The expert was particularly interested in rhythms related to sleep behavior and how certain events may disrupt sleep patterns. For instance, the buildup of stress due to uncontrollable, external factors may cause lost sleep, which may have health ramifications.

We discussed how a smartphone may collect data indicative of disruptions and breakages in patterns. Some of her suggestions included conceptualizing smartphone data as *channels* of information, which may provide important clues about a person's *contextual* information. Examples of channels are the state of the smartphone including screen locked, battery charging, apps being used and its GPS location. In order to derive how rhythmic a person is in their daily routines, it may be useful to find rhythms, disruptions and breaks in these channels.

Given that human behavior is complex and the smartphone channels used to make behavioral inferences can become overwhelming, IVA can be useful for making sense of, and correlating these information channels. Given that a correlation in these channels may be meaningful, (the expert suggested that correlating darkness with a lack of sound and screen being locked might be a useful method of detecting sleep). Such meaningful linkages and correlations across different channels would be difficult to show using non-visual statistics. IVA may powerfully augment an analyst's ability to make sense of people rhythms by intuitively overlaying and correlating complex channel data to increase its interpretability. The expert suggested two broad goals that she would have as someone analyzing smartphone collected human behavior data:

- **G1: Discover overall levels of behavioral rhythms:** and times during which breaks occur
  - Synthesizing an overall numerical measure or *score* to capture and quantify a person's bio-behavioral rhythm, which can then be represented visually to reveal deviations from normalcy.
- **G2: Explain and contextualize causative factors:** that led to deviations from normal rhythms across multiple channels:
  - The multiple channels of smartphone-gathered data can provide a multi-faceted view into a person's behaviors. The expert wanted several linked views of different channels such as a person's geo-location or the intensity of their smartphone interactions, which may explain the reasons for disruptions and breakage in their patterns. Such multi-view context may also enable the disambiguation of harmful disruptions in behavior patterns such as staying up all night (a sign of depression) versus benign pattern disruptions such as a person who is traveling for vacation.

We also discussed the specific tasks that the analyst would like to be able to perform to achieve the goals described above. Given our collective knowledge of the data in this domain, we devised the following tasks to achieve the goals described.

- **T1, Population-level meta-analysis**: Get a quick overview of the level of rhythmicity of the behavior patterns of all participants in a study to quickly find the ones with the most and least rhythms.
- **T2, Anomalous HBR identification**: Identify participants with significant disruptions in their behavioral rhythms quickly.
- **T3, Cross-channel HBR exploration**: Examine and contextualize individual participants' rhythmicity levels across multiple "channels". For instance, visualize physical activity vs geo-location disruptions.

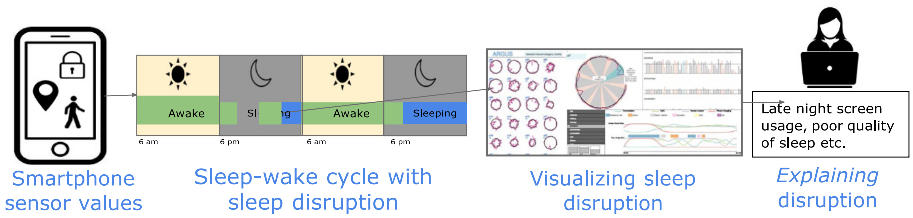

Fig. 1. The workflow of detecting rhythm deviations and then visualizing them using ARGUS.

- **T4, HBR contextualization**: Highlight contextual factors that might have bearing on rhythm such as day of the week, or weekday vs weekend.
- **T5, Raw sensor value drill-down**: Visualizing the values of sensor readings, which may hold important clues about what a person did around the time of disruptions in their HBR. For instance, the amount of time they spent interacting with their smartphone's screen.

Extant work has shown that smartphone channels such as phone state measurements are indicative of bio-behavioral trends that have significant implications for a person's health (Mohr et al., 2017). Abdullah et al. (2014) were able to predict research participants' sleeping habits using data about whether their smartphone's screen was locked. Other works have used ambient light sensor data from smartphones to detect the occurrence of sleep (Min et al., 2014; Chen et al., 2013). These channels' ability to detect very important bio-behaviors (sleep for instance), make their co-occurrence, lack of co-occurrence and disruptions in those channels very interesting.

4. ARGUS

To achieve the goals and tasks identified above, we researched and developed *ARGUS*, an interactive multi-pane visualization tool. Since the data that *ARGUS* would be dealing with is highly multi-variate and heterogeneous, we adopted the "Visual Analytics Mantra" proposed by Keim et al. (2008), which states: "Analyze First - Show the Important - Zoom, Filter and Analyze Further - Details on Demand". Here we describe our Rhythm Deviation Score for quantifying the level of regularity of HBRs, and visualizations we created and the rationale behind our choices.

4.1. Rhythmicity deviation score

We now expound on our novel Rhythmicity Deviation Score, a single score we synthesized to quantify the degree of regularity of a person's circadian rhythm based on data gathered from their smartphone sensors. The visualization panes of *ARGUS* then visualize this score along with contextual information. Our Rhythmicity Deviation Score is based on the Lomb–Scargle periodogram (Lomb, 1976; Scargle, 1982), a classic method for finding periodicity in irregularly-sampled data.

Definition: Channel

A *channel* C_i is given by a time series: $\big((c_{i,0},t_{i,0}),(c_{i,1},t_{i,1}),...,(c_{i,k},t_{i,k})\big)$, such that each $c_{i,j}\in\{0,1\}$.

We define "Channels" as sequences of binary variables that can represent smartphone-inferred or self-labeled/self-reported behavioral indicators such as physical activity (instances of time where the user was walking, sitting, etc.) or objective sensor readings such as whether the smartphone was locked, connected to a wireless network, or charging. In order to quantify changes in channel behaviors we must first define the *occurrence ratio*, the length of time for which the channel was a positive instance (i.e., subject *were* walking) over a certain time scale.

Definition: Lomb-Scargle Periodogram

Let $X=\big((y_0,t_0),(y_1,t_1),...,(y_k,t_k)\big)$ be a time series, such that $\sum_{i=0}^{k} y_i = 0$. The Lomb-Scargle periodogram, P_{LS}, of X for frequency ω is given by:

$$\mathrm{P}_{LS}(\omega)=\frac{1}{\sum_{i=0}^{k} y_i^2}\left\{\frac{[\sum_i y_i\, cos\,(\omega(t_i-\hat{\tau}))]^2}{\sum_i cos\,(\omega(t_i-\hat{\tau}))}+\frac{[\sum_i y_i\, sin\,(\omega(t_i-\hat{\tau}))]^2}{\sum_i sin\,(\omega(t_i-\hat{\tau}))}\right\},$$

where $\hat{\tau}$ is a time delay parameter such that the sinusoids are mutually orthogonal at each sample time t_i. $\hat{\tau}$ is given by:

$$tan\,(2\omega\hat{\tau})=\frac{\sum_i sin\,(2\omega t_i)}{\sum_i cos\,(2\omega t_i)}$$

An example of the Lomb–Scargle periodogram for 1 channel (Sleeping) is shown in Fig. 3. The peak of the periodogram occurs at 1 day, indicating that this user is relatively cyclic in their sleep habits on a 24-h cycle.

In order to use the Lomb–Scargle periodogram to identify disruptions in user behavior, we apply the periodogram on each individual *channel* of user behavior.

Definition: Occurrence Ratio

The *occurrence ratio* O_r of channel C_i for day D_n is the ratio of positive instances of C_i that occurred during day D_n. Let $C_{i,D_n}=\{(c_{i,j},t_{i,j}|t_{i,j}\in D_n)\}$. Then, the occurrence ratio is given by:

$$O_r(C_i,D_n)=\frac{\sum_{c_{i,j}\in C_{i,D_n}} c_{i,j}}{\|C_{i,D_n}\|}$$

As we are interested in investigating the users' circadian rhythm, which is a 24-h cycle, we choose 1 day to be the time scale over which this value is calculated. We define the Occurrence Ratio as the proportion of the day for which the channel is positive. The denominator is the total number of instances which would include both positive (1s) and negatives (0s). For instance, this would represent the proportion of time spent in the dark, or spent performing some reported activity. We can also define the *average occurrence ratio*, in order to typify the user's usual behavior.

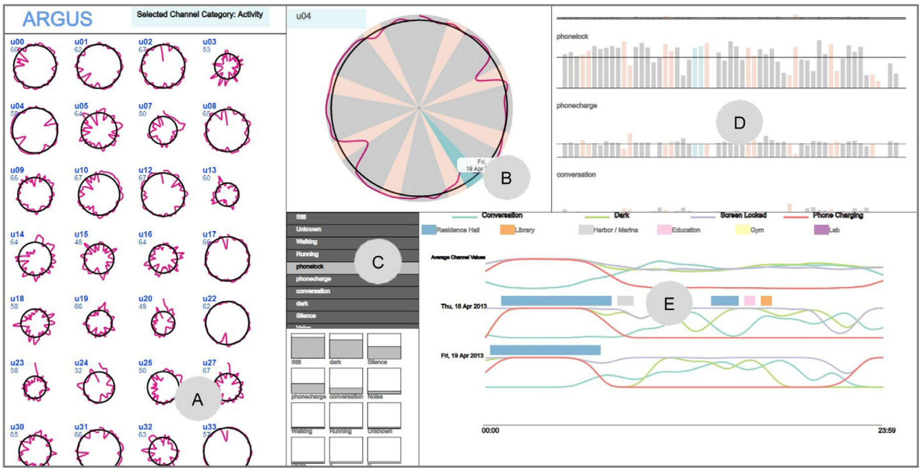

Fig. 2. ARGUS: A multi-pane visualization framework to discover and explain deviations in bio-behavioral rhythms. **A: Eyes of ARGUS**. This view provides a quick overview of the levels and breaks in bio-behavioral rhythms across all participants. **B: Magnified Eye of Argus**. A user can select an eye to magnify and gain a clearer look into the rhythm changes. The lighter color represents weekends and the rest are weekdays. **C: Co-occurrence View**. Provide an overview of how consistently channel values co-occur to enable analysts to explain behavior in terms of a lack of co-occurrence of frequently co-occurring channels. **D: Duration View**. A summary for how much a channel had positive values throughout the day. **E: Explainability View**. Visualizing and linking multiple channels to gain a greater understanding about the causes of breaks in rhythms.

Definition: Average Occurrence Ratio

The *average occurrence ratio* $\overline{O}_r$ of channel C_i is the average ratio of positive instances of channel C_i over all days in the set of days $\mathcal{D}$:

$$\overline{O}_r(C_i) = \frac{\sum_j O_r(C_i, D_j)}{\|\mathcal{D}\|}$$

We can now define the *circadian rhythm* of a given channel:

Definition: Circadian Rhythm

The *circadian rhythm* of channel C_i measures how cyclic C_i is on a frequency of 1 day. The circadian rhythm is defined as:

$$R(C_i) = \frac{\int_{\frac{1}{D}-\Delta t_1}^{\frac{1}{D}+\Delta t_1} P_{LS}^{C_i}(\omega)\, d\omega}{\int_{\frac{1}{D}-\Delta t_2}^{\frac{1}{D}+\Delta t_2} P_{LS}^{C_i}(\omega)\, d\omega},$$

such that $\Delta t_1 < \Delta t_2$ and $P_{LS}^{C_i}$ is the Lomb-Scargle periodogram of C_i.

This definition is based off of Wang *et al.'s* definition of circadian rhythm (Wang et al., 2018), though ours differs in that we use the Lomb–Scargle periodogram rather than the *power spectral density* as we deal with unevenly sampled data. While this definition may seem obtuse, it has an intuitive basis. The circadian rhythm is simply the integral of the periodogram for a small region around the frequency associated with 24 h divided by a similar integral taken over a larger range of values. Intuitively, if nearly all the power of the periodogram is concentrated at the 24 h mark (that is, their rhythm is nearly perfectly described by a 24-h cycle), then this value should be close to 1. Otherwise, this ratio goes to 0 (indicating that the user is *not* rhythmic on a 24 h cycle for the given channel). We follow the convention established by Wang et al. (2018), and set Δt_1 to be $\frac{1}{2}$ of an hour, and Δt_2 to be 12 h. In the Circadian Rhythm definition, D corresponds to the length of 1 day. The D_j in Average Occurrence Ratio definition is some particular day with length D, and the D in the average occurrence ratio definition is the set of all days.

Having defined the circadian rhythm, we now define disruptions in the circadian rhythm.

Definition: Channel Rhythm Disruption

The *channel rhythm disruption* CD of C_i for day D_j measures the difference in occurrence ratio of C_i for D_j and the average occurrence ratio of C_i, weighted by the circadian rhythm of C_i. CD is given by:

$$CD(C_i, D_j) = R(C_i) \cdot \|O_r(C_i, D_j) - \overline{O}_r(C_i)\|$$

The *channel rhythm disruption* is simply the change in behavior of a particular channel (that is, changes in the duration of positive instances of this channel) weighted by the circadian rhythm of this channel. The reason why we weight the change in behavior by the circadian rhythm is that we only want to identify *meaningful* disruptions in behavior. For instance, assume that the behavior of the channel C_i differs from the average for a given day. However, the average behavior is meaningless if the user is not rhythmic for that channel. By weighing the change in behavior by the circadian rhythm of that channel, the deviations in non-rhythmic channels will result in only a small channel disruption score.

Additionally, we wish to quantify disruptions in user behavior on a scale larger than single channels. For this reason, we define *channel categories*.

Definition: Channel Category

A *channel category* G_k is a set of channels. Any given pair of channel categories have no elements in common. Let $\mathcal{G}$ be the set of all channel categories and let $\mathcal{C}$ be the set of all channels. Then, for a pair of channel categories G_k and G_j,

$$G_k \cap G_j = \varnothing \; \forall \; G_k \neq G_j \in \mathcal{G}$$

Additionally, we require that all channels belong to a category. Thus:

$$\bigcup_{G_k \in \mathcal{G}} \bigcup_{C_j \in G_k} C_j = \mathcal{C}.$$

For our purpose, we group our channels into 3 categories: objective sensors, geo-location, and activities. More details about the channels in each category are given in Table 2.

Finally, we can now define the *rhythm deviation*, which is the metric we use to identify individuals who have experienced significant changes in behavior.

Definition: Rhythm Deviation

The *rhythm deviation* RD of category G_i for day D_j measures the average *channel rhythm disruption* for each channel in G_i on day D_j.

$$RD(G_i, D_j) = \frac{\sum_{C_k \in G_i} CD(C_k, D_j)}{\|G_i\|}$$

Next, we will discuss our IVA framework for investigating rhythm deviation and channel rhythm deviations.

4.2. Eyes of ARGUS

The pane shown in Fig. 1A, contains the Eyes of ARGUS (EA). Every glyph plots a person's average Rhythm Deviation (RD) score as a black circle (a circle with a larger circumference represents more overall rhythm) against the daily RD scores ordered in a clockwise direction. We average the RDS for every day and all channels for the selected channel category (selectable by the analyst) and encode it as the black circle — the higher the overall RDS, the lower the circle diameter. The positive (outward) deviations in the purple line represent days which adhere closely to the detected rhythm in the selected channel whereas the negative (inward) deviations represent rhythm disruptions. The strength of deviations is encoded by inwards distance from the black circle i.e. the more the purple line is towards the center, the higher the level of deviation (T2). Towards the top left of every EA is the user id and the number of days of participation. EA shows the deviations from rhythm, one circle per user, for all the days of participation.

This glyph is based off the Z-glyph developed by Cao et al. (2018) to effectively visualize *deviations* from the norm. Their analysis showed significant improvements in discerning deviations when using the z-glyph over traditional line glyphs. This is particularly applicable for this problem as deviations from normal rhythmicity or punctures in rhythms are exactly what ARGUS tries to highlight. Fig. 2A gives the analyst a quick overview of the level of bio-behavioral rhythmicity of users in general (G1,T1) and any interesting users that they might want to explore further (T2). The z-glyph family implemented both linear and star (circular or radial) metaphors. We decided to use the radial glyph because the results from a user study by Cao et al. (2018) in the same work established that the study participants preferred the radial glyph metaphor over the linear glyph metaphor in terms of efficiency and user comfort. Both the radial and linear metaphors had comparable results for accuracy of reading and both glyphs outperform the traditional glyphs. We gave preference to user comfortability which is important as ARGUS targets health experts who will likely not be experts in working with data visualizations.

Clicking on any EA will magnify it in the Magnified Eye of ARGUS (MEA) pane (Fig. 2B). The underlying circle has the maximum possible circumference (i.e. if a person is perfectly rhythmic). The circle is divided into sectors with each sector representing one day of participation. The days are ordered in a clockwise direction so that contiguous days are next to each other. ARGUS can work with non-contiguous days of data collection as well. This is an important consideration since there are cases where the study participant may turn off data collection on their phone. The sectors with the lighter color are weekends, while the gray ones are weekdays (G2,T4). The user is guided by the RD score and can select multiple days by clicking on the sectors (selecting a day changes its color to) to view them in more detail in Channel Duration view (Fig. 2D) and the Explainability View (Fig. 2E). Multiple day views are useful as events from a preceding day may have had an affect on the current day (G2).

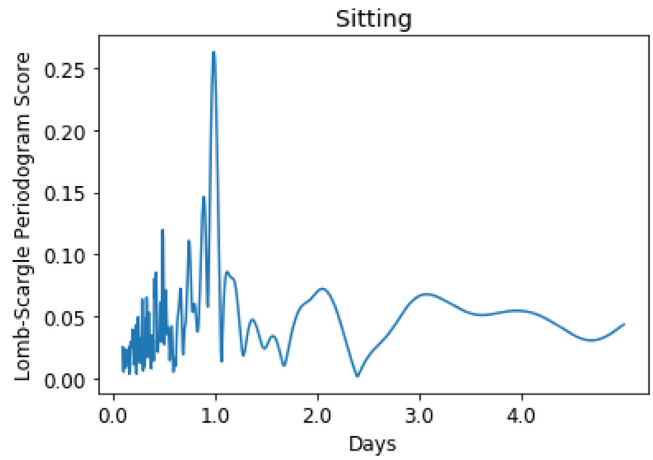

Fig. 3. The graph of the Lomb–Scargle periodogram of a particular user's sleep data. The *x*-axis represents the time scale (in days), and the *y*-axis represents the value of the Lomb–Scargle periodogram evaluated at frequencies corresponding to the time scales. We see a peak in the periodogram near 1 day, indicating that this user is relatively cyclic in their sleep habits on a 24-h bases.

The user can also select the channel category to visualize in the EA and MEA by clicking on the "Selected Channel Category" bar at the top of Fig. 2A (G2,T3). The mean rhythm score across the three channel categories is available for visualization. The specific channels for the 2 datasets are described in Table 2. The "Sensors" and "Geolocation" channels are similar across the two datasets used in that they are passively sensed by the phone. The "Activities" channels differ between the two datasets. For StudentLife (Wang et al., 2014), the activities such as walking or still are predicted by the phone with an on-device detection and classification algorithm whereas the activities in the Study1b dataset (our locally collected dataset), are self-labeled by the study participants i.e. passively gathered sensor is annotated with activity labels provided by the user.

4.3. Duration view

The Duration View (DV) (Fig. 2D) shows the overall "duration" of occurrences of a channel for every day of the participant's data. Duration means the amount of time during each day that the channel was "on". For instance, the duration of time a phone was plugged in or being present in a particular geo-location cluster or labeling/self-reporting the duration of some activity happening on a smartphone such as being asleep for a few hours. The panel containing the Duration View is scrollable and the duration across days is shown separately for every channel present (across the different channel categories), as bar plots. Every vertical bar represents a single day of participation and the height of the bars represents the overall duration per day for the channel. The horizontal line for every channel is a visual indicator for the mean duration. It follows the same coloring scheme as the pie layout in the MEA. In the example, in Figure 2D, the user's phone is locked throughout the day generally more than their phone is charging. This is meant to aid the analyst in explaining rhythm disruptions by providing the overall duration of various levels of channel disruption (G2, T2, T5).

Table 2
Description of data channels across the 2 datasets. Among the differences between datasets is that activities are inferred in StudentLife while they are user-provided in Study1b. Additionally, Study1b contains a richer set of activities.

Dataset	Channel	Inferences/sensing/self-reports	Frequency	Health rationale for collecting sensor from literature	Rhythm Detection literature
StudentLife	Activities (inferred)	Still Walking Running Unknown	Gathered for 1 min between 3 min intervals	Depression (Madan et al., 2011; Walker et al., 2020), Sedentary Behavior (Kerr et al., 2016)	Huang et al. (2018), Wang et al. (2018) and Matthews et al. (2016)
	Sensors	Phone locked Phone charging Phone in dark	Gathered whenever interval >1 h detected	Fatigue (Mohr et al., 2017) Sleep cycles (Saeb et al., 2017; Abdullah et al., 2014, 2017; Wang et al., 2014; Chen et al., 2013) Concentration (Mohr et al., 2017; Dingler et al., 2017)	Huang et al. (2018), Wang et al. (2018) and Matthews et al. (2016)
		Conversation	Gathered for 1 between 3 min intervals (continue gathering if detected)	Stress, Depression and Affective State (Wang et al., 2014)	
	Geolocation	GPS coordinates	Every 10 min	Depression (Saeb et al., 2016), Social Anxiety (Boukhechba et al., 2018a)	Saeb et al. (2016), Canzian and Musolesi (2015) Boukhechba et al. (2018a), Rashid et al. (2020), Xu et al. (2019)
Study1b	Activities (labeled)	Phone on table, facing down Stairs — Going Down Sleeping Stairs — Going Up Laying down Phone in bag Phone in pocket Typing Walking Phone one table, facing up Exercising Phone in hand Sitting Running Bathroom Jogging Exercising Standing	Sensor data gathered for 20 s between 1 min intervals and labels applied study participant	Sedentary Behavior (Kerr et al., 2016), Sleep (Ciman and Wac, 2019; Wang et al., 2014) Fatigue (Mohr et al., 2017) Mental disorders like schizophrenia (Ben-Zeev et al., 2017; Buck et al., 2019)	Huang et al. (2018), Wang et al. (2018) and Matthews et al. (2016)
	Sensors	Phone locked Charging WiFi connected	Every minute	Sleep cycles (Abdullah et al., 2017, 2014; Chen et al., 2013) Concentration (Mohr et al., 2017; Dingler et al., 2017) Infectious Diseases (Madan et al., 2011)	Huang et al. (2018), Wang et al. (2018) and Matthews et al. (2016)
	Geolocation	GPS coordinates		Depression (Saeb et al., 2016; Xu et al., 2019), Social Anxiety (Boukhechba et al., 2018a)	Saeb et al. (2016) and Canzian and Musolesi (2015) (Boukhechba et al., 2018a; Rashid et al., 2020; Xu et al., 2019)

4.4. *Co-occurrence view*

The co-occurrence prevalence between certain channels is also interesting for researchers as a break in co-occurrence may indicate or explain a break in behavioral rhythm. Clicking on a user's EA shows the Co-occurrence View (Fig. 2C). This view has a list of all channels available for the selected user. Clicking on a channel bar shows in an ordered format (left to right, top to bottom) the *most commonly co-occurring channels* bars, that is channels that were "on" at the same time as the selected channel. For example, being in the dark coinciding commonly with the phone being charged simultaneously. The gray fill in the bars is proportional to the frequency of co-occurrence. For instance, for this person, the phone being locked mostly coincided i.e. happened at the same time as the person was detected by their phone as being "still", being in the "dark" and being in "silence". This is meant to let the analyst understand further the causes of a break in rhythms as the absence of channel co-occurrence among mostly co-occurring channels in interesting (G2,T5). To link the occurrence of co-occurring channels, hovering over a commonly co-occurring bar highlights the durations of the selected channel and the hovered over the channel in the Explainability view. An example of how this occurs is illustrated in Fig. 4.

4.5. *Explainability view*

The Explainability View (EV) (Fig. 2E) aims to provide a finer day-level view of the collected data to assist analysts in figuring out potential causes of a change in a subject's rhythm (G2,T5). The lines show the variation in the channel data available for the phone data. The height of the lines represents if the channels' average channel values per hour. The first plot shows the average channel values across the entire period of collection. Every plot after that represents a specific day chosen in the MEA. For instance on 14th April (Fig. 2E), the phone was charging through the night into the morning. As this is laid out on a horizontal timeline against a common scale, changes over hours are easily noticeable. The colored bars over the lines represent the durations

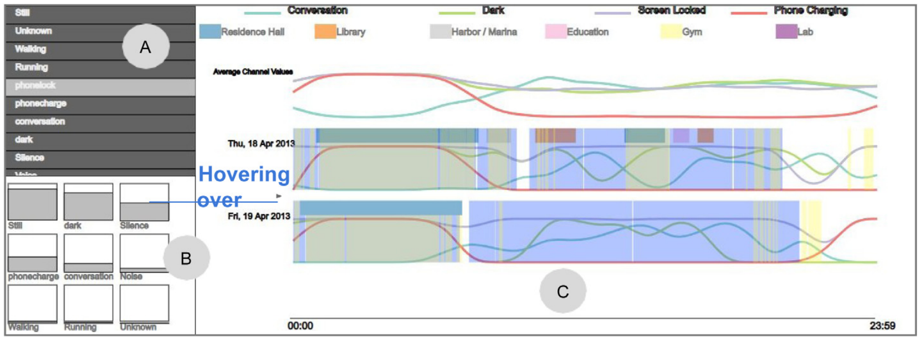

Fig. 4. In the co-occurrence View, the user can select a channel in A to see its most commonly co-occurring channel bars (shown in B), ordered from left to right, top to bottom in terms of frequency of co-occurrence. The vertical gray fill in bars in B is proportional to frequency of co-occurrence with A. Hovering over any of the bars in B highlights the duration of the channel selected in A across the Explainability View in light blue while the hovered over bar's channel gets highlighted in light yellow . The durations in the Explainability View with a light green overlay are when the 2 channels co-occur.

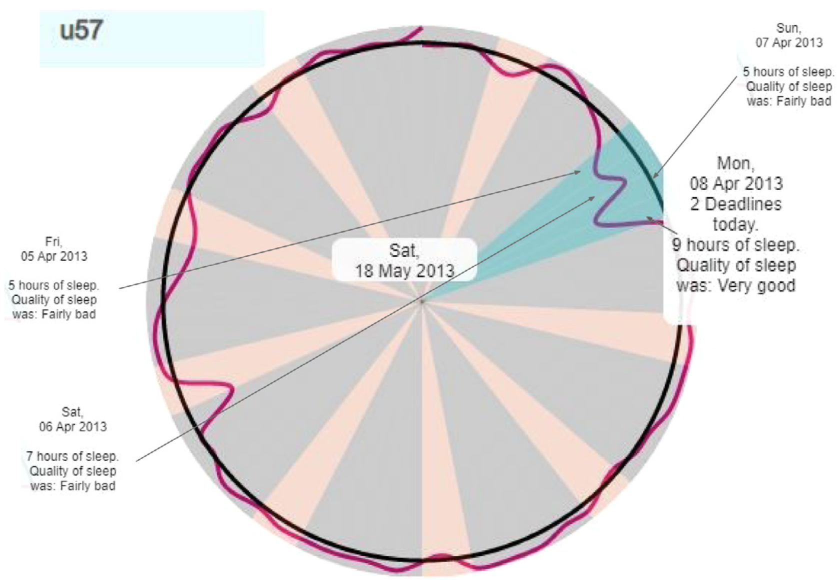

Fig. 5. Hovering over a day sector brings up a tool-tip that shows the human readable date and any other day related information (if any provided). The analyst can clearly see a disruption in rhythm and the reduction of sleep duration and quality leading up to a day with two deadlines.

of time for which they were in the same geographic cluster. Both the datasets we utilize in our case studies contained geo-locations of the participants throughout the day (whenever available). We ran a clustering algorithm called DBSCAN (Ester et al., 1996) to find geo-clusters for the participants. Behavioral rhythms, including patterns of presence at different locations, differ across people due to several factors such as differences in schedules and social requirements (Vetter, 2020; Saeb et al., 2016). In addition, even for data gathering among co-located people such as students at the same institution or employees at the same workplace, location presence may vary widely. Therefore we utilized a within-subjects approach for clustering, meaning every participant's geo-location data was clustered separately from other participants. The granularity of the clustering was set at 300 m. The clusters are encoded with colored bars (legend shown at the top of the EV. Given limited visual real estate we only show the top 6 clusters in which the participant was present for most of the time. The legend (Fig. 2E) shows the colors for the lines and the clusters. The colors were selected using a 10-class qualitative palette from ColorBrewer (Brewer and Harrower, 2010) to ensure that they were discernible. The human-understandable categories for the clusters were gathered by running the cluster coordinates against the Foursquare API (Foursqaure, 0000).

5. Illustrative use cases

To illustrate the usability of ARGUS, we introduce Emma, a graduate student in psychology who specializes in human behavioral rhythms and their effects on human health, especially for college students. Emma has access to two different real-world datasets that she is able to visualize using *ARGUS*.

5.1. Dataset 1: StudentLife

The first dataset we used for evaluating *ARGUS* is an open source dataset gathered from a smartphone sensing project called StudentLife (Wang et al., 2014). Smartphone sensor data was collected and analyzed to infer various participant behaviors including their GPA and physical state (e.g. still vs walking). The audio of scenes the user visited was also analyzed to make inferences such as whether the person was in a silent environment versus a noisy environment, or conversing. The dataset contains information for 49 students at Dartmouth College in the USA, who were enrolled in a specific mobile computing class for a 10-week academic term. This dataset records only significant durations (greater than or equal to 1 h) when the phone was in a dark environment, charging and the screen was being locked.

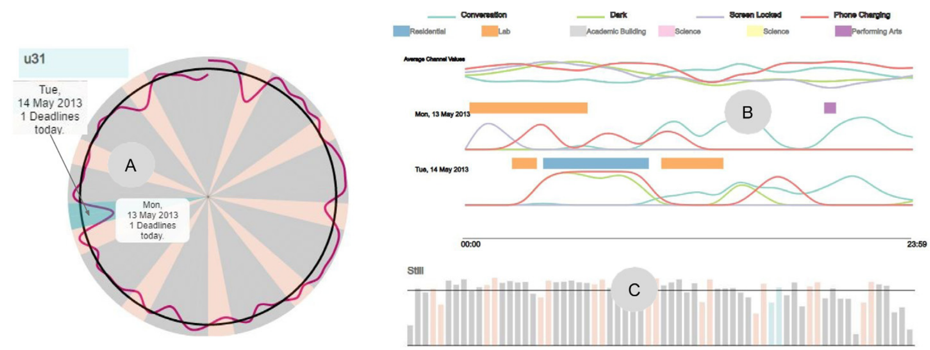

Fig. 6. May 13th and May 14th have disruptions in rhythm and both days have deadlines. Exploring these days further in the Explainability View reveals that the participant's dark, conversation, screen locked and phone charging channels are all off along with the fact that they are in a lab from around the beginning of day on May 13th.

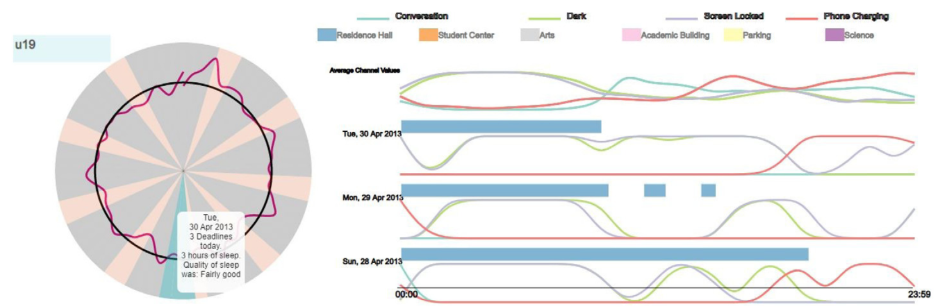

Fig. 7. There are 3 Deadlines on 30th April. In the preceding days, we is much less conversation than average. Such breaks in social activities are highly interesting for Emma.

The data contains geo-location (whenever available) of the device throughout the day.

The StudentLife project also collected subject responses to daily mental health questionnaires with wellness information such as their sleep duration ("How many hours did you sleep last night?"), sleep quality ("How would you rate your overall sleep last night?"), and stress levels etc. The answers to such questions were given on varying scales, which were provided in the documentation for the StudentLife (Wang et al., 2014) dataset and transcribed in human- understandable form (E.g. "Very Good", "Fairly Good", "Fairly Bad" for quality of sleep) for interpretable viewing on ARGUS. Data was gathered from the beginning of a 10-week academic term and the entire collection time period was approximately 10 weeks. Analysis of this dataset provides a clearer understanding of how student behavior changes over the course of an academic term.

Emma's visualization of the bio-rhythms of StudentLife students using *ARGUS* :

- *Identifying and contextualizing bio-rhythm disruptions caused by deadlines:* Emma is particularly interested in exploring changes in student behavior around more stressful times such as project due dates and deadlines (G1). The students were asked to provide the academic deadlines that occurred during their the days they participated in the StudentLife study. Emma takes a look at the eyes to see if there is any participant that sticks out (T1). She notices that participant u57 has a large inward spike (T2), indicating a large deviation of their bio-rhythm from normalcy. She clicks on it to magnify it in the MEA (Fig. 5). When she hovers over the day sectors for which the RD score was high, she noticed that the participant had low sleep duration and quality for three straight days (April 5th, 6th, and 7th) leading up to April 8th, which had two deadlines (G2, T4). After these two deadlines passed on the 8th, their sleep duration and quality, as well as their Rhythm Deviation score all improved. This leads Emma to believe that the stress caused by imminent deadlines disrupted the participant's bio-rhythm (G2). Similarly, as these deadlines passed, participants bio-rhythms returned to normal. Visual overlay of multiple panes human-understandable data along with objective rhythmicity scores and objective calculations in *ARGUS*, made this insight easy for Emma.
- *Explanations of bio-rhythm shifts using the Explainability View (EV):* After observing the effect that deadlines had on students' sleep patterns, Emma wants to see if she can observe a reduction in the quantity and quality of their sleep during days when participants did not respond to sleep questions. While exploring the data for u31, she notices a disruption in bio-rhythm towards the end of the term (T2). Hovering over the 2 days during which the bio-rhythm was disrupted revealed that there was a deadline on both days (Fig. 6A). Emma clicked on the 2 day sectors to view them in the Explainability View (EV) and the duration average view. In the duration view, she notices that the participant had lower levels of the "Still" state on May 13th (Fig. 6C). The participant had much lower levels of screen locked and being in a dark environment in EV (Fig. 6B). She also noticed that the participant was in a "Lab" in the very early hours of the morning. The participant was again recorded as being in the lab for the early hours of the morning and then being in a "Residential" location where they plugged their phone and were in a dark environment for a significant amount of time. The detailed and comprehensive system of overlaying various channel data in *ARGUS* enabled Emma to pinpoint, contextualize and understand a potentially concerning disruption in the students' bio-rhythm.
- *Detecting changes in other sensor channels:* Emma is also interested in the social behaviors of students around stressful

times. She notices u19 who does not have a high overall rhythm in their data (T2). As she explores their data, she notices a day with three deadlines. She explores the day and the days leading up to it in the EV and notices a steep drop in the amount of conversation this student was having (Fig. 7) (T4, T5). Emma is interested in discovering changes in social behavior rhythms caused by academic stress (G2).

- *Exploring the relationship between geo-location and bio-rhythms:* Small disruptions in students' bio-rhythm during a term may not be a major cause for concern (Vetter, 2020). Emma wants to investigate how students' geo-locations affect the rhythmicity of their bio-scores. She selects "Geo-location" (T3) in the Rhythm Selector (in the top pane in Fig. 2A) and notices u46 had a large deviation in their geo-location rhythm. Based on the color of the underlying slides, Emma notices that this large deviation occurred on the weekend (T4). She clicks on the non-rhythmic days (Friday, Saturday, and Sunday) and notices in the EV that the participant had a small geo-location recording showing that they were at a "hotel" and more geo-location readings for the 2 weekend days (Fig. 8). This leads Emma to believe that this was not a cause for major concern as it was an isolated incident related to traveling. The visual overlay of these various channels allowed Emma to disambiguate a potentially concerning rhythm disruption as merely being an innocuous one.

5.2. Smartphone sensor data gathered locally from our campus — Study1b

The second dataset we utilized to evaluate *ARGUS* was gathered around our own campus and will be referred to as Study1b. It contains smartphone sensor data for 103 people. Our approach was different because we did not constrain ourselves to a term like StudentLife and gathered data over a number of time different periods. Our participant population was also more diverse demographically and included teen-aged undergraduates to middle-aged campus office workers. We also had a shorter average period of participation (two weeks). We used a modified version of the ExtraSensory Android application, developed by Vaizman et al. (2018a). The application gathered sensor data for 20 s of every minute. Unlike StudentLife, this app did not suggest any inferred activity and just let the user provide activity labels for themselves in the wild as they lived their lives. Users could provide 18 different labels for activities such as "Walking" and "Sitting", as well as phone location such as "Phone in Pocket" or "Phone in Hand". The participants varied in terms of the number of labels they provided. The application collects several similar channels such as screen locked, battery charging and geo-location.

Emma's visualization of the bio-rhythms of study1b participants using *ARGUS*:

Emma visualizes this dataset in ARGUS. As this study was conducted in the wild, subjects had to continuously label their smartphone-sensed data with ground-truth labels of their actual activity to facilitate supervised machine learning later. Emma believes that students' bio-rhythms affected the quality/accuracy of labels they provided and used *ARGUS* to explore this hypothesis.

- *Investigating the effectiveness of analyzing sensor rhythm values in the absence of human provided ground truth labels:* She analyzes participant 47450B. She notices quickly in the Duration View (DV) that the participant has done an inconsistent job of providing self-reported labels for their data. This means that she will have to rely on objective sensor values to make sense of this data. She clicks on the channel rhythm selector (top pane of Fig. 2A) to select "Sensors" as the channel category (T3) as these channels do not require human labeling. She notices a day where the rhythm is off (Fig. 9). She visualizes it in the EV and notices that the user was in the "Frat House" cluster (a type of on campus residence common in universities in the USA). She notices in the DV that the participant provided no labels for "Lying down" or "Sleeping". She clicks on the "Lying down" bar in the co-occurrence view and notices that the top co-occurring positive values for "Lying down" are "Sleeping", "Phone on table" and "battery" (Fig. 9B). She hovers over the co-occurring bars for "Sleeping" and "Phone on table" and notices these labels were also not provided. She hovers over the battery and notices the light yellow overlay (Fig. 9C) for "battery" but no light blue overlay for "Lying down". She notices that the screen was also locked and that the rhythm disruption for objective sensor values in this day was caused by other deviations later on in the day, which may not be that interesting (T4,G2). Overlaying and linking this multi-faceted data allowed her to dismiss this day from her concern, which may not have been possible using traditional statistical analyses.
- *Detecting erroneous labels based on unlikely co-occurrence using ARGUS*: Emma is also aware that some people may have carelessly provided labels (van Berkel et al., 2020) which would make it difficult for her to accurately determine rhythms and deviations therein. She chooses the sensors category (T3) for the rhythms and notices a user 1AACA1 who has some bio-rhythm disruptions. She clicks on the sector for September 8th and she notices in the EV that the user had their screen unlocked for some time after midnight. She clicks the "Sleeping" bar in the co-occurrence view and can see that the most commonly co-occurring channels are "Phone on table", "WiFi" and "battery" (Fig. 10). Hovering over "'Phone on table" shows that the two channels do not co-occur on this day. Further, sleeping and phone usage while the screen was unlocked is unlikely to have occurred. This view leads Emma to believe that this may have been an instance of mislabeling and there was in fact a disruption in this person's bio-rhythm (T4, G2). Calculating the rhythm score separately for various objective sensor values enabled this as humans are error prone.

6. Evaluation

Our evaluation of ARGUS was two-fold: First, we invited the same expert, a professor in psychology to interact with the final version of ARGUS and give feedback on our use cases and the feasibility of our approach in contrast to purely autonomous models. For the second part, we invited 5 graduate students (with a computer science-related background) to interact with ARGUS and provide feedback about the understandability and ease of use of the tool itself.

The expert, who helped lead the goal and task analysis, was asked to go through the same use cases as Emma. After a brief tutorial, she was asked to interact with ARGUS and not to constrain herself to only the use cases provided. She liked our use of the Z-glyph, noting that it was easy for her to discern HBR disruptions, using the EA and the MEA. She appreciated the ability to juxtapose human provided, valuable information such as academic schedules, and types of locations that gave her the ability to discern "predictable" i.e., around academic deadlines vs "unpredictable" breaks that require further contextualization. She noted how "everyone has their weird behavior at times, especially students which is not necessarily concerning overall". Presenting multiple channels of smartphone-sensed data in linked views

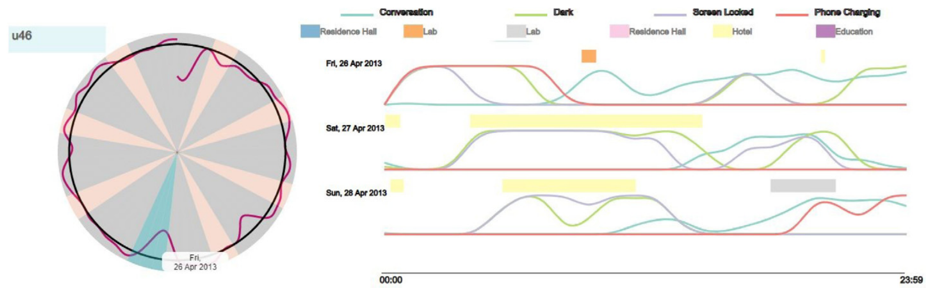

Fig. 8. There is a major disruption in bio-rhythm starting from Apr 26th, which is a Friday. Towards the end of this day, the available geo-location indicates that this person was at a hotel. The subsequent two days which are Saturday and Sunday, there are longer time periods of geo-location indicating that they are at a hotel.

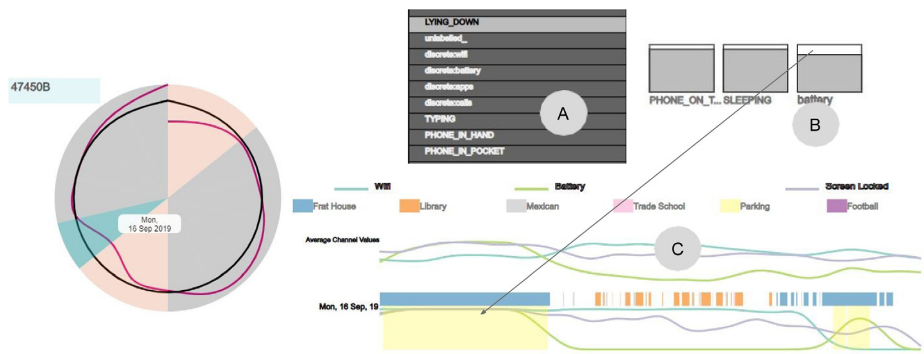

Fig. 9. Clicking the "Lying down" bar shows the most frequently co-occurring channels which are "Phone on table", "Sleeping" and "battery". Hovering over the "battery" bar shows the occurrence of that channel but also see no coincidence with "Lying down". However, the screen is locked throughout the highlighted period indicating a high probability of "Sleeping". Showing co-occurrences like these highlights the issues with such datasets and enables analysts make smarter decisions.

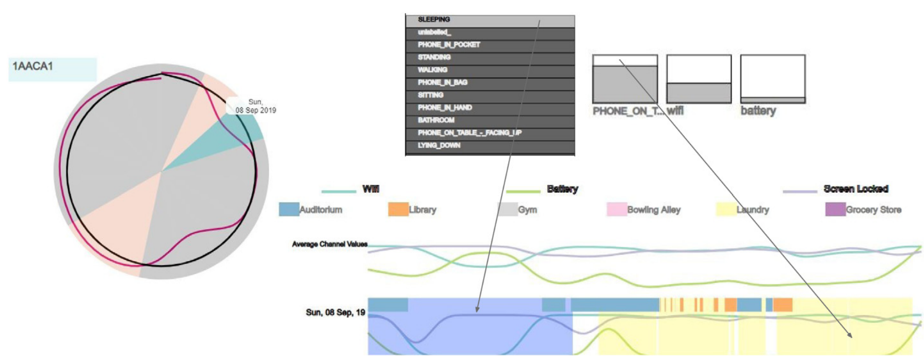

Fig. 10. Study participants may also do a poor job of self-labeling smartphone gathered data. This person usually co-labels "Phone on table" and "Sleeping" but not for this day. In addition, there is time period after the start of midnight for which the screen was unlocked that is unlikely to co-occur with "Sleeping". IVA makes these important aspects of the data human understandable.

helped her understand those breaks. Overall, she deemed our use cases to be relevant and found the ARGUS interactive visual analytics approach useful.

We invited 5 volunteers to interact with ARGUS and provide feedback. The volunteers were all graduate students, with some computational background. They were given a short tutorial about the project, datasets and ARGUS. They were then asked to interact with ARGUS and go through the same use cases that Emma did. At the end of the evaluation session, they were asked to fill out a questionnaire about the understandability and the ease of use of ARGUS' visual metaphors. We were interested in getting their feedback on the ability to complete the core tasks that the expert had summarized for us. Questions 1–5 (11) correspond to assessing the capabilities of ARGUS to achieve Tasks (T) 1–5. The questionnaire and results are presented in Figs. 11 and 12.

The evaluators gave the highest scores for Questions 1 and 2 which were meant to quantify the ability of a user to discern and gauge deviations in HBRs, using our implementation of the Z-glyph (Cao et al., 2018). This replicated findings from Cao et al. (2018) who also quantified and reported significant ease in the viewers' ability to discern deviations from the norm. For Question 3, the evaluators generally wanted a clearer explanation of "channels", along with labels in the Duration View (Fig. 2D that were indicative of units that were being presented. For instance, minutes vs. hours of being still, minutes vs. hours of being in dark, etc. which would then enable them to apply their own human intuition to better understand the data. These modifications will be added to future iterations of ARGUS. The high scores for Questions 4 and 5 show that the evaluators were able to utilize the interactive features in ARGUS to gain a clearer human-understandable view and assign *semantic* meaning to such objective smartphone data and human provided reports. Overall, the evaluators found

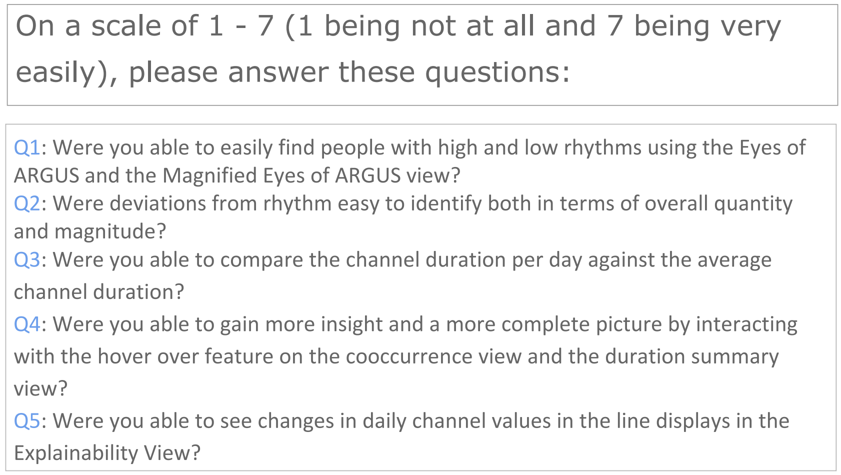
On a scale of 1 - 7 (1 being not at all and 7 being very easily), please answer these questions:

Q1: Were you able to easily find people with high and low rhythms using the Eyes of ARGUS and the Magnified Eyes of ARGUS view?
Q2: Were deviations from rhythm easy to identify both in terms of overall quantity and magnitude?
Q3: Were you able to compare the channel duration per day against the average channel duration?
Q4: Were you able to gain more insight and a more complete picture by interacting with the hover over feature on the cooccurrence view and the duration summary view?
Q5: Were you able to see changes in daily channel values in the line displays in the Explainability View?

Fig. 11. The questionnaire to evaluate ARGUS. We use a standard 7 point Likert scale (Likert, 1932) for the possible responses. Each question was designed to assess the effectiveness and ease of use of each of our visual metaphors.

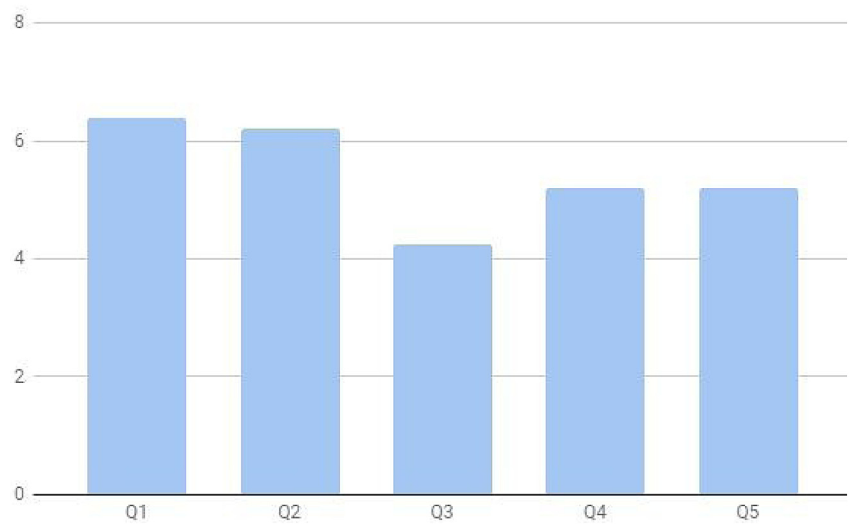

Fig. 12. Questionnaire results.

ARGUS intuitive and were generally able to understand the visual metaphors that we designed.

7. Limitations and future work

While promising, the work we have presented has several limitations including:

- *Limits of person-level monitoring:* As the number of participants increases in size, along with longer durations of time and more complex symptom labels, person-level monitoring of smartphone-sensed behavioral rhythms may become unfeasible. One approach of grouping people together for cohort analysis might be to sort people into *chronotypes* (Roenneberg, 2012) which is basically a person's sleep-wake cycle and the habit of going to bed and waking up at particular times each day. In addition, data aggregation along with longer time scales (for instance weeks instead of days) may allow analysts to accommodate longer time periods of data collection.
- *Missing human-provided ground truth activity and context labels* (as in the case of our Study1b dataset) may also become an issue (van Berkel et al., 2020) as larger deployments of studies may not be able to depend upon participants to provide accurate labels. In addition, longer and larger studies may also be subjected to more stringent privacy requirements which may limit the amount and type of information that is gathered such requiring the anonymization of geo-coordinates. In such cases, additional visual contextualization may be necessary to enable analysts to establish ground truth from partial labels. We plan will investigate using existing machine learning-based human behavior models on objective and anonymized sensor data like accelerometer and gyroscope data collected to detect and visualize users' contexts and inferred health patterns.

8. Conclusion

In this paper, we presented *ARGUS*, a visual analytics framework that allows analysts to not only identify disruptions in smartphone-gathered Human Bio-behavioral Rhythms (HBRs) but also to contextualize and explain them. To guide our designs, we conducted a detailed goal and task analysis with an expert smartphone-sensed health to understand the use cases that they would have while analyzing such data. We devised an intuitive Rhythm Deviation Score (RDS) that quantified the degree of rhythmicity in participants' bio-rhythm, which was then visualized using a glyph visual metaphor that enabled easy identification of disruptions in bio-rhythms. *ARGUS* provided additional overlays including multi-sensor channel and geo-location overlays and multiple linked visualization panes, which facilitated contextualization and reasoning about participants' bio-rhythm scores. ARGUS integrates IVA support for population-level HBR meta-analyses, easy identification of significant HBR disruptions, cross-channel exploration and contextualization of individual participant HBRs RDSs, and visualization of corresponding raw sensor values.

We provided an extensive walk-through of illustrative use cases to show how multiple linked panes provided a clearer look into the occurrences and causes of disruptions in bio-behavioral rhythms. In addition, results from evaluation sessions with experts and non-experts show that ARGUS is well-suited for presenting smartphone-sensed bio-behavioral data and can aid analysts in meaningful analysis.

CRediT authorship contribution statement

Hamid Mansoor: Writing - original draft, Writing - review & editing, Visualization, Software, Conceptualization. **Walter Gerych:** Writing - original draft, Writing - review & editing, Software, Conceptualization, Data curation. **Abdulaziz Alajaji:** Writing - original draft, Writing - review & editing, Conceptualization. **Luke Buquicchio:** Writing - original draft, Writing - review & editing, Conceptualization. **Kavin Chandrasekaran:** Writing - original draft, Writing - review & editing, Conceptualization. **Emmanuel Agu:** Writing - original draft, Writing - review & editing, Conceptualization, Supervision, Funding acquisition. **Elke Rundensteiner:** Writing - original draft, Writing - review & editing, Conceptualization, Supervision, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We would like to thank Prof. Angela Rodriguez for her insightful feedback. This material is based on research sponsored by DARPA, USA under agreement number FA8750-18-2-0077. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government.

References

Abdullah, S.S., 2020. Visual analytics of electronic health records with a focus on acute kidney injury.
Abdullah, S., Matthews, M., Murnane, E.L., Gay, G., Choudhury, T., 2014. 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. ACM, pp. 673–684.
Abdullah, S., Murnane, E.L., Matthews, M., Choudhury, T., 2017. Circadian computing: sensing, modeling, and maintaining biological rhythms. In: Mobile Health. Springer, pp. 35–58.
Ben-Zeev, D., Brian, R., Wang, R., Wang, W., Campbell, A.T., Aung, M.S., Merrill, M., Tseng, V.W., Choudhury, T., Hauser, M., et al., 2017. Crosscheck: Integrating self-report, behavioral sensing, and smartphone use to identify digital indicators of psychotic relapse. Psychiatr. Rehabil. J. 40 (3), 266.
van Berkel, N., Goncalves, J., Wac, K., Hosio, S., Cox, A.L., 2020. Human accuracy in mobile data collection.
Boukhechba, M., Chow, P., Fua, K., Teachman, B.A., Barnes, L.E., 2018a. Predicting social anxiety from global positioning system traces of college students: feasibility study. JMIR Ment. Health 5 (3), e10101.
Boukhechba, M., Daros, A.R., Fua, K., Chow, P.I., Teachman, B.A., Barnes, L.E., 2018b. Demonicsalmon: Monitoring mental health and social interactions of college students using smartphones. Smart Health 9, 192–203.
Brewer, C.A., Harrower, M., 2010. Colorbrewer 2.0: color advice for cartography. The Pennsylvania State University, http://colorbrewer2.org/ (Accessed 6 (02) (2009)).
Buck, B., Scherer, E., Brian, R., Wang, R., Wang, W., Campbell, A., Choud-hury, T., Hauser, M., Kane, J.M., Ben-Zeev, D., 2019. Relationships between smartphone social behavior and relapse in schizophrenia: a preliminary report. Schizophr. Res. 208, 167–172.
Calabrese, F., Ferrari, L., Blondel, V.D., 2015. Urban sensing using mobile phone network data: a survey of research. Acm Compu. Surv. (Csur) 47 (2), 25.
Canzian, L., Musolesi, M., 2015. 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, pp. 1293–1304.
Cao, N., Lin, Y.-R., Gotz, D., Du, F., 2018. Z-glyph: Visualizing outliers in multivariate data. Inf. Vis. 17 (1), 22–40.
Cao, N., Shi, C., Lin, S., Lu, J., Lin, Y.-R., Lin, C.-Y., 2015. Targetvue: Visual analysis of anomalous user behaviors in online communication systems. IEEE Trans. Vis. Comput. Graphics 22 (1), 280–289.
Chen, Z., Lin, M., Chen, F., Lane, N.D., Cardone, G., Wang, R., Li, T., Chen, Y., Choudhury, T., Campbell, A.T., 2013. Unobtrusive sleep monitoring using smartphones. In: Proceedings of the 7th International Conference on Pervasive Computing Technologies for Healthcare. ICST (Institute for Computer Sciences, Social-Informatics and ...), pp. 145–152.
Choe, E.K., Lee, B., Kay, M., Pratt, W., Kientz, J.A., 2015. Sleeptight: low- burden, self-monitoring technology for capturing and reflecting on sleep behaviors. In: Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing, pp. 121–132.
Ciman, M., Wac, K., 2019. Smartphones as sleep duration sensors: validation of the isensesleep algorithm. JMIR MHealth UHealth 7 (5), e11930.
Cuttone, A., 2017. Data Mining and Visualization of Large Human Behavior Data Sets (Ph.D. thesis, Ph. D. Thesis). TechnicalU-niversity of Denmark (DTU).
Dingler, T., Schmidt, A., Machulla, T., 2017. Building cognition-aware systems: A mobile toolkit for extracting time-of-day fluctuations of cognitive performance. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 1 (3), 1–15.
van den Elzen, S., Holten, D., Blaas, J., van Wijk, J.J., 2013. Reordering massive sequence views: Enabling temporal and structural analysis of dynamic networks. In: 2013 IEEE Pacific Visualization Symposium (PacificVis). IEEE, pp. 33–40.
Ester, M., Kriegel, H.-P., Sander, J., Xu, X., et al., 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. In: Kdd, Vol. 96. pp. 226–231.
Eurovis, 2020. [link]. URL https://conferences.eg.org/egev20/.
Fischer, D., Vetter, C., Roenneberg, T., 2016. A novel method to visualise and quantify circadian misalignment. Sci. Rep. 6, 38601.
Foursqaure, 0000. [link]. URL https://developer.foursquare.com/.
Garcia Caballero, H.S., Westenberg, M.A., Gebre, B., van Wijk, J.J., 2019. V-awake: A visual analytics approach for correcting sleep predictions from deep learning models. In: Computer Graphics Forum, Vol. 38. Wiley Online Library, pp. 1–12.
Gaultney, J.F., 2010. The prevalence of sleep disorders in college students: impact on academic performance. J. Amer. Coll. Health 59 (2), 91–97.
Geissmann, Q., Rodriguez, L.G., Beckwith, E.J., Gilestro, G.F., 2019. Rethomics: An r framework to analyse high-throughput behavioural data. PLoS One 14 (1), e0209331.
Gerych, W., Agu, E., Rundensteiner, E., 2019. Classifying depression in imbalanced datasets using an autoencoder-based anomaly detection approach. In: 2019 IEEE 13th International Conference on Semantic Computing (ICSC). IEEE, pp. 124–127.
Gupta, A., Heng, T., Shaw, C., Li, L., Feehan, L., 2018. Designing pervasive technology for physical activity self-management in arthritis patients. In: Proceedings of the 12th EAI International Conference on Pervasive Computing Technologies for Healthcare, pp. 1–10..
Gupta, A., Heng, T., Shaw, C., Li, L., Feehan, L., 2018. Towards developing an e-coach to support arthritis patients in maintaining a physically active lifestyle. In: Proceedings of the 12th EAI International Conference on Per- vasive Computing Technologies for Healthcare, pp. 392–395.
Gupta, A., Tong, X., Shaw, C., Li, L., Feehan, L., 2017. Fitviz: A personal informatics tool for self-management of rheumatoid arthritis. In: International Conference on Human-Computer Interaction. Springer, pp. 232–240.
Heng, T.B., Gupta, A., Shaw, C., 2018. Fitviz-ad: A non-intrusive reminder to encourage non-sedentary behaviour. Electron. Imaging 2018 (1), 332-1-332-12.

Huang, Q., Cohen, D., Komarzynski, S., Li, X.-M., Innominato, P., Lévi, F., Finkenstädt, B., 2018. Hidden markov models for monitoring circadian rhyth- micity in telemetric activity data. J. R. Soc. Interface 15 (139), 20170885.
Kakar, T., Qin, X., Rundensteiner, E.A., Harrison, L., Sahoo, S.K., De, S., 2019a. Diva: Exploration and validation of hypothesized drug-drug interactions. In: Computer Graphics Forum, Vol. 38. Wiley Online Library, pp. 95–106.
Kakar, T., Qin, X., Tapply, C.M., Spring, O., Murphy, D., Yun, D., Rundensteiner, E.A., Harrison, L., La, T., Sahoo, S.K., et al., 2019b. Designing a visual analytics system for medication error screening and detection. In: International Joint Conference on Computer Vision, Imaging and Computer Graphics. Springer, pp. 285–312.
Keim, D.A., Mansmann, F., Schneidewind, J., Thomas, J., Ziegler, H., 2008. Visual analytics: Scope and challenges. In: Visual Data Mining. Springer, pp. 76–90.
Kerr, J., Patterson, R.E., Ellis, K., Godbole, S., Johnson, E., Lanckriet, G., Staudenmayer, J., 2016. Objective assessment of physical activity: classifiers for public health. Med. Sci. Sports Exerc. 48 (5), 951.
Koven, J., Felix, C., Siadati, H., Jakobsson, M., Bertini, E., 2018. Lessons learned developing a visual analytics solution for investigative analysis of scamming activities. IEEE Trans. Vis. Comput. Graphics 25 (1), 225–234.
Kreitzman, L., Foster, R., 2011. The Rhythms of Life: The Biological Clocks that Control the Daily Lives of Every Living Thing. Profile books.
Le, T., Reeder, B., Chung, J., Thompson, H., Demiris, G., 2014. Design of smart home sensor visualizations for older adults. Technol. Health Care 22 (4), 657–666.
Liang, Z., Ploderer, B., Liu, W., Nagata, Y., Bailey, J., Kulik, L., Li, Y., 2016. Sleepexplorer: a visualization tool to make sense of correlations between personal sleep data and contextual factors. Pers. Ubiquitous Comput. 20 (6), 985–1000.
Likert, R., 1932. A technique for the measurement of attitudes. Arch. Psychol..
Lomb, N.R., 1976. Least-squares frequency analysis of unequally spaced data. Astrophys. Space Sci. 39 (2), 447–462.
Madan, A., Cebrian, M., Moturu, S., Farrahi, K., et al., 2011. Sensing the'' health state'' of a community. IEEE Pervasive Comput. 11 (4), 36–45.
Malik, S., Du, F., Monroe, M., Onukwugha, E., Plaisant, C., B, 2015. Shneider- man, Cohort comparison of event sequences with balanced integration of visual analytics and statistics. In: Proceedings of the 20th International Conference on Intelligent User Interfaces, pp. 38–49.
Mansoor, H., Gerych, W., Buquicchio, L., Alajaji, A., Chandrasekaran, K., Agu, E., Rundensteiner, E., 2020a. ARGUS: Interactive visual analytics frame- work for the discovery of disruptions in bio-behavioral rhythms. In: Kerren, A., Garth, C., Marai, G.E. (Eds.), EuroVis 2020 - Short Papers. The Eurographics Association, http://dx.doi.org/10.2312/evs.20201043.
Mansoor, H., Gerych, W., Buquicchio, L., Alajaji, A., Chandrasekaran, K., Agu, E., Rundensteiner, E., 2020b. Intosis: Interactive observation of smart- phone inferred symptoms for in-the-wild data. In: 2020 IEEE International Conference on Big Data (Big Data). pp. 4882–4891. http://dx.doi.org/10.1109/BigData50022.2020.9378147.
Mansoor, H., Gerych, W., Buquicchio, L., Chandrasekaran, K., Agu, E., Rundensteiner, E., 2019a. Comex: Identifying mislabeled human behavioral con- text data using visual analytics. In: 2019 IEEE 43rd Annual Computer Software and Applications Conference (COMPSAC), Vol. 2. http://dx.doi.org/10.1109/COMPSAC.2019.10212.
Mansoor, H., Gerych, W., Buquicchio, L., Chandrasekaran, K., Agu, E., Rundensteiner, E., 2019b. Delfi: Mislabelled human context detection using multifeature similarity linking. In: 2019 IEEE Visualization in Data Science (VDS). http://dx.doi.org/10.1109/VDS48975.2019.8973382.
Matthews, M., Abdullah, S., Murnane, E., Voida, S., Choudhury, T., Gay, G., Frank, E., 2016. Development and evaluation of a smartphone-based measure of social rhythms for bipolar disorder. Assessment 23 (4), 472–483.
Meyer, T.E., Monroe, M., Plaisant, C., Lan, R., Wongsuphasawat, K., Coster, T.S., Gold, S., Millstein, J., Shneiderman, B., 2013. Visualizing Patterns of Drug Prescriptions with Eventflow: A Pilot Study of Asthma Medications in the Military Health System. Tech. rep., OFFICE OF THE SURGEON GEN-ERAL (ARMY) FALLS CHURCH VA.
Min, J.-K., Doryab, A., Wiese, J., Amini, S., Zimmerman, J., Hong, J.I., 2014. Toss'n'turn: smartphone as sleep and sleep quality detector. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM, pp. 477–486.
Mohr, D.C., Zhang, M., Schueller, S.M., 2017. Personal sensing: understand- ing mental health using ubiquitous sensors and machine learning. Annu. Rev. Clin. Psychol. 13, 23–47.
Nguyen, P.H., Henkin, R., Chen, S., Andrienko, N., Andrienko, G., Thon-nard, O., Turkay, C., 2020. Vasabi: Hierarchical user profiles for interactive visual user behaviour analytics. IEEE Trans. Vis. Comput. Graph. 26 (1), 77–86. http://dx.doi.org/10.1109/TVCG.2019.2934609.
Nguyen, P.H., Turkay, C., Andrienko, G., Andrienko, N., Thonnard, O., Zouaoui, J., 2018. Understanding user behaviour through action sequences: from the usual to the unusual. IEEE Trans. Vis. Comput. Graph..
Ohayon, M.M., Lemoine, P., Arnaud-Briant, V., Dreyfus, M., 2002. Prevalence and consequences of sleep disorders in a shift worker population. J. Psychos. Res. 53 (1), 577–583.
Onnela, J.-P., Rauch, S.L., 2016. Harnessing smartphone-based digital phenotyping to enhance behavioral and mental health. Neuropsychopharmacology 41 (7), 1691–1696.
Payandeh, S., Chiu, E., 2019. On visualization of movements for monitoring older adults. In: Computer Graphics International Conference. Springer, pp. 316–326.
Plaisant, C., Mushlin, R., Snyder, A., Li, J., Heller, D., Shneiderman, B., 2003. Lifelines: using visualization to enhance navigation and analysis of patient records. In: The Craft of Information Visualization. Elsevier, pp. 308–312.
Polack, P.J., Sharmin, M., de Barbaro, K., Kahng, M., Chen, S.-T., Chau, D.H., 2017. Exploratory visual analytics of mobile health data: Sensemaking challenges and opportunities. In: Mobile Health. Springer, pp. 349–360.
Polack Jr, P.J., Chen, S.-T., Kahng, M., Barbaro, K.D., Basole, R., Sharmin, M., Chau, D.H., 2018. Chronodes: Interactive multifocus exploration of event sequences. ACM Trans. Interact. Intell. Syst. 8 (1), 1–21. http://dx.doi.org/10.1145/3152888.
Pu, J., Xu, P., Qu, H., Cui, W., Liu, S., Ni, L., 2011. Visual analysis of people's mobility pattern from mobile phone data. In: Proceedings of the 2011 Visual Information Communication-International Symposium. ACM, p. 13.
Rabbi, M., Ali, S., Choudhury, T., Berke, E., 2011. Passive and in-situ assessment of mental and physical well-being using mobile sensors. In: Proceedings of the 13th International Conference on Ubiquitous Computing. ACM, pp. 385–394.
Rashid, H., Mendu, S., Daniel, K.E., Beltzer, M.L., Teachman, B.A., Boukhechba, M., Barnes, L.E., 2020. Predicting subjective measures of social anx- iety from sparsely collected mobile sensor data. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 4 (3), 1–24.
Resnick, P., Carton, S., Park, S., Shen, Y., Zeffer, N., 2014. Rumorlens: A system for analyzing the impact of rumors and corrections in social media. In: Proc. Computational Journalism Conference, Vol. 5.
Roenneberg, T., 2012. Internal Time: Chronotypes, Social Jet Lag, and Why You're So Tired. Harvard University Press.
Roenneberg, T., Allebrandt, K.V., Merrow, M., Vetter, C., 2012. Social jetlag and obesity. Curr. Biol. 22 (10), 939–943.
Saeb, S., Cybulski, T.R., Schueller, S.M., Kording, K.P., Mohr, D.C., 2017. Scalable passive sleep monitoring using mobile phones: opportunities and obstacles. J. Med. Internet Res. 19 (4), e118.
Saeb, S., Lattie, E.G., Schueller, S.M., Kording, K.P., Mohr, D.C., 2016. The relationship between mobile phone location sensor data and depressive symptom severity. PeerJ 4, e2537.
Scargle, J.D., 1982. Studies in astronomical time series analysis. ii-statistical aspects of spectral analysis of unevenly spaced data. Astrophys. J. 263, 835–853.
Senaratne, H., Mueller, M., Behrisch, M., Lalanne, F., Bustos-Jiménez, J., Schneidewind, J., Keim, D., Schreck, T., 2017. Urban mobility analysis with mobile network data: a visual analytics approach. IEEE Trans. Intell. Transp. Syst. 19 (5), 1537–1546.
Shen, Z., Ma, K.-L., 2008. Mobivis: A visualization system for exploring mobile data. In: 2008 IEEE Pacific Visualization Symposium. IEEE, pp. 175–182.
Vaizman, Y., Ellis, K., Lanckriet, G., Weibel, N., 2018a. Extrasensory app: Data collection in-the-wild with rich user interface to self-report behavior. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. ACM, p. 554.
Vaizman, Y., Weibel, N., Lanckriet, G., 2018b. Context recognition in-the-wild: Unified model for multi-modal sensors and multi-label classification. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 1 (4), 168.
Vetter, C., 2020. Circadian disruption: What do we actually mean?. Eur. J. Neurosci..
Walker, W.H., Walton, J.C., DeVries, A.C., Nelson, R.J., 2020. Circadian rhythm disruption and mental health. Transl. Psychiatry 10 (1), 1–13.
Wang, R., Chen, F., Chen, Z., Li, T., Harari, G., Tignor, S., Zhou, X., Ben-Zeev, D., Campbell, A.T., 2014. Studentlife: assessing mental health, academic performance and behavioral trends of college students using smartphones. In: Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing. ACM, pp. 3–14.
Wang, W., Harari, G.M., Wang, R., Müller, S.R., Mirjafari, S., Masaba, K., Campbell, A.T., 2018. Sensing behavioral change over time: Using within-person variability features from mobile sensing to predict personality traits. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 2 (3), 141.
Xu, X., Chikersal, P., Doryab, A., Villalba, D.K., Dutcher, J.M., Tum-minia, M.J., Althoff, T., Cohen, S., Creswell, K.G., Creswell, J.D., et al., 2019. Leveraging routine behavior and contextually-filtered features for depression detection among college students. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 3 (3), 1–33.
Yan, R., Liu, X., Dutcher, J., Tumminia, M., Villalba, D., Cohen, S., Creswell, D., Creswell, K., Mankoff, J., Dey, A., et al., 2020. Corhythmo: A compu- tational framework for modeling biobehavioral rhythms from mobile and wearable data streams, bioRxiv.
Zhang, Y., Chanana, K., Dunne, C., 2018. Idmvis: Temporal event sequence visualization for type 1 diabetes treatment decision support. IEEE Trans. Vis. Comput. Graphics 25 (1), 512–522.