



**University of  
Zurich**<sup>UZH</sup>

**Zurich Open Repository and  
Archive**

University of Zurich  
University Library  
Strickhofstrasse 39  
CH-8057 Zurich  
[www.zora.uzh.ch](http://www.zora.uzh.ch)

---

Year: 2021

---

## **Observing and predicting knowledge worker stress, focus and awakeness in the wild**

Soto, Mauricio ; Satterfield, Chris ; Fritz, Thomas ; Murphy, Gail C ; Shepherd, David C ; Kraft, Nicholas

**Abstract:** Knowledge workers face many challenges in the workplace: work is fragmented, disruptions are constant, tasks are complex, and work hours can be long. These challenges can affect knowledge workers' stress, focus and awakeness, and in turn their interaction with the digital environment, the quality of work performed and their productivity in general. We report on a field study with 14 knowledge workers over an eight-week period in which we investigated, using experience sampling, how the workers experience stress and awakeness over time. During this field study, we also collected biometric data including heart- and skin-related measures, which we then used to investigate if it is possible to predict stress, focus and awakeness, in the moment. We observed and report on various trends in knowledge worker stress and awakeness levels over several weeks, finding that people tend to have certain "baseline" levels for these aspects. Moreover, we found that days with high levels of stress tend to cluster, similarly as the days with low awakeness. We further show that machine learning models can be built from the data of a single minimally invasive device to predict stress, focus, and awakeness. Overall, we found that our models were capable of large improvements in precision and recall in comparison to a random classifier for stress (25.9% increase over random for precision, 4.2% for recall) and awakeness (52.4% increase in precision, 40.8% in recall). The abstract concept of focus proved to be the hardest to predict (26.0% increase in precision, 27.8% decrease in recall).

DOI: <https://doi.org/10.1016/j.ijhcs.2020.102560>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-216315>

Journal Article

Accepted Version

Originally published at:

Soto, Mauricio; Satterfield, Chris; Fritz, Thomas; Murphy, Gail C; Shepherd, David C; Kraft, Nicholas (2021). Observing and predicting knowledge worker stress, focus and awakeness in the wild. *International Journal of Human-Computer Studies*, 146:102560.

DOI: <https://doi.org/10.1016/j.ijhcs.2020.102560>

# Observing and Predicting Knowledge Worker Stress, Focus and Awakeness in the Wild

Mauricio Soto\*

*Hitachi ABB Power Grids, Raleigh, North Carolina*

Chris Satterfield\*

*The University of British Columbia, Vancouver, British Columbia*

Thomas Fritz\*

*University of Zurich, Zurich, Switzerland*

Gail C. Murphy\*

*The University of British Columbia, Vancouver, British Columbia*

David Shepherd\*

*Virginia Commonwealth University, Richmond, Virginia*

Nicholas Kraft\*

*UserVoice, Raleigh, North Carolina*

---

## Abstract

Knowledge workers face many challenges in the workplace: work is fragmented, disruptions are constant, tasks are complex, and work hours can be long. These challenges can affect knowledge workers' stress, focus and awakeness, and in turn their interaction with the digital environment, the quality of work performed and their productivity in general. We report on a field study with 14 knowledge workers over an eight-week period in which we investigated, using experience sampling, how the workers experience stress and awakeness over time. During this field study, we also collected biometric data including heart- and skin-related measures, which we then used to investigate if it is possible to predict stress, focus and awakeness, in the moment. We observed and report on various trends in knowledge worker stress and awakeness levels over several weeks, finding that people tend to have certain "baseline" levels for these aspects. Moreover, we found that days with high levels of stress tend to cluster, similarly as the days with low awakeness. We further show that machine learning models can be built from the data of a single minimally invasive device to predict stress, focus, and awakeness. Overall, we found that our models were capable of large improvements in precision and recall in comparison to a random classifier for stress (25.9% increase over random for precision, 4.2% for recall) and awakeness (52.4% increase in precision, 40.8% in recall). The abstract concept of focus proved to be the hardest to predict (26.0% increase in precision, 27.8% decrease in recall).

*Keywords:* Biometrics, Psycho-Physiological Data, Stress, Awakeness, Focus, Computer Interaction, Ubiquitous Computing, Empirical Study, User Centered Design

---

## 1. Introduction

*‘The most valuable asset of a 21st-century institution (whether business or non-business) will be its knowledge workers and their productivity’* [1]. Knowledge workers constantly face challenges, such as a high work fragmentation, continuous disruptions and distractions, highly complex and demanding tasks, and long working hours [2, 3, 4]. These challenges, amongst others, can lead to stress in the workplace. Stress is an ever-growing concern as it can lead to fatigue, burnout and various other illnesses, ultimately resulting in work absences and marked productivity losses [5, 6, 7].

Given the importance of understanding stress and its relationship and effect on work, a number of studies have been conducted using a variety of methods. Minimally invasive studies have shown the benefit of sensing autonomic nervous system (ANS) activity by analyzing certain biometric (*aka.* psycho-physiological) signals of the human body, such as skin temperature [8]. These approaches can be used in long-term field studies, opening up the question of what stress looks like in the wild. In contrast, other studies have studied stress using more invasive techniques, such as measuring cortisol as a stress biomarker [9]. These studies are best suited for laboratory environments.

Our work builds on this ability to non-invasively monitor stress by performing an eight-week study in the workplace with 14 knowledge workers examining stress as well as two related human aspects: focus and awakeness (i.e., wakefulness). We collected data on these two additional aspects because stress and the challenges that contribute to stress levels can also affect sleep. A lack of sleep can affect the ability of knowledge workers to stay awake and focused at work which, in turn, can result in undesired consequences on work and productivity [10, 11, 12].

In this paper, we investigate two research questions. Our first question examines how knowledge workers experience stress and awakeness in their workplace over a long period of time. Our second question explores the possibility of using biometric data to predict whether a knowledge worker is experiencing stress, is focused, and is awake at a given moment in time. A better understanding of how knowledge workers experience stress, focus and awakeness can help inform the design of workplaces to manage and alleviate stress, and to encourage focus and awakeness. An ability to predict stress, focus and awakeness in the moment can enable the development of digital tools to support knowledge workers in their work by better managing disruptions, automatically adjusting lighting to reduce sleepiness, or by avoiding stress, amongst other possibilities.

To the best of our knowledge, this is the longest in-situ study performed in a real-world environment of knowledge workers analyzing an extensive collection of biometric signals with experience samples examining real-time prediction of stress, focus and awakeness in the workplace. Previous studies are predominantly controlled lab studies, which used only a subset of the relevant

---

\*These authors contributed equally.

*Email addresses:* mauricio.soto-gonzalez1@hitachi-powergrids.com (Mauricio Soto), cds00@cs.ubc.ca (Chris Satterfield), fritz@ifi.uzh.ch (Thomas Fritz), murphy@cs.ubc.ca (Gail C. Murphy), shepherd@vcu.edu (David Shepherd), nkraft@gmail.com (Nicholas Kraft)

biometric signals we captured, much more limited in duration, or not focused on a real-world environment [13, 14, 15, 16, 17, 18, 19, 20].

The participants in our study perform various job functions for a research and development group within a single large corporation. For this study, they wore a novel, state-of-the-art biometric armband comprised of precise sensors<sup>1</sup> which captured heart-, respiration- and skin-related measures with low invasiveness. This modality was chosen to ease longitudinal deployment in the field. We then used machine learning to create classifiers and analyze their ability to predict stress, focus and awakeness levels based on the biometric signs gathered by the sensors. Considering all of these three human aspects helps us to consider whether one aspect might be easier to detect than another. If one aspect is easier to predict, it might serve as a proxy or an indicator of the presence (or absence) of other aspects.

In our analysis, we identify several trends in day-to-day stress levels that emerged over the course of our eight-week study. The results of our analysis show that biometric signals collected from a single minimally invasive device can be used to predict stress and its related aspects accurately, with the abstract concept of focus (predictably) being the hardest to detect. Our results further determine that knowledge workers' self-reported levels of stress, focus and awakeness, and their physiological manifestation and prediction can vary substantially between individuals. The three main contributions of our work are:

- A qualitative examination of how knowledge worker stress and awakeness behaves and fluctuates in the wild based on an eight-week field study with 14 office workers.
- The creation and analysis of measures for the automatic monitoring of knowledge workers' stress, focus and awakeness in the workplace based on a machine learning model trained on biometric data and experience samples.
- A discussion on the impact of applying this research to improve the interaction of knowledge workers with their digital environment leading to an improvement in their productivity and well-being, besides a reflection on aspects that can be further improved in future studies.

## 2. Related Work

The study and prediction approach used in this paper is related to previous studies of stress, focus, and awakeness and studies using biometrics to predict them. We consider related work in each of these categories in turn.

### 2.1. Stress

Much previous work relates to identifying and mitigating stress in an office environment. Previous studies have measured stress by taking one of two possible approaches. The first approach is to measure plasma catecholamine and cortisol as stress biomarkers [9]. This approach is impractical for use over prolonged periods of time, as in our eight-week study. Further, this approach is imprecise because of the delay from the stress stimulation to the stress response, which may take from minutes to hours [22, 23].

The second approach, which we have chosen for our study, is to measure autonomic nervous system (ANS) activity by analyzing biometric signals of the human body, such as blood pressure,

---

<sup>1</sup>Biovotion Everion [21]

heartbeat, and skin temperature [8, 24, 25]. In particular, changes in heart rate variability are associated with cognitive and emotional stress [26, 27]. This second approach has been used successfully by several past studies [28, 29, 30].

Zaman et al. [31] performed a study in which they measure stress and productivity in short sessions with 63 participants using biometric sensors (thermal facial camera, wrist EDA, chest breathing sensor, and facial camera). Different from this study, our corpus is gathered from real-world office workers. Also, our study spans for eight weeks allowing us to analyze the behavior of our participants through the weeks and providing a more extensive in-depth analysis of our participants over time.

Hovsepian et al. [32] have worked to obtain a standard for continuous stress assessment. They used sensors to conduct a seven-day lab study with 26 participants, as well as a field study with 20 participants. They found that their model showed significant improvement over simple heart rate variability measurements. However, this model requires the use of a suite of invasive sensors that would be impractical for a study the length of ours.

Hernandez et al. [33] investigated the use of a pressure-sensitive keyboard and a capacitive mouse as non-intrusive means for measuring computer users' stress levels. They found participants' exhibited significantly increased typing pressure and mouse contact when in stressful conditions. McDuff et al. [26] experimented with a camera to measure photoplethysmographic signals indicative of cognitive stress. Vizer et al. [34] used keystroke and linguistic features to automatically measure stress levels in response to cognitive and physical stress conditions. All of these studies were performed in a controlled lab setting over a short duration and the results have yet to be replicated in the field.

Kocielnik et al. [35] developed a framework for unobtrusive and continuous measurement of stress in real life conditions. They equipped university staff members with a wristband sensor and combined this data with information from the participants' calendars over the course of four weeks. They observed that the data they collected reflected well the participants' perceptions of their own stress levels. However, this work focused mostly on providing retrospective information to users so they can work to improve their own stress balance. They did not attempt to make observations about the big picture of participants' stress profiles or make predictions in real time. Also using a wristband, Hernandez et al. [36] studied the stress level during calls in a call center over a seven day period. They collected skin conductance measures and examined the interpersonal variability of reporting stress.

Our study stands out from these works by nature of its length and focus on knowledge worker stress in everyday office life, using unobtrusive measures. To the best of our knowledge there is no longitudinal study that attempts to explain and predict knowledge worker stress that comes close to the length of our own study. The eight week duration gives us authority to speak on the nature of day-to-day fluctuations in knowledge worker stress levels.

Rather than trying to measure or predict stress, several studies have induced stress and examined its effect on work performance, motivation and others. For instance, Sarsenbayeva et al. [37] induced stress and studied how it affects mobile interaction, showing that it can reduce completion time and accuracy during target acquisition tasks. Evans and Johnson [38] investigated the correlation between noise in the workplace and stress levels. They found that workers exposed to open-office noise showed aftereffects that indicate motivational deficits but found no difference in cortisol levels. The population for their experiment comprised 40 female clerical workers, who were randomly assigned to a control condition or to three-hour low-intensity noise room designed to simulate typical open-office noise levels.

Stress in the workplace and its effects on service providers has also been analyzed in call

centers [36] and in the context of the perceived imbalance between resources and demands [39]. These studies considered several factors such as personality traits, career-related goals and attitudes, as well as life outside of work, and examined their correlation with stress levels and burnout.

Finally, there is also some work that already used proprietary stress measures for other types of classification. For example, Mirjafari et al. [40] used a proprietary stress measure by Garmin together with other data collected through the sensing of mobile devices and built machine learning models to differentiate between low and high performing workers.

## 2.2. Focus

Focus refers to the allocation of limited cognitive processing resources [41]. Mark et al. [12] studied attentional states, including focus, for workplace activities by analyzing the digital activity of 32 information workers in situ for 5 days. They found that boredom is highest in the early afternoon and focus peaks in the middle of the afternoon. They also found that doing work that requires focus correlates with stress, while rote work correlates with happiness.

Interruptions in the office are a common barriers keeping workers from sustaining focus on their work related activities, particularly when the interruptions occur at inopportune moments. Such interruptions may include emails, alerts, or interactions with co-workers [2, 42, 43]. Interruptions in inopportune times can have negative effects that range from higher error rate and lower overall performance to an increase in stress and frustration [44, 45, 3]. External interruptions may cause workers to enter a “chain of distraction” [43]. This chain is composed by stages of preparation, diversion, resumption and recovery that result in time away from an ongoing task. Since interruptions can have a large impact on the focus and productivity of office workers, several studies have examined the prediction of interruptibility—the availability for interruptions—using a variety of features, including computer interaction and biometrics [46, 47, 48, 49, 16, 50]. Most of these studies were again conducted for small and controlled tasks over shorter periods of time.

Other studied constructs that relate to focus in the workplace include cognitive absorption, cognitive engagement, flow, and mindfulness. Cognitive absorption describes periods of time in which a person experiences total immersion in an activity. This state is also accompanied by a sense of deep enjoyment, a feeling of control, curiosity, and not realizing the passing of time. It has been associated with ease of use and perceived usefulness of information technology [51]. Cognitive engagement is described [52] as a period of strong focus in an activity without the feeling of a sense of control of the situation. Flow [53], and mindfulness [54, 55] are psychological states that describe periods of prolonged attention and total immersion in an activity. Flow occurs when a person is focused on an activity that requires high challenge and high use of the person’s skills, whereas mindfulness is characterized by being aware of fine detail, affording the capacity to discover and manage unexpected events.

## 2.3. Awakeness

Sleepiness (lack of awakeness) and its associated risk of serious injury to passengers has been studied in the context of automobile accidents [10, 56]. These studies show a strong association between the level of acute driver sleepiness and the risk of injury crash. Connor et al. [10] conducted a population-based case study using the Stanford sleepiness scale, which is similar to a seven-point Likert scale and describes seven different levels of sleepiness from “Could not stay awake, sleep onset was imminent” (1) to “Felt active, wide awake” (7). Nordbakke and Sagberg [56] show that drivers are well aware of various factors influencing the risk of falling

asleep while driving. Drivers also have good knowledge of the most effective measures to prevent falling asleep at the wheel. However, most of drivers continue driving even when recognizing sleepiness signals, due to the desire to arrive at a reasonable time, the length of the drive, or pre-planned commitments.

#### *2.4. Biometrics*

Our study investigates the prediction of stress, focus, and awokeness using two different types of measurements, and biometric signals over eight weeks in a real-life office setting. Existing work [13, 14, 15, 16, 17, 18, 19, 20]—some of it already mentioned above— analyzes a broad array of biometric signals and correlates them with individual’s cognitive states and processes. For example, Zuger et al. [50] used biometrics to sense interruptibility in the office [50]. Biometric signals have also been studied in the context of technology users. For example, Parnin [18] analyzes electromyography to measure sub-vocal utterances, and how these might be correlated with the programmer’s perceived difficulty of programming tasks. Similarly, biometrics have been used to measure code difficulty by using biometric sensors [57] and using Near Infrared Spectroscopy to measure developer’s cerebral blood flow [19].

Eye tracking technology [58, 59, 60] and brain activity [61, 62] have been used in previous studies to examine different tasks in an office environment. Eye tracking has been used to analyze memory load and processing load by inspecting task-evoked pupillary response and pupil size [63]. Similar studies have shown high correlation between pupil size and mental workload of subtasks [63] and cognitive load [64]. Brain activity has been associated with different mental states [65] by analyzing specific frequency bands (alpha, beta, gamma, delta, and theta) using electroencephalography (EEG). The increase or decrease of some of these frequencies is correlated with attentional demand and working memory load [66, 67]. In contrast to studies that use eye tracking or EEG, we focused on a less invasive technology that can be applied in a real world scenario. Similarly, we considered adding additional metrics such as the Depression, Anxiety and Stress Scale (DASS) and the Perceived Stress Scale (PSS) [68] but opted for the ones presented in the paper after several discussions with experts in the area, and after piloting surveys to maximize participant compliance.

### **3. Field Study**

We conducted an eight-week field study with 14 participants using experience sampling and biometric sensors to investigate how knowledge workers experience stress and awokeness over time and the feasibility of predicting stress, focus, and awokeness using biometric signals.

#### *3.1. Participants*

We recruited a group of 14 professionals via personal contacts from a large power and automation company. The group is diverse in terms of age, work experience, gender, and work responsibilities. All participants work primarily in an office environment, though half of the participants spend at least 10% (and up to 50%) of their time in a laboratory environment. Office workers are a population that generalizes to a variety of contexts, including part-time laboratory workers, and guarantees that our participants have varying work patterns that include different levels of computer usage, as well as different levels of activity in both individual and collaborative tasks.

From the 14 participants, 11 are male and 3 are female. The average participant age is 40, with 5 in the age range 25-34, 7 in the age range 35-44, and 2 in the age range 55+. The average number of years of professional experience of the participants is 12, with 2 having less than 5 years, 10 having 5-15 years, and 2 having more than 25 years. All participants work for a research organization within the company, but their job functions span line management, laboratory science, scientific research, technology evaluation, and software development.

### 3.2. Procedure

We performed a study over the course of eight weeks. This study includes the collection of the participant's biometric data and self-reporting surveys during their work hours. We also collected computer interaction data, but given privacy concerns of some participants chose not to use this data (see Section 3.3.2). We informed participants about the study purpose and procedure, handed out biometric sensors, introduced and explained the self-reporting to the participants, handed out consent forms and ensured informed consent from each participant for the study.

After the initial setup, participants were asked to fill out three surveys each day for the following eight weeks (see Section 3.3.3) and wear the biometric sensor during their work hours. At the end of the eight weeks, we collected the biometric sensors and performed a short follow-up interview on the study and the participants' experiences.

Keeping participants engaged in an eight-week study can be difficult. In the second month of the study, largely to incentivize participants to continue, we offered participants two one-hour Tai Chi classes per week (each in the middle of the day) and asked them to attend one class a week. The choice of Tai Chi provided a link to mindfulness, a topic of growing interest in communities. Consultations we held with researchers in psychiatry and psychology suggested Tai Chi as a good link to mindfulness and an intervention that might alleviate stress. By recording attendance of participants at the Tai Chi sessions, we are able to analyze the impact of the sessions on the results. As a brief summary, the impact of the sessions was minimal; we present an analysis of the effects in Section 6.

### 3.3. Data Collection

We collected three datasets from each participant as described below:

#### 3.3.1. Biometric Sensors

Figure 1 illustrates Biovotion's Everion, which we used to track the biometric signals of the study participants. The Everion is worn on the upper arm and provides continuous monitoring of certain biometric measurements.<sup>2</sup> Previous studies [50, 13, 14, 15, 16, 17] have used similar devices [69, 70, 71] to capture psycho-physiological and biometric measurements for shorter periods of time or capturing a smaller number of measurements. A comparative study has shown that the Everion can be used as a valid proxy for HRV metrics for knowledge workers [72].

Table 1 lists the biometrics measurements that we collected using the Everion. Each measurement is collected once per second, and each recorded observation has an associated timestamp and quality rating. Data collected by the Everion is uploaded to a server, from which we downloaded the data for use in our study. We chose these biometric measurements based on previous research which indicates their potential (see references in Table 1) and the availability and feasibility of collecting these measurements with low invasiveness over a long duration.

---

<sup>2</sup><https://biovotion.zendesk.com/hc/en-us/articles/213613165>





Figure 1: We used Biovotion’s Everion to collect biometric measurements from the participants

Biometric Measurement	Units of Measure
<b>Physical Activity</b>	[73, 74]
Intensity of motion	(No unit)
Energy Expenditure	Calories per second (cal/s)
Step counter	Steps
<b>Heart</b>	[75, 14, 76, 77]
Heart rate	Beats per Minute (bpm)
Blood pulse wave	(No Unit)
Heart rate variability: RMSSD	Milliseconds (ms)
Blood oxygenation	Percent (%)
Blood perfusion	(No unit)
<b>Skin</b>	[14, 77]
Galvanic skin response	kOhm
Skin temperature	Degrees Celsius (°C)
<b>Respiration</b>	[76, 14, 77, 78]
Respiratory rate	Breaths per Minute (bpm)

Table 1: Biometric measurements captured by the Everion, organized by category and with references to previous works using similar data. (*RMSSD denotes the root mean square of successive heartbeat interval differences*)

### 3.3.2. Computer Interaction Data

To gain a better understanding of our participants day-to-day work activities, we asked participants to install an open source computer interaction monitor [79]. The monitor ran in the background on participants’ computers and tracked the active windows, as well as the keyboard and mouse activity. Four of our participants opted out of this part of the study for privacy reasons (S6, S10, S12, and S13).

### 3.3.3. Surveys

Following guidelines from previous studies [80, 81, 82] and following the preferences of extensive user piloting, we sent via text message a survey request to each participant two times per workday. Pilot participants preferred the usage of text message, in part, due to them being

accessible and noticeable anywhere in the office.

We sent the first request at a random time between 9am and 11am and sent the second request at a random time between 1pm and 3pm. We randomized the request times to avoid either establishing or observing a standard behavioral pattern. That is, we did not want the participants to plan for the arrival of the survey request at a set time, and we did not want the survey request to overlap with a set daily behavior (e.g., coffee break every day at 2:30pm). Similarly, we avoided using tools which allow for too much freedom in response time [83] since this would discourage participation in stressed time frames and would bias the corpus. The same survey was sent each time consisting of the following three questions:

1. How awake are you right now?
2. How stressed do you feel right now?
3. How focused on work are you right now?

We used the phrase “right now” to capture each aspect in the moment (so as to permit later prediction of each aspect based on biometric data). The wording of the questions is based on a previous survey of individuals in an organizational context [84]. The use of awakesness (rather than sleepiness) in Question 1 is inspired by previous work [85] and to some extent also captures the “arousal” aspect of the affective space [86].

One last survey was sent at the end of the day at 4:25pm, which asked the four different questions detailed below:

1. How awake have you been today?
2. How stressed did you feel today?
3. How productive have you been today?
4. How do you feel about your workday?

Following guidelines from similar previous studies [87, 88], we asked the participants to respond to each question using a 5-point Likert scale ranging from 1 (not at all awake/stressed/focused) to 5 (extremely awake/stressed/focused). Each participant response, as stored by Survey Gizmo, comprised the date, the time at which the response was initiated, the time at which the survey was submitted, the unique identifier for the participant, and the responses submitted by the participant.

We did not ask our participants about their focus levels in the end of the day survey, as focus is more of an in-the-moment aspect than stress and awakesness.

#### **4. Observed Trends Over Time in Stress and Awakesness Levels**

To gain insights into how knowledge workers experience stress and awakesness over an extended period of work, we examined the end of day survey responses collected from each participant to see if any identifiable trends emerged. As we noted in the last section, we did not ask participants about focus in the end of day surveys as focus is an aspect relevant at a particular moment in time rather than an aspect for an extended period of work. We used data from 13 of the 14 participants - we excluded one participant from this analysis as they experienced atypical stress levels in the latter half of the study due to factors outside of our control.

##### *4.1. Stress Levels*

Overall, we identified three prominent characteristics in the stress levels.

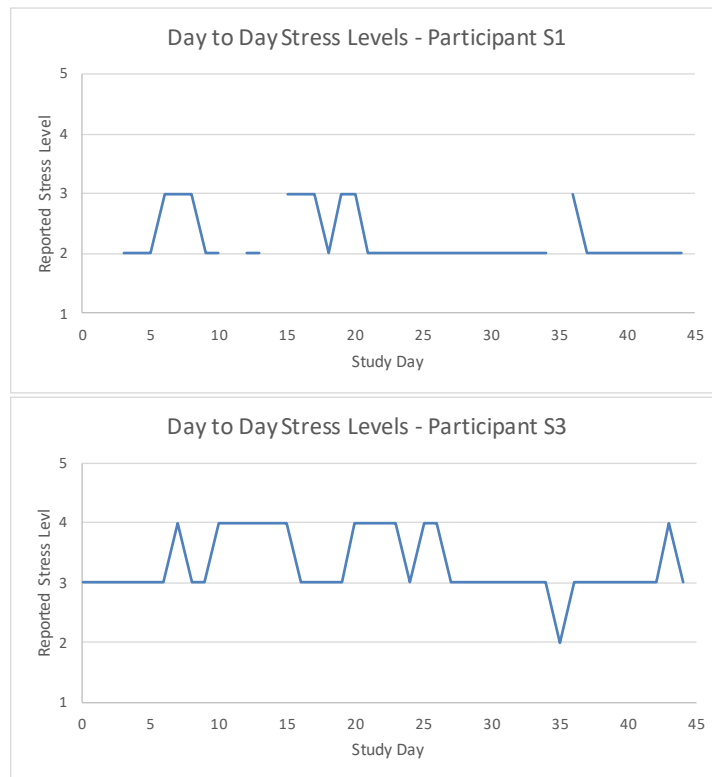


Figure 2: The day-to-day stress levels reported by two participants (S1 and S3) are shown. The y-axis represents values on the 5-point Likert scale we asked participants to respond with ranging from 1/Not at all stressed to 5/Extremely stressed. The x-axis represents the day of the study on which the value was recorded (from 0-45). Some gaps are present in the chart for S1 as they did not report their stress level on those days.

#### 4.1.1. Baseline Stress Levels

Common amongst all participants was a trend to select one stress rating far more frequently than any other. We will refer to this value as the participant's baseline stress level. All but one participant reported their perceived stress level for the day as their baseline stress level more than 50% of the time. In total, the baseline values made up 65% of the reported values collected from participants. Interestingly, while participants sometimes saw periods of sustained increases in stress, lasting as many as 6 consecutive workdays in the most extreme case, participants would always return to their baseline stress level at some point.

The baseline stress level varied significantly between participants. Seven participants (54%) reported feeling average stress levels most frequently (rating 3 on our scale), while five (38%) reported feeling little stress (rating 2) and one (8%) reported feeling no stress at all (rating 1). Figure 2 illustrates these points using data from two participants, showing the tendency to report and return to baseline stress levels, as well as a distinct difference in baseline stress level (rating 2 for S1 vs rating 3 for S3). Gaps in the chart for S1 represent days for which the participant did not report their stress level.

Variable	Stress			Awakeness		
	F.E.E.	p-value	C.I. (95%)	F.E.E.	p-value	C.I. (95%)
Monday	-0.033	0.745	(-0.234, 0.167)	-0.058	0.497	(-0.227, 0.110)
Tuesday	0.028	0.773	(-0.164, 0.221)	0.036	0.662	(-0.126, 0.198)
Wednesday	0.163	0.100	(-0.031, 0.357)	-0.018	0.830	(-0.181, 0.145)
Thursday	0.022	0.821	(-0.169, 0.213)	0.076	0.927	(-0.085, 0.238)
Friday	0.082	0.473	(-0.142, 0.306)	-0.178	0.025	(-0.334, -0.022)
Month End	0.044	0.687	(-0.171, 0.259)	-0.082	0.374	(-0.263, 0.099)

Table 2: Fixed effects estimates (F.E.E.), 95% confidence intervals (C.I.) and associated p-values for the explanatory variables (day of week and proximity to month end) that we examined in our linear mixed model analysis.

#### 4.1.2. Stressful Days Tend to Cluster

Accounting for the variance between participants perceived stress baselines, we consider a stressful day to be one that represents a deviation of one or more stress levels above the participant's baseline. Of the 93 stressful days we observed in total, we found that 39 (41%) of these days occurred in groupings of two or more consecutive stressful workdays. The most common size of these groups was two workdays, while the largest group we observed was six workdays. The day after a stressful day is much more likely to be a stressful day as compared to any other day with a 0.55 average increase over baseline, compared to 0.02 average increase over baseline.

#### 4.1.3. Extreme Changes in Stress Levels are Rare

After accounting for each participant's perceived stress baseline, we examined the frequency of deviations from the baseline. We found that participants were far more likely to report a stress level that was within one point of their baseline, than to report a stress level 2 or more points away. These extreme deviations represented only 15% of all reported values, which differed from the participant's baseline. As well, the majority (78%) of these deviations came from just two participants. This suggests that some people may be less resilient to the stress of the workplace than others. For most participants, extremely stressful days were few and far between.

#### 4.2. Awakeness

We applied the same analyses described above to the self-reported awakeness levels of our participants. Compared with the reported stress levels, we observed the same trend of participants reporting one baseline far more commonly than any other. We did not find there to be a significant correlation between reported stress and awakeness levels. Overall, participant's awakeness levels fluctuated significantly less than their stress levels (73.1% of reports were at the baseline level, compared to 65.1% for stress,  $p < 0.05$ ). Participants were unlikely to experience days with heightened (above baseline) awakeness. Such days made up only 6.0% of the total observed workdays across all participant. Most of these days came from one participant, S2, who reported 15 heightened awakeness days compared to the next highest, S13 with four. Large deviations (>1 point deviation) from each participants baseline awakeness levels were extremely uncommon, accounting for only nine (2.1%) of our total observations. Similarly to what we observed with high stress days, low awakeness days frequently came in clusters of two or three days in a row. Given that the immediately preceding day was a low awakeness day; a given day was 228.2% more likely to be a low awakeness day than when considering any day at random ( $p < 0.0001$ ).

### 4.3. Explaining Fluctuations

In an attempt to explain some of the fluctuations in stress and awakesness that our participants were experiencing, we created a linear mixed model with the self-reported daily stress and awakesness levels as dependent variables and the participants as random effects. We experimented with day of the week and proximity to beginning or end of month as possible explanatory variables. Ultimately, the analysis showed that none of the variables that we examined had a significant explanatory power with respect to our participants perceived stress levels. For awakesness, we found that there was a small (fixed effects estimate: -0.178) yet significant decrease in awakesness levels on Fridays in particular. Table 2 shows some of the detailed results of these analyses. These results show the difficulty of explaining a person's stress and awakesness via simple measures, and point to the need for additional instrumentation and data collection if we are to successfully understand and make predictions about these human aspects in the workplace.

## 5. Predicting Stress, Focus and Awakesness in the Moment

To investigate whether stress, focus and awakesness can be predicted in the moment based on biometric measures, we investigated classifiers trained for each individual and across all participants. We report on the effectiveness of these classifiers and the features that are important in predicting stress, focus and awakesness.

### 5.1. Data Preparation

In a machine learning context, data preparation and utilization is an essential part of the proposed solution. To prepare the collected data for use in training and testing our proposed machine learning models, we performed the following steps.

#### 5.1.1. Data Cleaning

All data recorded by the Everion is associated with a quality score ranging from 0-1 that is calculated using proprietary methods. In accordance with the recommendations of Biovotion, to prevent our results from being effected by erroneous data we set a quality threshold of 0.5 and discarded any data gathered which had a quality rating below this threshold.

#### 5.1.2. Data Linking

We linked the collected biometric data and survey responses for each participant. Linking the data is necessary to construct training and test datasets for use in creating and evaluating machine learning models.

To link the data, we look back one hour from the start time of each survey response for available biometric data. For example, if a participant started a survey response at 11:05am, we look for biometric data from between 10:05am to 11:05am. If no biometric data was recorded in the hour time window, we exclude the survey response from the dataset. Otherwise, we consider the survey response to have associated biometric data.

There are several reasons for a survey response to lack associated biometric data:

- The participant was not wearing the Everion in the hour before beginning the survey.
- The Everion was not recording data in the hour before the participant began the survey (e.g., due to low battery).

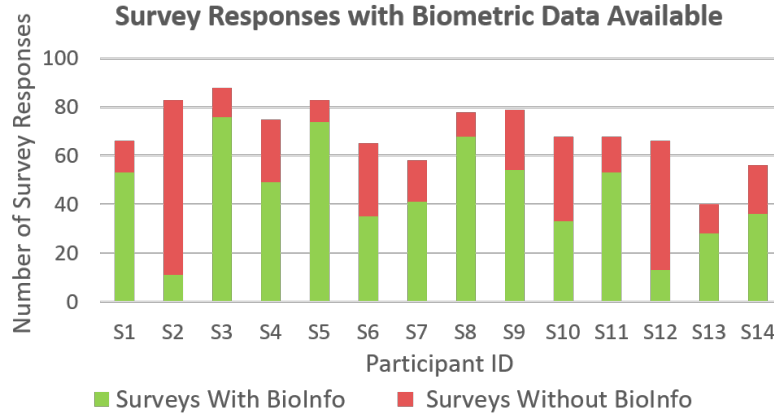


Figure 3: The figure shows the biometric data available per each participant. Green sections represent survey responses with biometric data available, red sections responses with no biometric data available

- Biometric data was not being uploaded successfully to the server.

Figure 3 illustrates the number of survey responses with associated biometric data for each study participant. The total number of responses per participant is affected by their response rate and by the number days out-of-office (e.g., vacations, holidays, etc.). Participant S2 and S12 have particularly low numbers of usable survey responses. In each of these cases, the issue related to biometric data not being uploaded successfully to the server.

### 5.1.3. Feature Extraction

We extracted features from the biometric data to provide as input to machine learning models. Previous studies [89, 16] identify time windows as an important factor that impacts the prediction accuracy of a classifier. We considered many time windows from the literature on biometric analysis [50], ranging from 10 seconds to 3 hours. Specifically, we considered the following time windows: *10sec, 20sec, 30sec, 45sec, 1min, 2min, 3min, 5min, 7.5min, 10min, 20min, 30min, 45min, 1hour, 2hour, 3hour*.

From the start time of each survey response, we look back the amount of time that corresponds to each time window and we create features for all of the biometric data available in that time window. For example, if a participant started a survey response at 11:05am, for the 30min time window, we create features using all of the available biometric data from 10:35am to 11:05am. If there is a large portion ( $\geq 50\%$ ) of data missing (either because of a recording issue or because of low quality data) from the time window considered, then the time window is marked as missing. In this case, features are imputed based on the mean of other samples of the same feature for that participant. This is an effective and commonly used technique, which is preferable to the alternative of deletion as it preserves our already small sample size [90]. For each time window, we calculate 10 statistical measurements from the biometric data to create 10 distinct features. Specifically, the 10 statistical measurements are: mean, standard deviation, variance, median, 25<sup>th</sup> percentile, 75<sup>th</sup> percentile, interquartile range, maximum, minimum, and range. Thus, for each survey response, we generate a large number of corresponding features based on three factors: biometric measurement, time window, and statistical measurement. In addition to these biometric features, we also considered the time of day in which the questions were asked. These features

		Distributions		
Participant	# Responses	Stress	Focus	Awakeness
S1	52			
S2	10			
S3	76			
S4	48			
S5	74			
S6	34			
S7	41			
S8	68			
S9	54			
S10	33			
S11	53			
S12	13			
S13	27			
S14	36			
All	619			

Table 3: The distribution of the responses of each participant to the three questions asked during the day are shown. Each bar in the histograms represent one of the five values on the 5-point Likert scale we asked participants to respond with, where the far left side of the histograms are 1/Not at all, and the far right sides are 5/Extremely

are created to predict the responses described by the ground truth. To attempt to account for inter-participant differences, we normalized all features on a per-participant basis.

#### 5.1.4. Response Transformations

Table 3 illustrates the distribution of responses from each participant for each of the three survey questions (listed in Section 3.3.3). The figure shows that there is a notable imbalance in the distribution of the self-reported responses provided by the participants. Most participants did not use all five points of the five-point Likert scale in their responses, and the distributions tend to skew toward one side or the other, depending on the question. Given this distribution and based on our earlier observation that the participants tended to adhere to a baseline reporting level for stress and awakensness, we elected to simplify the problem from five classes to two. This transformation enables us to more easily represent patterns in the data, such as when a participant fluctuates from a normal to high stress level. To perform this transformation, we began by calculating the median response value for each participant and each question. We classified each response below the median as 0 (‘negative’) and each response above the median as 1 (‘positive’). The distribution for the stress question skewed left, so we included the median values in the ‘positive’ class (i.e. ‘stressed’), while the distributions for focus and awakensness skewed right, so we included those median values in the ‘negative’ class (i.e. ‘not focused’, ‘not awake’).

With this method, we transformed the survey data into a two-point scale, representing negative or positive responses for each of the three human aspects of interests (e.g., not stressed or stressed).

#### 5.1.5. Oversampling

Even after binarizing the responses as described in the previous section, we found the distribution of responses was still quite imbalanced for many of our participants. This can be seen in the distribution columns in Table 4. To mitigate this effect, we applied random oversampling to our training sets, which artificially rebalances the dataset by creating randomly replicated data in

the minority class. This technique has commonly been used in previous studies on unbalanced datasets [91, 92].

### 5.2. Selecting a Classifier Algorithm

Many different algorithms can be used to build a classifier. To select an algorithm, we compared multiple classifiers using the popular machine learning library scikit-learn [93], evaluating each one by using leave-one-out cross validation. Our analysis showed that random forest outperforms all other classifiers, including Naïve Bayes, decision trees, support vector machine, and a multilayer perceptron neural network. For the remainder of this paper, we refer to a random forest classifier.

Participant	Stress				Focus				Awakeness			
	Accuracy	Precision of class 'stressed'	Recall of class 'stressed'	Distribution	Accuracy	Precision of class 'not focused'	Recall of class 'not focused'	Distribution	Accuracy	Precision of class 'not awake'	Recall of class 'not awake'	Distribution
S1	0.750	0.111	0.166	9/52	0.609	0.193	0.143	14/52	0.628	0.396	0.396	16/52
S2	0.866	0.000	0.000	1/10	0.133	0.063	0.083	4/10	0.367	0.000	0.000	4/10
S3	0.842	0.500	0.222	12/76	0.644	0.476	0.385	26/76	0.965	0.000	0.000	2/76
S4	0.506	0.460	0.440	22/48	0.861	0.000	0.000	4/48	0.701	0.617	0.537	18/48
S5	0.919	0.200	0.067	5/74	0.802	0.196	0.127	11/74	0.658	0.517	0.372	26/74
S6	0.824	0.400	0.400	5/34	0.451	0.143	0.112	12/34	0.902	0.400	0.222	3/34
S7	0.976	0.000	0.000	1/41	0.642	0.571	0.334	16/41	0.934	0.000	0.000	2/41
S8	0.603	0.566	0.548	31/68	0.745	0.531	0.315	18/68	0.755	0.355	0.432	15/68
S9	0.629	0.324	0.229	16/54	0.451	0.348	0.334	23/54	0.741	0.250	0.083	12/54
S10	0.970	0.000	0.000	1/33	0.768	0.611	0.407	9/33	0.848	0.000	0.000	3/33
S11	0.667	0.614	0.530	22/53	0.711	0.056	0.034	10/53	0.792	0.000	0.000	8/53
S12	0.744	0.600	0.500	4/13	0.769	1.000	0.250	4/13	1.000	-	-	0/13
S13	0.926	0.000	0.000	2/27	0.963	0.000	0.000	1/27	1.000	-	-	0/27
S14	0.527	0.526	0.555	18/36	0.834	0.200	0.166	4/36	0.963	0.625	0.833	2/36
Overall	0.768	0.322	0.261	149/619	0.670	0.313	0.192	156/619	0.804	0.263	0.240	111/619
Baseline - Majority	0.735	0.000	0.000		0.743	0.000	0.000		0.833	0.000	0.000	
Baseline - Random	0.716	0.256	0.251		0.650	0.249	0.266		0.755	0.173	0.170	
Improvement Over Random (%)	7.249	25.850	4.188		3.146	25.990	-27.754		6.468	52.368	40.828	
General	0.543	0.246	0.436		0.521	0.268	0.519		0.554	0.189	0.450	
General Baseline - Majority	0.759	0.000	0.000		0.748	0.000	0.000		0.821	0.000	0.000	
General Baseline - Random	0.635	0.255	0.268		0.585	0.198	0.212		0.687	0.151	0.162	

Table 4: Results of predictions using the individual models. The distribution columns show the proportion of the minority class out of the total number of responses for each of the three variables. The baseline rows represent the averaged results of our baseline classifiers. The general row shows the averaged results of our models trained on all participants.

### 5.3. Individual Classifiers

Since peoples' experience of stress, focus, and awakeness (as well as their physiological manifestations) can vary substantially (e.g., [36]), we first trained and evaluated individual classifiers for each participant (as opposed to a general one for all participants) using leave-one-out cross validation. The results of our analysis are reported in Table 4. For our analysis, we report values of accuracy, one of the most commonly used metric to compare performance, as well as precision and recall of the classes of interest: 'stressed', 'not focused', and 'not awake'. Since the imbalance in the data can lead to high accuracy values if a classifier always just predicts the most likely/frequent class while ignoring the class of higher importance and interest, precision and recall of the class of interest are also important to consider [92, 94, 36].

Besides our results, for the purposes of comparison, we also present two commonly used baselines in this table - a majority classifier, which always predicts the larger of the two classes



considered, and a stratified random classifier which randomly chooses between the two classes, but with a proportional bias towards the larger class.

For some users (i.e., S1. as seen in Table 3), the imbalance in their data was so extreme that even after adjusting by oversampling, we were not able to create a reasonable classifier. These scenarios are difficult to predict, as any classifier will not have enough variance in its training data for the 'stressed' situation to adequately distinguish it from the non-stressed case.

Overall, we were able to use extracted physiological features to predict all three aspects with reasonable accuracy, precision, and recall. We present a comparison between the averaged results of our individual classifiers and those of the baseline stratified random classifier in the "Improvement Over Random" row of Table 4. This is calculated as the difference between the overall average and the baseline results, divided by the baseline results (for example,  $\frac{Acc_{Overall} - Acc_{Random}}{Acc_{Random}}$ ). We do not compare our results with the majority classifier directly as this classifier achieved precision and recall scores of zero when predicting stress, lack of focus, and lack of awakeness, making a meaningful comparison unfeasible. The improvement percentages demonstrate that the predictions made by our classifiers are much better than random after correcting for the imbalance in our dataset.

While the individually trained classifiers improved on average across all participants upon the baseline in all cases except in recall of 'not focused', the improvement was substantially higher for awakeness (52.4% improvement in precision, 40.8% in recall, and 6.5% in accuracy) than for stress or focus. In addition, the performance of the individually trained classifiers varied greatly across participants. While some participants showed a large improvement, for others the baseline performed much better than the individually trained classifier. For instance, for predicting 'stressed', the individual classifiers improved upon the baseline for S4, S6, S8, S11, S12, and S14 with a maximum improvement of 152.0% in precision and 111.1% in recall for S12, while they did worse for S1, S3, S5, S7, S9, S10, and S13, and in the worst cases did not correctly predict a single instance of 'stressed'. Typically, users that have the lowest precision and recall values are those where the data is the most imbalanced.

#### 5.4. Feature Selection and Importance

There are a large variety of features that can be (and have been) calculated in previous research for each of the basic measurements listed in Table 1, such as the mean, standard deviation, maximum, and interquartile range. In addition, each of these metrics can be combined with the various time windows captured of a basic measurement, resulting in a large feature space. To reduce the feature space, we experimented with multiple feature selection methods, including selecting the top k highest correlated features by various metrics such as mutual information, Pearson's correlation coefficient, ANOVA's F-value, as well as wrapper methods such as recursive feature elimination, optimizing mean decrease accuracy by iteratively permuting features, and only selecting features that exceed a certain Gini importance threshold. We found that all methods produced similar results with respect to accuracy, precision, and recall for the individual models. Ultimately, we elected not to utilize any feature selection in order to simplify our approach, as we found there to be minimal differences in performance between the techniques, and the random forest algorithm is capable of (and robust for) handling datasets with many features.

Overall, the features that were selected as the important ones for the individual models based on the random forest algorithm varied greatly across participants. Yet, some feature categories were considered to be important more frequently than others. Table 5 shows the averaged Gini importance for the feature categories used for predicting stress. Heart rate variability proved to

be an important measure for all of the aspects of interest, ranking as the most important feature category for both stress and awakensness, and the second most important one for focus. This is not surprising, as heart rate variability and skin temperature have been shown in several previous studies to be possible indicators for stress levels [27, 26, 8]. We also found blood pulse wave to be an important indicator for both focus and awakensness, but less important for stress, while respiration rate was important in stress and focus but not awakensness. Besides these mentioned feature categories, there was great variation in which measures were important to which of the three aspects. This shows that there is a clear benefit to having multiple biometric streams available for predicting stress, focus and awakensness.

Feature Category	Stress	Focus	Awakensness
Heart Rate Variability	18.3%	13%	13.6%
Blood Pulse Wave	10%	14.2%	13.1%
Heart Rate	8.7%	12.6%	10.3%
Skin Temperature	15.7%	9.8%	10%
Galvanic Skin Response	6.6%	8.2%	5.1%
Respiration Rate	14.8%	12.7%	10%
Oxygen Saturation	5.6%	4.3%	2%
Energy Expenditure	6%	7.7%	4.8%
Activity	4.6%	7.8%	7.6%
Steps	1.7%	0.8%	0.9%
Time of Day	0%	0.1%	0.5%

Table 5: The averaged Gini importance of each feature category, per response variable.

### 5.5. Individual vs. General Model

Individual models are trained specifically for each individual and thus require a data collection period before they are capable of making accurate predictions. On the other hand, the idea of general models is to be able to train them on already collected data and then to be able to apply them even to new and unseen individuals, thus overcoming the cold-start problem. Given the large individual differences in biometrics, training a general model to achieve an adequate accuracy for new individuals is not necessarily possible.

To examine the performance of a general model for our participants, we trained three general models, one for focus, one for awakensness and one for stress. We roughly followed the same procedure as for the individual models. Due to the larger amount of data available in the general case, we used the more common random undersampling, which randomly selects elements in the majority class to exclude from the dataset, instead of random oversampling to balance the distribution of the dataset. The models were trained on the datasets of 13 of the 14 participants, and then evaluated on the dataset of the last (leave-one-participant-out cross-validation), repeating this process for all 14 participants.

The bottom three rows of Table 4 present the averaged performance results for this approach in terms of accuracy, precision, and recall, as well as the results of baseline stratified random and majority classifiers following the same leave-one-group-out cross-validation procedure. Although the averaged precision and recall are comparable or better than those of the averaged individual results, this was at the cost of a large decrease in overall accuracy. Upon closer investigation into the performance of the general model when testing on each participant, we found that individually

trained models for each participant performed much better than a general model trained over all participants. Using stress as an example, for participant S12, for whom we saw the greatest increase compared to the baseline in individual models, the general model was unable to predict a single instance of ‘stressed’ correctly. This is consistent with our expectations because biometric features are highly specific to individuals.

## 6. Discussion

Humans experience stress, focus and awakesness in different ways. In this paper, we have attempted to study these mental states in the workplace using both participant self-reports and biometric data. We discuss implications from our study for the workplace, including ways in which the information might inform digitally-controlled or digitally-informed parts of the workplace. We also discuss challenges imposed by the data and possible future paths of research.

### 6.1. *Implications for Workplaces*

Being able to accurately recognize periods of high-stress in knowledge workers could enable more respectful workplaces. For example, an ability to sense and predict stress in the moment could help companies prevent or de-escalate potentially dangerous situations in the workplace (e.g., confrontation between co-workers). Building an understanding of stress and focus over time and in the moment could also help create workplaces that are more conducive to enabling knowledge workers to be more productive. For example, this information could be used to feed an awareness dashboard of a team’s stress level, and avoid digital interruptions in high-stress or high-focus periods similar to previous studies [95]. Building an understanding of awakesness and focus could also enable the creation of workplaces that are conducive to workers producing higher-quality work. For example, if awakesness or focus decreases, they might be enhanced by adapting lighting in the workplace or scheduling breaks to prevent focus loss.

Currently, the cost of biometric sensors and necessary infrastructure, such as automated light and sound systems for adjusting the environment, makes our approach most appropriate for high-value workspaces, such as control rooms, command centers, or dispatch offices. However, as standard office settings become more personalizable (e.g., via adjustable desks, lighting, and sound showers) and sensor costs decrease, our approach could be applied to any office environment, and thus could impact a large percentage of modern workers.

As in modern cars, temperature and lighting could be regulated on a per-person basis, which would allow the environment to react to the person’s current state and to maximize each person’s preferences and productivity (e.g., preferences of men and women in temperature [96]). We believe that our results present a good step to more in situ usage of biometric sensing. Future studies can build on the evidence, and for example, reduce the effort required for the experience sampling from continuously rating stress levels to using biometrics as a ground truth with occasional validation of the predicted stress levels, or identify how to balance needs across a group of office workers and how to handle conflicting levels between different group members.

### 6.2. *The Effect of Tai Chi on Stress*

As described in our study procedure, due to our focus on stress, we offered participants an opportunity to be part of Tai Chi classes. The two primary reasons for this intervention were to keep participants motivated to continue the self-reports over the long study period, and to offer a technique that might help reduce stress. After consulting with other researchers, we decided to

offer Tai Chi in the last month of the study, leaving the first month of the study unaltered. While participants were not required to take a Tai Chi class, they all took one per week except for two participants that opted out of the last two weeks of classes. Four participants explicitly stated that they liked the Tai Chi sessions or that they were “excellent”. These classes were performed once per week with a duration of one hour.

While this was not the focus of our study, we performed a secondary analysis to examine whether the Tai Chi classes had any effect on the participants’ stress levels in the workdays directly following the Tai Chi session. For this, we build a linear mixed model with the self-reported daily stress level as dependent variables and the participants as random effects. We found that Tai Chi attendance contributed a small amount to decreased stress in the work week immediately following the Tai Chi session (slope of -0.188,  $p < 0.05$ ). However, we did not find a connection between attendance and stress on the day of the session suggesting that Tai Chi might have long-term effects, but is not conducive at relieving stress close in time to the intervention. Overall, the Tai Chi thus had a small impact on the collected data of the second month of the study, which poses a threat to the validity to our observed trends for this period of time. However, since knowledge workers might attend these kinds of classes on their own doing, we believe that this is negligible, and the analysis rather provides a weak indication that this kind of stress intervention might in fact help to reduce stress in the wild.

### *6.3. Ground Truth and Self-Reporting*

Studying mental states, such as stress, awakeness and focus, requires collecting a valid ground truth from each participant. We spent considerable time when designing the study determining the exact questions to ask of participants, consulting experts in the area, and basing the questions and wording on previous research and studies. Despite the care taken, it is possible that the gathered data lacks reliability and validity. Some have questioned the reliability and validity of self-reports as we used in our study due to subjective biases, lack of care in reporting, and the highly individual nature of reporting aspects such as stress [36, 32].

In addition, in contexts such as the workplace, as in our study, participants might be afraid to genuinely report levels of aspects, such as sleepiness. Hence, there is a chance that the self-reports we gathered do not adequately reflect the ground truth of the underlying variables under investigation. It could even be the case that certain biometrics might represent a more accurate ground truth of the studied phenomena than the self-reports. This suggests that a more confirmative study rather than an inquiry study could be a better approach, and we will explore such routes in future work.

### *6.4. Imbalanced Data*

Study participants provided highly imbalanced data in their survey responses, with most participants only taking advantage of a subset of the Likert-scale values and the data points mostly being clustered around the middle of the scale, as can be seen in Table 3. While some of the imbalance is expected due to certain classes, such as ‘not stressed’, being more common in the workplace, this imbalance also provides challenges in the training and assessment of a machine learning classifier, as also found by others, e.g. [97]. We addressed this for the training by oversampling in case of few data samples for the individual models and undersampling in case of a general model where more data was available. Rebalancing the dataset using such techniques is a common and effective practice [98]. Alternative techniques such as SMOTE [99] exist and can perform better than those we employed, however they are impractical considering the limited amount of data we have to work with.

Oversampling the dataset as we did may also lead to an increased risk of model overfitting. However, we believe the benefits of rebalancing the dataset outweigh this possible downside. In light of the imbalance in the data, the results we achieved with our models are encouraging. For the assessment of the classifiers' performance we addressed the imbalance by not just presenting accuracy, but also by focusing on prediction and recall to examine the classifier's performance in predicting the infrequent (yet more important) cases, such as when a user is struggling to stay awake and an intervention or warning might be needed most.

#### *6.5. Predicting Stress with Computer Interaction Data*

Knowledge workers, a focus of our work and study, often spend a large amount of time each day interacting with information on their computer at work. An interesting direction for future study is to consider whether this computer interaction data, which can be gathered non-invasively as work occurs, could serve to sense and predict stress, focus and awakeness. Features that could be investigated include keystrokes per minute, mouse clicks per minute and changes in the active window title.

### **7. Threats to Validity**

There are numerous threats to validity to our study.

#### *7.1. External Validity*

The results of our study may not generalize to a broader population of office workers. To mitigate this risk, we collected participants from a wide variety of departments with different age ranges, genders, work experience, and working in different positions therefore providing evidence that our approach's performance is comprised of and can generalize to a wide range of knowledge workers.

Secondly, our results may not generalize to a different office environment. We conducted this study in a typical office environment, similar to many among technology workers across the world. These office environments control for a series of variables to make them standard worldwide such as controlled temperature and lighting.

#### *7.2. Internal Validity*

This study tries to find correlations between biometric features and the human aspects of stress, focus, and awakeness. Nonetheless, biometric signals are influenced by far more variables than the ones this study comprehends. Therefore, trying to draw a strong causality between the biometric features and the aspects would be inaccurate. To mitigate this risk, we collected the data in a rote environment and in a regular manner to minimize the number of external causes that may affect each participant's biometric signals.

It is possible that the amount of data collected is not sufficient to draw valid conclusions. To address this threat, we collected a data for an eight-week period, which is 400% longer than the longest previous studies [50, 100].

Due to the small amount of data available to us for the purposes of this study, there is a risk that our models may be overfitting to some degree. Such overfitting indicates that the results presented in our work may be less than optimal, however any overfitting applies strictly to the training data and given that there is no overlap between the data used for training and testing, we do not believe this invalidates our results. We leave further optimization of our approach to future researchers.

### 7.3. Construct Validity

A threat to the study is that there are other factors that might either influence the human aspects of interest or that were considered but are unrelated biometric signals. To mitigate this risk, we used a state-of-the-art device that captures a large number of highly accurate biometric signals. We collected the most commonly analyzed biometrics that historically have shown correlation with the studied human aspects of stress, focus and awakeness. In an effort to maintain this research applicable to real-world environments, we picked the already existing Everion device, even when, as a trade-off, we could not capture more descriptive and more intrusive signals such as SDNN, SCL, SCR, eye tracking, or brain activity.

A future more thorough statistical analysis of the relationships between the aspects of interest, such as stress, and the physiological data may further provide deeper insights into the data and how it might be used in prediction.

## 8. Conclusions

Stress, awakeness, and focus at work are highly relevant aspects when it comes to productivity and well-being at the workplace. In this paper, we presented the results of a study with 14 professional knowledge workers in their workplace over an eight-week period to better understand how workers experience these human aspects over time and to examine the ability of biometrics to predict these aspects. The longitudinal and in-situ placement of the study support and extend previous work. Based on daily collected survey responses, we observed that although participants sometimes saw periods of sustained stress or sleepiness, they would always return to their baseline reporting level at some point. We also observed that stress levels seldom spiked, but when they did rise, the rise in stress tended to last more than a day. In addition to the survey responses, we continually collected biometric data with which we were able to create a model that is able to predict user stress, sleepiness, and lack of focus with small improvements in accuracy (from 3.1% to 7.2%, depending on the aspect in question), and moderate improvements in precision (25.9% - 52.4%) in comparison to a stratified random classifier. While the precision and recall scores we report are still low overall, this is a difficult problem to solve and our improvements indicate the potential for future researchers to build upon.

These results open up new opportunities to help increase knowledge workers' productivity and well-being, ranging from instantaneously taking action to prevent potentially risky situations and prevent accidents due to a lack of focus or awakeness, all the way to recommending interventions to reduce stress if it becomes more chronic.

- [1] P. F. Drucker, Knowledge-worker productivity: The biggest challenge, *California management review* 41 (2) (1999) 79–94.
- [2] V. M. González, G. Mark, Constant, constant, multi-tasking craziness: managing multiple working spheres, in: *Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM, 2004, pp. 113–120.
- [3] G. Mark, D. Gudith, U. Klocke, The cost of interrupted work: more speed and stress, in: *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, ACM, 2008, pp. 107–110.
- [4] M. Czerwinski, E. Horvitz, S. Willhite, A diary study of task switching and interruptions, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '04, ACM, 2004, pp. 175–182.
- [5] G. R. J. Hockey, Compensatory control in the regulation of human performance under stress and high workload: A cognitive-energetical framework, *Biological Psychology* 45 (1) (1997) 73 – 93.
- [6] C. Setz, B. Amrich, J. Schumm, R. L. Marca, G. Troster, U. Ehlert, Discriminating stress from cognitive load using a wearable eda device, *IEEE Transactions on Information Technology in Biomedicine* 14 (2) (2010) 410–417.
- [7] E. F. for the Improvement of Living, W. Conditions, Work-related stress, Tech. rep., <https://www.eurofound.europa.eu/printpdf/publications/report/2010/work-related-stress> (2010).

- [8] H. Kataoka, H. Kano, H. Yoshida, A. Saijo, M. Y. . M. Osumi, Development of a skin temperature measuring system for non-contact stress evaluation, in: IEEE Engineering in Medicine and Biology Society, Vol. 20, 2000.
- [9] J. R. Piazza, D. M. Almeida, N. O. Dmitrieva, L. C. Klein, Frontiers in the use of biomarkers of health in research on stress and aging, in: 35th Annual International Conference of the IEEE EMBS, Vol. 65B (5), 2010, pp. 513–525.
- [10] J. Connor, S. Ameratunga, R. Norton, E. Robinson, R. Dunn, J. Bailey, I. Civil, R. Jackson, Driver sleepiness and risk of serious injury to car occupants: population based case control study, in: BMJ, 2002.
- [11] S. Cohen, R. C. Kessler, L. U. Gordon, Measuring stress: A guide for health and social scientists, Oxford University Press on Demand, 1997.
- [12] G. Mark, S. T. Iqbal, M. Czerwinski, P. Johns, Bored Mondays and focused afternoons: the rhythm of attention and online activity in the workplace, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ACM, 2014, pp. 3025–3034.
- [13] A. Sano, R. W. Picard, Stress recognition using wearable sensors and mobile phones, in: Humaine Association Conference on Affective Computing and Intelligent Interaction (ACII), IEEE, 2013, pp. 671–676.
- [14] J. A. Healey, R. W. Picard, Detecting stress during real-world driving tasks using physiological sensors, Transactions on intelligent transportation systems 6 (2) (2005) 156–166.
- [15] J. Wijsman, B. Grundlehner, H. Liu, H. Hermens, J. Penders, Towards mental stress detection using wearable physiological sensors, in: Engineering in Medicine and Biology Society, IEEE, 2011, pp. 1798–1801.
- [16] M. Züger, T. Fritz, Interruptibility of software developers and its prediction using psycho-physiological sensors, in: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, ACM, 2015, pp. 2981–2990.
- [17] N. Goyal, S. R. Fussell, Intelligent interruption management using electro dermal activity based physiological sensor for collaborative sensemaking, Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies 1 (3) (2017) 52.
- [18] C. Parin, Subvocalization - toward hearing the inner thoughts of developers, in: Proceedings of International Conference on Program Comprehension, 2011.
- [19] T. Nakagawa, Y. Kamei, H. Uwano, A. Monden, K. Matsumoto, D. M. German, Quantifying programmers' mental workload during program comprehension based on cerebral blood flow measurement: A controlled experiment, in: Companion Proc. of ICSE, 2014.
- [20] S. Radevski, H. Hata, K. Matsumoto, Real-time monitoring of neural state in assessing and improving software developers' productivity, Proceedings of Connected Health: Applications, Systems and Engineering Technologies.
- [21] Biovotion, Everion, <https://www.biovotion.com/everion/>, [Online; accessed 9-July-2019] (2019).
- [22] T. Chandola, A. Heraclides, M. Kumari, Psychophysiological biomarkers of workplace stressors, in: Neurosci Biobehav Rev., Vol. 35, 2010, pp. 51–57.
- [23] D. H. Hellhammer, S. Wust, B. M. Kudielka, Salivary cortisol as a biomarker in stress research, in: Psychoneuroendocrinology, Vol. 34, 2009, pp. 163–171.
- [24] A. P. J. van Eekelen, J. H. Houtveen, G. A. Kerkhof, Circadian variation in base rate measures of cardiac autonomic activity, in: European Journal of Applied Physiology, Vol. 93, 2004, pp. 39–46.
- [25] M. Valentini, G. Parati, Variables influencing heart rate, in: Progress in Cardiovascular Diseases, Vol. 52, 2010, pp. 11–19.
- [26] D. J. McDuff, J. Hernandez, S. Gontarek, R. W. Picard, Cogcam: Contact-free measurement of cognitive stress during computer tasks with a digital camera, in: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, CHI '16, ACM, 2016, pp. 4000–4004.
- [27] R. K. Dishman, Y. Nakamura, M. E. Garcia, R. W. Thompson, A. L. Dunn, S. N. Blair, Heart rate variability, trait anxiety, and perceived stress among physically fit men and women, International Journal of Psychophysiology 37 (2) (2000) 121 – 133.
- [28] T. F. of the European Society of Cardiology the North American Society of Pacing Electrophysiology, Heart rate variability, standards of measurement, physiological interpretation, and clinical use, in: European Heart Journal, Vol. 93, 1996, pp. 1043–1065.
- [29] V. Gal, V. Vuksanovic, Heart rate variability in mental stress aloud, in: Medical Engineering and Physics, Vol. 29, 2007, pp. 344–349.
- [30] N. Montano, A. Porta, C. Cogliati, G. Costantino, E. Tobaldini, K. R. Casali, F. Iellamo, Heart rate variability explored in the frequency domain: A tool to investigate the link between heart and behavior, in: Neuroscience and Biobehavioral Reviews, Vol. 33, 2009, pp. 71–80.
- [31] S. Zaman, A. Wesley, D. R. D. Silva, P. Buddharaju, F. Akbar, G. Gao, G. Mark, R. Gutierrez-Osuna, I. Pavlidis, Stress and productivity patterns of interrupted, synergistic, and antagonistic office activities, in: Scientific data, 2019, pp. 1–18.
- [32] K. Hovsepian, M. al'Absi, E. Ertin, T. Kamarck, M. Nakajima, S. Kumar, cstress: Towards a gold standard for continuous stress assessment in the mobile environment, in: ACM international conference on Ubiquitous computing, 2015, pp. 493–504.

- [33] J. Hernandez, P. Paredes, A. Roseway, M. Czerwinski, Under pressure: Sensing stress of computer users, in: CHI Conference on Human Factors in Computing Systems, 2014.
- [34] L. M. Vizer, L. Zhou, A. Sears, Automated stress detection using keystroke and linguistic features: An exploratory study, *International Journal of Human-Computer Studies* 67 (10) (2009) 870–886. doi:10.1016/j.ijhcs.2009.07.005.  
URL <http://www.sciencedirect.com/science/article/pii/S1071581909000937>
- [35] R. Kocielnik, N. Sidorova, F. M. Maggi, M. Ouwerkerk, J. H. D. M. Westerink, Smart technologies for long-term stress monitoring at work, *Proceedings of the 26th IEEE International Symposium on Computer-Based Medical Systems (2013)* 53–58doi:10.1109/cbms.2013.6627764.
- [36] J. Hernandez, R. R. Morris, R. W. Picard, Call center stress recognition with person-specific models, in: *International Conference on Affective Computing and Intelligent Interaction*, 2011, pp. 125–134.
- [37] Z. Sarsenbayeva, N. van Berkel, D. Hettiachchi, W. Jiang, T. Dingler, E. Velloso, V. Kostakos, J. Gonçães, Measuring the effects of stress on mobile interaction, in: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2019, pp. 1–18.
- [38] G. W. Evans, D. Johnson, Stress and open-office noise, in: *Journal of Applied Psychology*, Vol. 25(5), 2000, pp. 779 – 783.
- [39] C. Cherniss, *Staff Burnout - Job Stress in the Human Services*, Sage Publications, Inc, 1980.
- [40] S. Mirjafari, K. Masaba, T. Grover, W. Wang, P. G. Audia, A. T. Campbell, N. V. Chawla, V. D. Swain, M. D. Choudhury, A. K. Dey, S. K. D’Mello, G. Gao, J. M. Gregg, K. Jagannath, K. Jiang, S. Lin, Q. Liu, G. J. Mark, G. J. Martinez, S. M. Mattingly, E. Moskal, R. Mulukutla, S. Nepal, K. A. Nies, M. D. Reddy, P. D. Robles-Granda, K. Saha, A. Sirigiri, A. D. Striegel, Differentiating higher and lower job performers in the workplace using mobile sensing, in: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2019, pp. 1–24.
- [41] J. R. Anderson, *Cognitive Psychology and Its Implications*, 2004.
- [42] J. Chong, R. Siino, Interruptions on software teams: a comparison of paired and solo programmers, in: *Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*, ACM, 2006, pp. 29–38.
- [43] S. T. Iqbal, E. Horvitz, Disruption and recovery of computing tasks: Field study, analysis and directions, in: *SIGCHI Conference on Human Factors in Computing Systems*, 2007, pp. 677–686.
- [44] B. P. Bailey, J. A. Konstan, J. V. Carlis, The effects of interruptions on task performance, annoyance, and anxiety in the user interface, in: *Proceedings of INTERACT*, Vol. 1, 2001, pp. 593–601.
- [45] M. Czerwinski, E. Cutrell, E. Horvitz, Instant messaging: Effects of relevance and timing, in: *People and computers XIV: Proceedings of HCI*, Vol. 2, British Computer Society, 2000, pp. 71–76.
- [46] J. Fogarty, A. J. Ko, H. H. Aung, E. Golden, K. P. Tang, S. E. Hudson, Examining task engagement in sensor-based statistical models of human interruptibility, in: *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, ACM, 2005, pp. 331–340.
- [47] S. T. Iqbal, B. P. Bailey, Effects of intelligent notification management on users and their tasks, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 2008, pp. 93–102.
- [48] D. Chen, R. Vertegaal, Using mental load for managing interruptions in physiologically attentive user interfaces, in: *CHI’04 extended abstracts on Human factors in computing systems*, ACM, 2004, pp. 1513–1516.
- [49] B. P. Bailey, S. T. Iqbal, Understanding changes in mental workload during execution of goal-directed tasks and its application for interruption management, *ACM Transactions on Computer-Human Interaction (TOCHI)* 14 (4) (2008) 21.
- [50] M. Zuger, S. Muller, A. Meyer, T. Fritz, Sensing interruptibility in the office: A field study on the use of biometric and computer interaction sensor, in: *Conference on Human Factors in Computing Systems*, 2018.
- [51] R. Agarwal, E. Karahanna, Time flies when you’re having fun: Cognitive absorption and beliefs about information technology usage, in: *MIS Quarterly*, Vol. 24 (4), 2000, pp. 665–694.
- [52] J. Webster, H. Ho, Audience engagement in multimedia presentations, in: *Data Base for the Advancement in Information Systems*, Vol. 28(2), 1997, pp. 63–77.
- [53] M. Csikszentmihalyi, *Flow: The Psychology of Optimal Experience*, 1990.
- [54] K. E. Weick, K. M. Sutcliffe, Mindfulness and the quality of organizational attention, in: *Organization Science*, Vol. 17, 2006, pp. 514–524.
- [55] E. Dane, Paying attention to mindfulness and its effect on task performance in the workplace, in: *Journal of Management*, Vol. 37 (4), 2011, pp. 997–1018.
- [56] S. Nordbakke, F. Sagberg, Sleepy at the wheel: Knowledge, symptoms and behaviour among car drivers, in: *Transportation Research Part F: Traffic Psychology and Behaviour*, Vol. 10, 2007, pp. 1–10.
- [57] T. Fritz, A. Begel, S. C. Müller, S. Yigit-Elliott, M. Züger, Using psycho-physiological measures to assess task difficulty in software development, in: *Proceedings of the 36th International Conference on Software Engineering*, ACM, 2014, pp. 402–413.
- [58] R. Bednarik, M. Tukiainen, An eye-tracking methodology for characterizing program comprehension processes, in:



- Proc. of ETRA, 2006.
- [59] M. Crosby, J. Stelovsky, How do we read algorithms? a case study, *Computer* 23 (1).
  - [60] P. Rodeghero, C. McMillan, P. W. McBurney, N. Bosch, S. D’Mello, Improving automated source code summarization via an eye-tracking study of programmers, in: *International Convergence on Software Engineering*, 2014.
  - [61] Y. Ikutani, H. Uwano, Brain activity measurement during program comprehension with NIRS, in: *Proc. of SNPD*, 2014.
  - [62] J. Siegmund, C. Kästner, S. Apel, C. Parnin, A. Bethmann, T. Leich, G. Saake, A. Brechmann, Understanding source code with functional magnetic resonance imaging, in: *Proceedings of International Conference on Software Engineering*, 2014.
  - [63] J. Beatty, Task-evoked pupillary responses, processing load, and the structure of processing resources, in: *Psychological Bulletin*, Vol. 91(2), 1982.
  - [64] J. Klingner, Fixation-aligned pupillary response averaging, in: *Symposium on Eye-Tracking Research and Applications*, 2010, pp. 275–282.
  - [65] H. Berger, Über das elektrenkephalogramm des menschen, in: *European Archives of Psychiatry and Clinical Neuroscience*, Vol. 87, 1929, pp. 527–570.
  - [66] M. E. Smith, A. Gevins, Neurophysiologic monitoring of mental workload and fatigue during operation of a flight simulator, in: *The International Society for Optical Engineering*, 2005, pp. 116–126.
  - [67] M. B. Serman, C. A. Mann, D. A. Kaiser, Quantitative eeg patterns of differential in-flight workload, in: *6th Annual Workshop on Space Operations Applications and Research*, Vol. 2, 1993.
  - [68] R. Ferdous, V. Osmani, O. Mayora, Smartphone app usage as a predictor of perceived stress levels at workplace, in: *International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth)*, 2015.
  - [69] Y. Okada, T. Y. Yoto, T. aki Suzuki, S. Sakuragawa, H. Sakakibara, K. Shimoi, T. Sugiura, Wearable eeg recorder with acceleration sensors for monitoring daily stress\*: Office work simulation study, in: *Journal of Medical and Biological Engineering*, Vol. 34, 2011, pp. 420–426.
  - [70] P. Electro, Equine H7 heart rate sensor belt set, [https://www.polar.com/en/products/equine/accessories/equine\\_H7\\_heart\\_rate\\_sensor\\_belt\\_set](https://www.polar.com/en/products/equine/accessories/equine_H7_heart_rate_sensor_belt_set), [Online; accessed 9-July-2019] (2019).
  - [71] I. Fitbit, Fitbit Charge 2, <https://www.fitbit.com/de/charge2>, [Online; accessed 9-July-2019] (2019).
  - [72] L. Barrios, P. Oldrati, S. Santini, A. Lutterotti, Evaluating the accuracy of heart rate sensors based on photoplethysmography for in-the-wild analysis, in: *Proceedings of the 13th EAI International Conference on Pervasive Computing Technologies for Healthcare, PervasiveHealth’19, Association for Computing Machinery*, New York, NY, USA, 2019, p. 251–261. doi : 10.1145/3329189.3329215.  
URL <https://doi.org/10.1145/3329189.3329215>
  - [73] K. R. Fox, The influence of physical activity on mental well-being, *Public Health Nutrition* 2 (3) (1999) 411–418.
  - [74] S. G. Aldana, L. D. Sutton, B. H. Jacobson, M. G. Quirk, Relationships between leisure time physical activity and perceived stress, *Perceptual and Motor Skills* 82 (1) (1996) 315–321. doi : 10.2466/pms.1996.82.1.315.  
URL <https://doi.org/10.2466/pms.1996.82.1.315>
  - [75] E. Haapalainen, S. Kim, J. F. Forlizzi, A. K. Dey, Psycho-physiological measures for assessing cognitive load, in: *Proceedings of the 12th ACM international conference on Ubiquitous computing*, ACM, 2010, pp. 301–310.
  - [76] L. Mulder, Measurement and analysis methods of heart rate and respiration for use in applied environments, *Biological psychology* 34 (2) (1992) 205–236.
  - [77] A. Haag, S. Goronzy, P. Schaich, J. Williams, Emotion recognition using bio-sensors: First steps towards an automatic system, in: *Tutorial and research workshop on affective dialogue systems*, Springer, 2004, pp. 36–48.
  - [78] Y. Masaoka, I. Homma, Anxiety and respiratory patterns: their relationship during mental stress and physical load, *International Journal of Psychophysiology* 27 (2) (1997) 153–159.
  - [79] A. N. Meyer, G. C. Murphy, T. Zimmermann, T. Fritz, Design recommendations for self-monitoring in the workplace: Studies in software development, in: *Proceedings ACM Human Computer Interaction*, 2017, p. 24. doi : <https://doi.org/10.1145/3134714>.
  - [80] S. Lalle, C. Conati, G. Carenini, Predicting confusion in information visualization from eye tracking and interaction data, in: *International Joint Conference on Artificial Intelligence*, 2016.
  - [81] P. Panwar, C. M. Collins, Detecting negative emotion for mixed initiative visual analytics, in: *CHI Conference on Human Factors in Computing Systems*, 2018.
  - [82] Y. Luo, B. Lee, D. Y. Wohn, A. L. Rebar, D. E. Conroy, E. K. Choe, Time for break: Understanding information workers’ sedentary behavior through a break prompting system, in: *CHI Conference on Human Factors in Computing Systems*, 2018.
  - [83] A. T. Adams, E. L. Murnane, P. Adams, M. Elfenbein, P. F. Chang, S. Sannon, G. Gay, T. Choudhury, Keppi: A tangible user interface for self-reporting pain, in: *CHI Conference on Human Factors in Computing Systems*, 2018.
  - [84] P. Gloor, D. Oster, O. Raz, A. Pentland, D. Schoder, The virtual mirror: Reflecting on the social and psychological self to increase organizational creativity, *International Studies of Management & Organization* 40 (2) (2010) 74–94.

- [85] P. Wilhelm, D. Schoebi, Assessing mood in daily life: Structural validity to change, and reliability of a short-scale to measure three basic dimensions of mood, *European Journal of Psychological Assessment* 23 (4) (2007) 258–267.
- [86] J. Russell, A circumplex model of affect, *Journal of Personality and Social Psychology* 39 (6) (1980) 1161–1178.
- [87] J. Fogarty, A. J. Ko, H. H. Aung, E. Golden, K. P. Tang, S. E. Hudson, Examining task engagement in sensor-based statistical models of human interruptibility, in: *SIGCHI Conference on Human Factors in Computing Systems*, 2005, pp. 331–340.
- [88] T. Tanaka, K. Fujita, Study of user interruptibility estimation based on focused application switching, in: *Conference on Computer Supported Cooperative Work*, 2011, pp. 721–724.
- [89] P. V. A. Bernstein, A. Zurfluh, Interruptibility prediction using motion detection, in: *Workshop on Managing Context Information in Mobile and Pervasive Environments*, 2005.
- [90] G. Hawthorne, G. Hawthorne, P. Elliott, Imputing cross-sectional missing data: comparison of common techniques, *Australian & New Zealand Journal of Psychiatry* 39 (7) (2005) 583–590.
- [91] N. V. Chawla, N. Japkowicz, A. Kotcz, Editorial: Special issue on learning from imbalanced data sets, *SIGKDD Explor. Newsl.* 6 (1) (2004) 1–6. doi : 10 . 1145/1007730 . 1007733.
- [92] B. W. Yap, K. A. Rani, H. A. A. Rahman, S. Fong, Z. Khairudin, N. N. Abdullah, An application of oversampling, undersampling, bagging and boosting in handling imbalanced datasets, in: T. Herawan, M. M. Deris, J. Abawajy (Eds.), *Proceedings of the First International Conference on Advanced Data and Information Engineering (DaEng-2013)*, Springer Singapore, 2014, pp. 13–22.
- [93] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Clondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine learning in python, in: *Journal of Machine Learning Research* 12, 2011, pp. 2825–2830.
- [94] S. Bhattacharyya, S. Jha, K. Tharakunnel, J. C. Westland, Data mining for credit card fraud: A comparative study, *Decision Support Systems* 50 (3) (2011) 602 – 613.
- [95] M. Züger, C. Corley, A. N. Meyer, B. Li, T. Fritz, D. Shepherd, V. Augustine, P. Francis, N. Kraft, W. Snipes, Reducing interruptions at work: A large-scale field study of flowlight, in: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, ACM, 2017, pp. 61–72.
- [96] S. Karjalainen, Gender differences in thermal comfort and use of thermostats in everyday thermal environments, in: *Building and Environment*, Vol. 42, 2007, pp. 1594–1603.
- [97] A. Exler, A. Schankin, C. Klebsattel, M. Beigl, A wearable system for mood assessment considering smartphone features and data from mobile ecgs, in: *Pervasive and Ubiquitous Computing*, 2016, pp. 1153–1161.
- [98] P. Branco, L. Torgo, R. P. Ribeiro, A survey of predictive modeling on imbalanced domains, *ACM Computing Surveys (CSUR)* 49 (2) (2016) 1–50.
- [99] N. V. Chawla, K. W. Bowyer, L. O. Hall, W. P. Kegelmeyer, Smote: synthetic minority over-sampling technique, *Journal of artificial intelligence research* 16 (2002) 321–357.
- [100] S. Muller, T. Fritz, Using (bio)metrics to predict code quality online, in: *In Proceedings of the ICSE*, 2016.