Detecting Job Promotion in Information Workers Using Mobile Sensing

SUBIGYA NEPAL, Dartmouth College, USA
SHAYAN MIRJAFARI, Dartmouth College, USA
GONZALO J. MARTINEZ, University of Notre Dame, USA
PINO AUDIA, Dartmouth College, USA
AARON STRIEGEL, University of Notre Dame, USA
ANDREW T. CAMPBELL, Dartmouth College, USA

Most people desire promotions in the workplace. Typically, rising through the ranks comes with increased demands, better salary and higher status among peers. However, promoted workers have to deal with new challenges, such as, adjusting to new roles and responsibilities, which can in turn impact their physical and mental wellbeing. In this year long study, we use mobile sensing to track physiological and behavioral patterns of N=141 information workers who are promoted. We show that the workers experience a change in their physiological and behavioral patterns after promotion captured by passive sensing from phones, wearables and Bluetooth beacons. Furthermore, we use a random convolutions based approach to extract patterns from multivariate time series signals and evaluate the performance of different models to classify a worker's mobile sensing data as belonging to a promoted or non-promoted period with an AUC of 0.72. As a result, we report for the first time that mobile sensing can detect job promotion events by modeling physiological and behavioral changes of information workers in an objective manner.

CCS Concepts: • **Human-centered computing → Ubiquitous and mobile computing**; • **Applied computing**;

Additional Key Words and Phrases: Mobile Sensing, Job Promotion, Mobile Behavioral Pattern

ACM Reference Format:
Subigya Nepal, Shayan Mirjafari, Gonzalo J. Martinez, Pino Audia, Aaron Striegel, and Andrew T. Campbell. 2020. Detecting Job Promotion in Information Workers Using Mobile Sensing. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 4, 3, Article 113 (September 2020), 28 pages. https://doi.org/10.1145/3414118

1 INTRODUCTION

Promotions play an important role for organizations and individuals. For organizations, it is a way to keep employees committed and motivated towards the company goals by rewarding promoted workers with financial and status gains. For individuals, rising through the ranks leads to a boost in morale, wellbeing, and life satisfaction [37, 38]. However, promotions can be a mixed blessing for many – while they provide an increase in occupational status, financial reward, job autonomy, privilege and flexibility, they can often also be accompanied by added

Authors' addresses: Subigya Nepal, Dartmouth College, Hanover, NH, 03755, USA, sknepal@cs.dartmouth.edu; Shayan Mirjafari, Dartmouth College, Hanover, NH, 03755, USA; Gonzalo J. Martinez, University of Notre Dame, Notre Dame, IN, 46556, USA; Pino Audia, Dartmouth College, Hanover, NH, 03755, USA; Aaron Striegel, University of Notre Dame, Notre Dame, IN, 46556, USA; Andrew T. Campbell, Dartmouth College, Hanover, NH, 03755, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
© 2020 Association for Computing Machinery.
2474-9567/2020/9-ART113 $15.00
https://doi.org/10.1145/3414118

responsibility, longer working hours, stress and reduced work-life balance. Promotions are therefore, a win-some, lose-some game. While workers may win through the status gain, financial and personal growth, they may impact their psychological wellbeing and work-life balance. In addition, job change or role transitions have the ability to change people [1, 12, 41, 55]. The characteristics of a person may change as a result of change in occupational status [24]. As people transition in their new role after promotion, it might involve psychological, cognitive and behavioral adjustments. Workers might need to embrace their new role in the social network, learn new tasks, get familiar with new routines as well as cope with potentially new physical settings and surroundings [15, 40]. As such, these transitions can have positive as well as negative consequences on a promoted worker's life. Therefore, we believe studying the after effects of promotion is an important topic. Prior studies have linked workplace promotions with health, mental wellbeing, job satisfactions among other adjustment issues for workers [10, 21, 23, 36].

Despite their importance, we do not have any objective assessment of how promotion affects an individual's life. In the absence of literature using objective measures to capture the effects of promotion, we propose a passive sensing based approach as a means to assess workers' reactions to being promoted. We do not concern ourselves with how the promotions happen or why they do, because the subject has been studied thoroughly before [5, 22, 71, 88]. Rather, we focus on physiological and behavioral changes brought about as a result of promotion. More specifically, we look at the period before promotion (i.e., the non-promoted period) and after promotion (i.e., the promoted period) in order to examine differences. We study N=141 information workers across different industries who were promoted during the year long study using data from workers' mobile phones, wearables and Bluetooth beacons. We compare these physiological and behavioral changes before and after promotion taking into account gender and job performance. Unlike prior studies that use self-reported measures to assess the impact of workplace promotion [37, 38], we believe that passively collected data from mobile devices can provide an objective measure of promotion events. The contributions of this paper are as follows:

(1) To the best of our knowledge, we present the first workplace based mobile sensing study to investigate the objective behavioral and physiological changes that occur as a result of promotion. We consider N=141 information workers who are promoted in a year long study. We do this by performing a comparison of workers' physiological and behavioral patterns up to 60 days prior to promotion and 60 days after promotion.
(2) We study the changes in physiological and behavioral patterns associated with promotion considering gender and higher/lower job performers. Our findings show that there are gender differences associated with promotion events; for example, after promotion, female workers experience an increase in stress duration during working hours, but males do not.
(3) We explore the feasibility of using mobile sensing streams to detect job promotion events. Using a binary segmentation based changepoint detection method, we find that mobile sensing streams (e.g., sleep duration, step count) change abruptly during the promotion period, illustrating that the signals have the potential to capture promotion.
(4) We use ROCKET [16], a random convolutions based approach to extract patterns from multivariate time series signals. Using 1D convolutions, we extract features from the multivariate time series and evaluate the performance of different models to classify whether the given mobile sensing data belongs to a promoted period or a non-promoted period. For this analysis, we include non-promoted as well as promoted workers' data. We report an AUC of 0.72. Our findings show that passively collected mobile sensing data from phones, wearables and beacons is able to detect job promotion events.

The structure of the paper is as follows. We first describe the related work in Section 2 and then detail our study and data collection in Section 3. We discuss our analysis and results in Section 4 and Section 5, respectively.

Following this, we discuss the implication, insights and possible applications of our work in Section 6. In Section 7, we discuss the important issue of safeguarding workers' rights as new sensing technologies are developed in the future of work era [4]. Finally, we discuss the limitations of our work in Section 8 and present some concluding remarks in Section 9.

2 RELATED WORK

The majority of prior research related to workplace promotion is focused on an individual's likelihood for being promoted. This topic is studied in relation to both job performance and other related factors. For example, researchers find related factors, such as, personal characteristics, psychological attributes and education level are more related to being promoted than simply performance on the job [5, 22, 71, 88]. Researchers use personality characteristics, job attributes and psychological information to train machine learning models to predict whether an employee is likely to be promoted. Long et al. [44] use demographic (e.g., gender, date of birth, etc.) and job features (position, department, position type etc.) to predict if an employee (N=71132) is promoted or not using a Random Forest model reporting an AUC of 0.96. Other research [91] reports a correlation between work related interactions (i.e., workplace blogs, assigning tasks, downloading work related files, etc.) and online social connections with employee promotion and retention. Work related interaction is strongly predictive and correlates with promotion and retention. The authors [91] collect data from an internal social network platform used by the company and train a logistic regression model to predict promotion and resignation of N=104 employees.

There is little in the way of work examining transitions around promotion specifically pre-post promotion behavioral and physiological changes reported in this paper. Prior work investigating changes around job promotion purely use self-reports associated with general health, psychology, happiness and other metrics. Boyce et al. [10] use the British Household Panel Survey (BHPS) data collected annually over 16 years and report that promoted individuals suffer from deterioration of psychological wellbeing. Johnston et al. [36] use the Household, Income and Labour Dynamics in Australia (HILDA) survey across 9 years to analyze promotion effects. The authors [36] find that while promotion improves job security, pay perception and job satisfaction, there is negligible effect on general health and happiness. Similar to Boyce et al., the authors report that promotion negatively affects the mental health of workers. However, several studies show that working in low-ranked jobs also leads to poor mental wellbeing [21, 29, 73] and social functioning [72]. Because promotion is often followed by a rise in job rank, the findings from these studies suggest that promotion may in fact lead to an improvement in overall health. Along the same lines, researchers also find that promotion may reduce the probability of heart disease by up to 13 percent over a period of 15 years [3]. In a study of N=871 employees, de Lange et al. [15] find that after promotion, employees report an increase in job autonomy and work engagement. With respect to "adjustment" after promotion, Kramer et al. [40] show that employees need to create new communication relationships and skill sets to become comfortable in new roles. Other prior literature examines gender differences in behavior after promotion. Nyberg et al. [58] report that self-rated health decreases for both males and females after promotion. However, promoted females in the study report greater health decreases in comparison to males. Johnston et al. [36], on the other hand, report that promotion impacts the mental health of young males more severely than that of females [36].

There is a growing interest in future of work research. The IMWUT community is well placed to advance this new field of research given its history of developing passive sensing devices and behavioral studies. For instance, prior studies using passive sensing find associations between passive mobile sensing and behavioral markers of mental health [13, 32, 33, 66, 81, 83, 85] and personality [53, 84], among other things. A study by Obuchi et al. [59] goes on to show the promise of passive sensing data by linking it with brain functional connectivity. The authors obtain an F1 Score of 0.79 when classifying N=105 students' functional connectivity between two regions of the brain by just using mobile sensing data. Specific to future of work research, a work on job performance using mobile sensing [52] shows that mobile sensing can coarsely predict workplace performance. The authors collect mobile sensing and self-report data from N=554 workers in different industries across 60 days. The researchers train a machine learning model to identify if a worker is a higher or a lower performer using the self-reported labels and mobile sensing data, obtaining an AUC score of 0.83. In another work, Muralidhar et al. [54] study job performance of 100 students at a hospitality school by collecting their interview video. Participants are asked to perform a reception desk based role with "real" clients (i.e., the researchers) as a means to gauge how they would perform in the job. The researchers [54] annotate the recorded video of these interactions using five independent coders/raters. Using verbal and non-verbal features extracted, the researchers show that they are able to obtain an R^2 of 0.25 for inferring perceived performance and soft skills in the reception desk scenario. Also, Swain et al. [75] use passive sensing to characterize daily activities and personalities of N=603 information workers. They report that fewer location visits, batched phone-use, shorter desk-sessions and longer sleep duration are related with better organizational citizenship behavior and increased task proficiency among the workers.

While there is a growing interest in assessing the workplace and job performance particularly using traditional self-reports, there has been no work to the best of our knowledge that studies changes in behavior and physiology in the workplace after promotion using mobile sensing.

3 METHODOLOGY

In what follows, we discuss our study design, mobile sensing approach, ground-truth, demographic information of the workers in this year long study and the features used in the analysis.

3.1 Study Design

The Tesserae study [51] recruited N=750 information workers across different companies in the USA. All the participants are in the study for a year and respond to several surveys, as detailed in [51]. Individuals participating in the study are provided with a Garmin Vivosmart 3 [45] wearable, a continuous mobile sensing app based on StudentLife [82] and a set of Bluetooth beacons to be placed at their home, on their work desk, in their office and on their keychain. The study protocol is fully approved by the Institutional Review Board (IRB). While we discuss important and relevant information from the Tesserae project as it relates to promotion, please refer to Tesserae study [51] for full details of the study design, participants, data collection, etc. Participants in the Tesserae study are instructed to maintain data compliance percentage of 80% to be eligible for monetary remuneration. In this paper, we study N=141 workers from the complete cohort of N=750 who are promoted during the period of the study as well as the remaining 609 workers who are not promoted. Out of 141 promoted workers, 48 work at a multinational consultancy company, 44 work at a multinational technology company. The remaining 49 participants work in a software company, a university and at various other small companies.

3.2 Demographic Descriptors

With respect to demographic information of the promoted participants, 32% are female, 50% are male and the gender of remaining 18% of workers is unknown because they are associated with the blinded set of participants used in evaluation of the Tesserae study. 48 participants are under the age of 30, 66 are between 30 to 50 and 2 participants are above the age of 50. In terms of highest education attainment, 43% of the participants have a college degree and 33% have a postgraduate degree (Master's or Doctorate). The remaining participants have attended some form of college or graduate school (perhaps, programs that do not award degrees). Tenure wise, the majority of the participants have been with their current employer for two years or more. Figure 1 shows the distribution of this demographic information.

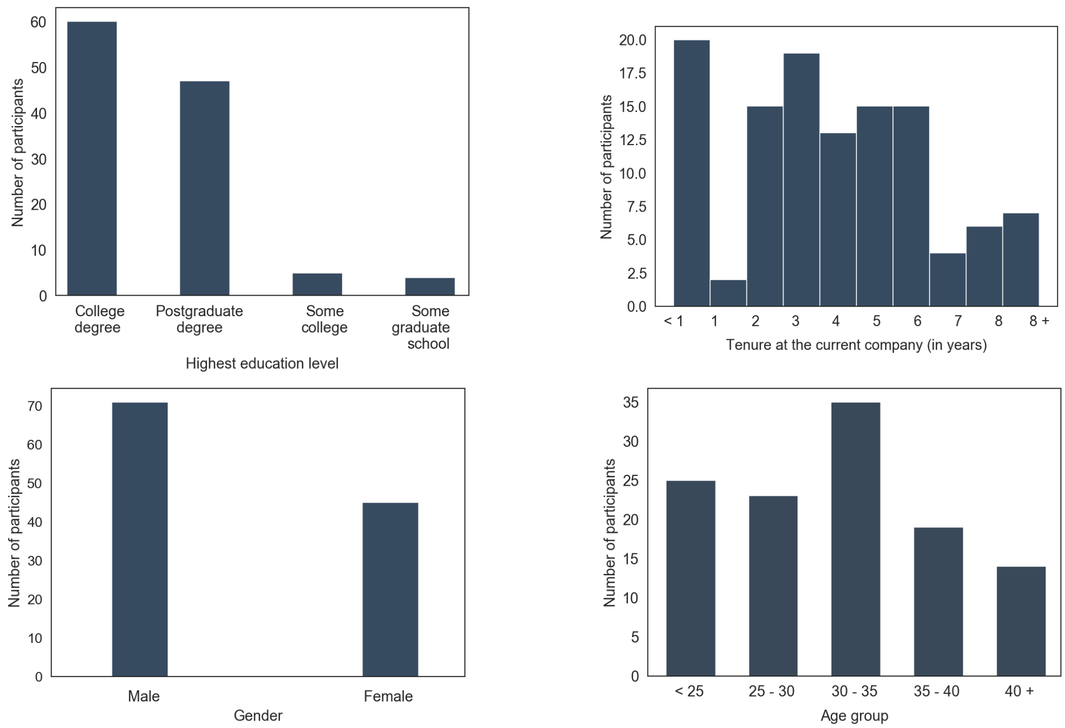

Fig. 1. Demographic information of the promoted participants.

3.3 Promotion Ground Truth

The ground truth for the analysis is collected through participant surveys and self-reports. At the end of the year long study, participants complete exit surveys. Two of these surveys specifically ask: *Were you recently promoted?* and *When were you promoted ? (month and year)*. Participants also go back through their calendar in order to submit significant life events (e.g., illness, marriage, vacation, promotion) that happened during the year long study period. Participants who choose to share the significant life events are compensated extra for doing so. The ground truth used for this paper uses responses to both the exit survey and life events survey associated with promotion. The exit survey asks for the month and year of the most recent promotion, while the life events survey asks for the exact day (from their calendar) of any significant event (in our case, promotion). As a result, we know the day of promotion for N=48 participants, while for the majority of the participants (i.e., N=93), we only know the month and the year of their promotion.

3.4 Mobile Sensing System

The mobile sensing app is installed on the participants iOS/Android phone and tracks participant's phone usage, location, physical activity and Bluetooth interactions with Bluetooth beacons (at work, home, office, keychain) and runs passively without any user interaction in the background of their phone. The mobile app also syncs with the Garmin wearable streaming data from the wearable to the phone. Through the wearable we collect various physiological and behavioral data, such as, sedentary duration, motion intensity, sleep, stress, steps and heart rate data. The Garmin wearable also connects to its own companion phone application which streams information to its web backend in order to compute daily summary information. Workplace activity and interactions related information are collected with the help of Bluetooth beacons (Gimbal). The passive sensing app implements the Gimbal's API to capture proximity or co-location of Bluetooth beacons. This acts as a proxy in our analysis for number and duration of interactions, breaks and time spent at desk/home/work. For full details on the collection system, see [51].

3.5 Features

The features that we use in this paper are mostly based on prior literature discussed in the Related Work section. We list the generated features in Table 1. We generate daily summary features of activity (e.g., sedentary duration, distance, steps), heart rate (e.g., average heart rate, heart rate variability), stress (e.g., average stress duration), workplace behavior (e.g., number of interactions, time spent at desk) and sleep (e.g., sleep duration, Rapid Eye Movement (REM) sleep duration, duration of wake-ups). Each participant is given Bluetooth beacons to place at their home, at their work desk and on their key-chain. A Bluetooth beacon is also placed at the entrance of their respective workplace. Based on whether the signals of the respective Bluetooth beacons are detected by the participant's phone (and for how long), we infer the time spent at work, time spent at home, time spent at their work desk, and number and duration of breaks away from their work desk. Workplace interactions are inferred based on the strength of the detected Bluetooth signals: if participant A's phone detects a Bluetooth signal of participant B with greater signal strength (and likewise participant B's phone detects participant A's Bluetooth signal with greater signal strength), we count that as an interaction (if it lasts for more than a threshold of a few minutes). With respect to stress measure, it is provided by the participant's Garmin wearable. It ranges from 'rest stress' with values from 0-25, 'low stress' with values from 25-50, 'medium stress' with values from 50-75, and finally 'high stress' with values from 75-100. According to the Garmin specification [76], inference of the stress measure is computed by taking into account participant's activity level, respiration rate, HR and HRV during baselines [76]. Typically, when there is an increase in HR, drop in HRV lower than the resting state baseline, and respiratory rate is low relative to the HR, the Sympathetic Nervous System (SNS) which is responsible for physiological responses related to heart rate dominates, leading to a stress state [76]. Sleep related data are collected via the wearable and the phone [49]. In addition to the total sleep duration, the wearable provides us duration relevant to several sleep stages such as awake period, deep, light and REM sleep period. We also calculate daily sleep debt based on the difference in the ideal amount of sleep a participant gets in the prior week.

We divide the day into several "epoch" periods to better understand a worker's physiological and behavioral sensing data over different periods of the day for modelling purposes: epoch 0 (representative of entire 24 hour day), epoch 1 (12 am - 9am; night time, typically when people sleep), epoch 2 (9am - 6pm; day time, typically when people work) and epoch 3 (6pm - 12am; evening time, typically when people leave work and go home or elsewhere). We also consider other time periods based on Bluetooth beacons, such as, time spent at their desk (while at work), time spent not at their desk (while at work), and time spent while not at work (i.e., before they arrive and after they leave). Features related to activity, stress, sleep, steps and heart rate allow us to analyze the effect of promotion on wellbeing, while workplace behavior, distance travelled and phone usage based features are important from the perspective of capturing changes in employee engagement, communication and socialization. Prior work has shown that promotion can impact all these aspects of wellbeing [10, 15, 36, 40], therefore, we use all of these features in our modelling and analysis. With regards to the units of the features, unless otherwise stated, all duration based features are measured in seconds and distance based features in meters. Exceptions to this are time spent at work and desk, which are measured in minutes and total sleep duration as well as sleep debt features, which are measured in hours. Heart rate uses the standard beats per minute unit and HRV is based on Root Mean Square of the Successive Differences (RMSSD) and Standard Deviation of N-N intervals (SDNN/SDANN) [46].

Table 1. Features generated from the mobile sensing system: from the phone, wearable and beacons.

Feature Category	Features
Activity	Sedentary duration, physical activity, motion intensity
Workplace activity	Time spent at desk, time spent at work, number of times participants leave their desk (for durations of 5, 15 and 30 minutes), arrival time to work, departure time from work, duration spent on breaks
Workplace interactions	Number of unique participants, duration spent on interactions, percent of time spent alone, percent of time spent in different interactions (interacting with only one person, with one or more person, with two or more person, or three or more person)
Distance	Total distance travelled, average distance from home, number of locations visited
Stress	Duration of high/medium/low stress, average stress levels
Sleep	Sleep duration, daily sleep debt, duration of deep/light/REM sleep, duration of wake-ups
Steps	Number of steps, walking/running durations
Heart rate	Average heart rate, heart rate variability
Phone usage	Number of unlocks, usage duration

4 ANALYSIS

We discuss our modeling and analysis of physiological and behavioral patterns associated with promotion periods in what follows. We present our results in the next section.

4.1 Analyzing Changes in Behavior and Physiology

We start our analysis by exploring the changes in behavior and physiology of the participants after promotion. We do this by performing a paired Wilcoxon signed rank test [86] on the sensing data of the participant up to 60 days before and 60 days after the month of promotion. We use Wilcoxon signed rank test because the data points are non-normally distributed. Wilcoxon signed rank test compares two related samples to detect whether there's a difference between them after an intervention (in our case, the intervention is promotion.) As mentioned in the Introduction section we select a 60 day epoch period for analysis in order to maximize the number of participants we can include in the analysis across the year based on the amount of data we have from all the promoted workers. We know the day of promotion for N=48 participants, while for the majority of the participants (i.e., N=93), we only know the month and the year of their promotion. Our analysis is based on N=141 of these participants who reported being promoted during the period of the study. Because of this, we exclude the month of the promotion from the analysis. Let us explain that last sentence. If a participant reported being promoted in February, for example, we exclude all their data for February from our analysis irrespective of which specific day they are promoted on. We do this because we do not know the exact day of promotion for a majority of the participants (i.e., N=93) and so the only way we can avoid data before promotion (i.e., non-promoted period) "leaking" into data after promotion (i.e., promoted period) is by removing data associated with the promotion month which we know for all N=141 of these participants. This means that non-promoted period includes data just before the month of promotion (i.e., until the end of January in our example of promoted in February) and promoted period includes the data after the month of promotion (i.e. from the start of March). We report result of the analysis using four criteria: 1) the overall behavioral and physiological change among all the promoted workers; 2) the difference in behavior and physiology between promoted workers after promotion and non-promoted workers; 3) the behavioral and physiological changes associated with gender; and finally 4) the changes associated with high and low performers [52].

4.2 Predictive Classification Model

For the final part of the analysis, we build a predictive classification model that utilizes the sensing data from workers' phone, wearable and beacons (associated features discussed in Table 1) to determine whether the data belongs to promoted or non-promoted periods. Our aim is to be able to detect promotion events by leveraging machine learning technique. We include both the promoted workers and non-promoted workers in this analysis.

For promoted workers, we label the data after the month of their promotion as belonging to promoted class, and data before the month of their promotion as belonging to non-promoted class. Similarly, in case of non-promoted workers, we label all their data as belonging to non-promoted class. Then, we design a model that can detect which period the sensing data belongs to. For the purposes of modeling, we treat the problem as a multivariate time series classification task. It could be the case that multiple sensing streams (e.g., arrival time to work, number of steps, sleep duration) interplay with each other and if we can capture these variations, we would be equipped with better cues to detect promotion based on the sensing data. For each stream, we create a daily summary time series spread over 7 days (i.e., a week) for each of the promoted and non-promoted periods. We make predictions on each weekly grouping of data and once we have the entire month's prediction (i.e., 4 weeks prediction), we use majority voting technique to identify whether that month belongs to a promoted or a non-promoted period. Our machine learning model makes use of a recent technique in time series classification which the authors report to have outperformed other state of the art time series classification algorithms [16].

The approach used by ROCKET (Random Convolutional Kernel Transform) [16] is influenced by the success of Convolutional Neural Network (CNN) [42]. Although CNNs are mostly known for their performance in image based tasks [26, 70], they also perform well when used in time series classification [18, 19]. Researchers believe that they could be effective in time series classification because the convolutional kernels capture patterns in the input time series [18]. A kernel refers to a matrix that is convoluted with the input time series through a sliding dot product [18]. The result is called feature map, another matrix, which is used for classification. Kernels can capture complex patterns and shapes reflected in the feature map. Feature maps are the result of applying the kernels to an input. In essence feature maps are a rendition of the input showing how much of the pattern represented by the kernel are present in the input time series. However, unlike CNNs where the convolution kernels are typically learned, ROCKET uses random convolutions. What this means is that, the kernel weights are not learned, they are randomly generated.

Random convolutions are not novel innovations; they have been used in prior works [34, 35, 62, 68]. In fact, researchers have suggested that random convolutions may be advantageous for small datasets where learning better kernel representation is difficult [34, 90]. ROCKET performs random convolutions over the time series data to extract patterns which are then used by a linear model to make predictions. This approach provides significant gains over the state of the art performance with huge reductions in time and computation power required [16]. ROCKET kernels have random length, dilation, padding, weights and bias [28]. The only hyperparameter for ROCKET is the number of random kernels to generate. For more detail on ROCKET, please refer [16]. In our implementation of ROCKET we extend the model to take into account the multivariate nature of our dataset. We do this by performing a multi-channel 1 dimensional convolution over the time series data. Multi-channel here refers to the number of sensing streams. We perform a one dimensional random convolution over all the sensing

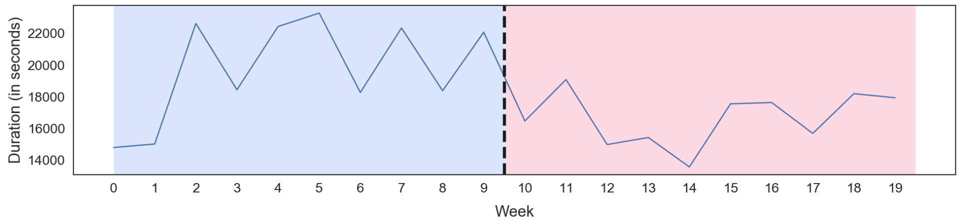

(a) Light sleep duration changepoint

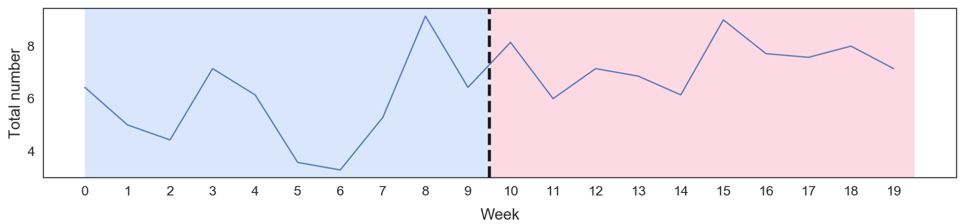

(b) Number of 15 minute breaks changepoint

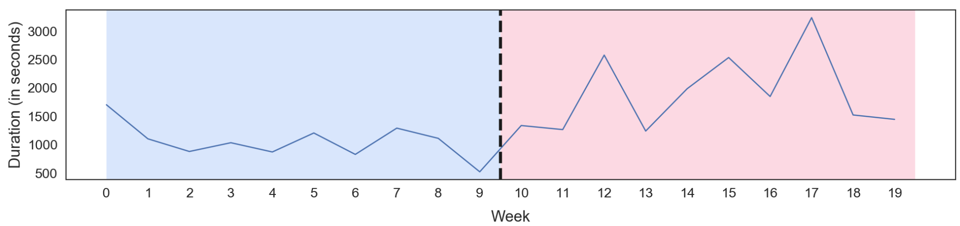

(c) Stress duration while at work changepoint

Fig. 2. Changepoint detection for several sensing streams (viz. sleep, breaks, stress). The figure shows changepoint detected for three different promoted workers for 'light sleep duration', 'number of 15 minute breaks' and 'stress duration at work' streams during one of the weeks in the month of their promotion. The vertical dashed line denotes the changepoint identified – the blue and red shaded regions represent the period before and after the changepoint, respectively. The sensing time series figures clearly show that there is a significant variation in the light sleep duration (i.e., decrease), number of 15 minute breaks (i.e., increase) and stress duration at work (i.e., increase) following the changepoint. Note, the changepoint time series selected are represented on many other examples and are just shown here to illustrate the visible change in physiological and behavioral patterns.

streams grouped by 7 days. Thereafter, we select features using Sequential Forward Selection (SFS) [20] approach for classification.

5 RESULT

In what follows, we present our results with regards to behavioral and physiological change in workers after promotion (i.e., promoted period). We report results associated with specific subgroups of interest including gender and high/low job performers. We also discuss our results on detecting job promotion using mobile sensing.

5.1 Behavioral and Physiological Change after Promotion

5.1.1 Within Promoted Workers. We provide initial insights into behavioral and physiological differences captured through mobile sensing by comparing features from up to 60 days period after the month of promotion and up to 60 days before promotion. We use paired Wilcoxon signed rank test for this purpose because it allows us to compare paired data for each promoted worker before and after promotion. We show the significant behavioral and physiological changes in Table 2. Note, all the reported changes are significant with p-value of less than 0.10 after correcting for multiple comparisons using false discovery rate using the Benjamini Hochberg (FDR-BH) procedure [6].

Table 2. Changes in behavior and physiology of workers after promotion. The following table lists significant changes in behavior and physiology we find for all the participants after promotion. Epochs refer to grouping of different periods of the day; epoch 0: 24 hours (whole day); epoch 1: 12 am - 9 am (night/early morning); epoch 2: 9 am - 6 pm (working hours); epoch 3: 6 pm - 12 am (evening). Note, all the reported changes are significant with p-value of less than 0.10 after correcting for multiple comparisons using false discovery rate using the Benjamini Hochberg (FDR-BH) procedure [6].

Category	**Period**	**Change in pattern**
Activity	Epoch 0	Total physically active duration reduces after promotion Time spent in vehicle decreases
	Epoch 2	Time spent in vehicle decreases
Workplace activity	Epoch 0	Increase in number of 15 minutes break Decrease in number of 30 minutes break
Workplace interactions	Epoch 0 / 2 / 3	Number of interactions decreases Duration of interactions decreases Number of unique participants in conversation decreases
Distance	Epoch 2	Fewer number of unique locations visited
Stress	Epoch 0-3	Decrease in high and low stress duration Increase in rest stress duration
	Epoch 0 / 2 / 3	Increase in stress variability
Sleep	At night	Increase in sleep duration Increase in REM sleep duration
Steps	Epoch 0 / 3	Decrease in step count
Heart rate	Epoch 1	Increase in heart rate
	Epoch 2 / 3	Decrease in heart rate
	At desk	Decrease in HRV while at desk
Phone usage	Epoch 1	More time spent on phone Increase in the number of unlocks
	Epoch 3	Increase in the number of unlocks
		Wilcoxon paired test; p-value < 0.10 (FDR-BH)

As shown in Table 2, we find that the workers are less physically active after being promoted. Given the fact that most of our cohort comprises information workers requiring knowledge work (i.e., problem solving, cognitive skills) rather than physical work, they likely spend longer periods of time sitting and working which results in less physical activity. This is corroborated in the table by a reduction in step count as promoted workers move through the ranks. On the other hand, we find that wellbeing related factors such as stress and sleep improve after promotion. Findings from prior research on promotion and wellbeing are inconclusive; some researchers show that promotion can lead to better health outcomes, while others find promotion leads to deteriorating health outcomes. In a panel data of N=2,681 workers, Johnston et al. [36] report finding small yet positive association with job stress and promotion (i.e., as people get promoted, their job stress also increases). The authors [36] do not find any significant effect of promotion on general health and happiness. Boyce et al. [10] find that people suffer significant deterioration in their psychological wellbeing after promotion. However, Karasek et al. [38] argue that higher occupational levels that are achieved as a result of promotion lead to reduced stress because employees have more autonomy and control of their work which helps them mitigate the high demands of their job. In another line of research, job satisfaction is shown to be strongly associated with employee wellbeing [69], and because promotion typically leads to an increase in job satisfaction [23, 67], it could be inferred that promotion leads to better wellbeing through the mediating role that it plays in increasing job satisfaction.

Mobile sensing (HRV from the wearable) and location information (from the beacons) allow us to measure deeper contextual information associated with workers not readily available to prior researchers; that is, we can analyse a worker's behavioral and physiological signals while they sit at their desk, for example. We find that although the stress reduces for promoted workers across different periods of the day, their heart rate variability (HRV) decreases while they are at their desk. Heart rate variability is a commonly accepted biomarker of stress and generally a lower HRV is associated with higher stress [39]. We believe that the decrease in HRV while at the work desk is potentially pointing to the fact that their stress is increased while they are working at their desk maybe because of the added responsibility of their new role after promotion. However, when considering different periods of the day (i.e., epochs) as well as the entire day, we find that promotion actually leads to positive health outcomes in terms of stress and sleep. Perhaps this is because the job rewards (job flexibility, control, autonomy, salary etc.) actually help people cope with the increased demands and responsibility and maximize their overall health when considered in relation with the entirety of the day or other periods other than while working at their desk.

In terms of workplace activity, we find that there is an increase in the number of breaks of shorter duration and a decrease in breaks of longer duration. It is not clear why this would be the case. One possible explanation is that the promoted workers are highly engaged at work as a result of changing demands in their new roles. Prior research shows that workplace promotion typically leads to an increase in employee performance and productivity [63, 67]. This might relate to another one of our findings – people visit fewer places during the working hours once promoted in comparison to the period before promotion. We posit that workers might be more engaged with new work activities as a result of promotion, which is discussed in the literature [15].

Workplace interaction with colleagues appears to go down after promotion, in both the numbers of interactions as well as the duration of interactions. This is an interesting finding. In addition, after promotion people tend to interact with fewer people. There are multiple interpretations of this result. For one, we observe that promoted employees reduce the number of longer breaks, which might imply that they are engaging more in work items and as a result reducing time spent interacting with other workers. Another interpretation of this result is that promoted workers simply experience changes in communication and engagement with former peers (e.g., individuals, group members, supervisors, managers). The literature indicates that non-promoted workers (for instance, after getting passed over for promotion) may be less attached to their co-workers, hence decreasing the overall interaction for promoted employees [27]. In a survey of N=20 promoted employees, Kramer et al. [40] find that employees experience a sense of social isolation from peers and supervisors after getting promoted.

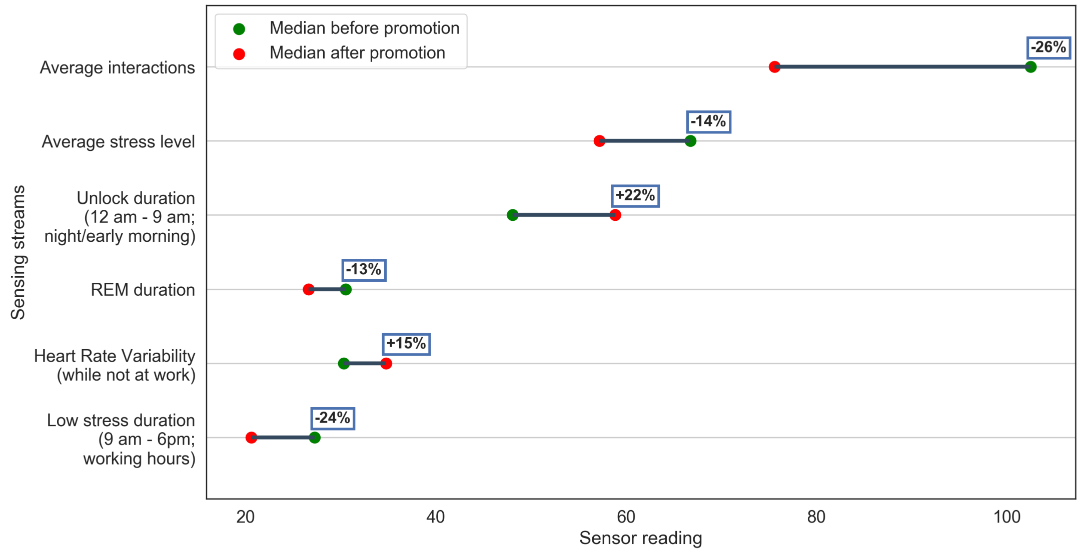

Fig. 3. Comparison of changes after promotion. The figure shows a number of the sensing streams along with the percentage change in them after promotion. Comparisons are shown between median values of the streams. The green dot represents the sensor reading before promotion and red dot, after promotion. We can see that some streams such as average stress level decrease, while others such as HRV increase. All the duration are represented in minutes.

After promotion, employees report feeling isolated because "they were no longer one of them". It makes sense that communication channels, conversational engagement and other workplace interaction are disrupted once an individual is promoted.

Finally, we find that after promotion, workers spend more time on their phone during the night/early morning period. We also capture more number of unlocks during night/early morning and evening periods, both of which are periods after work. Figure 3 shows a number of the significant features that change along with the percentage change after promotion. Comparisons are shown between median values of the streams in the figure. The green dot represents the sensor reading before promotion and red dot after promotion. We can observe that some streams such as average stress level decrease, while other streams such as HRV increase. Our results find significant common patterns across the cohorts we study. While the literature supports a number of our observations we do not claim that these observed behavioral and physiological changes are generalizable.

5.1.2 Between Promoted and Non-Promoted Workers. From the previous analysis, we learn that promoted workers' physiology and behavior changes after promotion. However, we only took within-person changes into account. In what follows, we study the significant differences in behavior that show up when we compare promoted participants with non-promoted participants (i.e., studying between-person differences).

Similar to our prior analysis, we observe up to 60 days of data after promotion period for promoted workers, and similarly 60 days of data for non-promoted workers. We then perform a between-groups comparison using Wilcoxon rank-sum test (also known as Mann-Whitney U test) [47]. The Wilcoxon rank-sum test is a non-parametric version of independent T-test, allowing us to compare two independent groups. We find several significant differences between the promoted and non-promoted workers, as listed in Table 3.

We find that promoted workers are less physically active than non-promoted workers. They have a comparatively lower physically active duration and take fewer steps throughout the day. Non-promoted workers seem to travel more and also venture out further away from their home in comparison to promoted workers. Furthermore, promoted workers spend more time by themselves and at their work desk. This supports our finding from the within-person analysis that we discussed before. As mentioned in Section 5.1.1, one reason behind being less physically active could be straightforward: the vast majority of our study participants are information workers, therefore, a rise through the rank might coincide with spending more time stationary at their desk or in meetings which result in less physical activity. We also find that promoted workers sleep longer and have lower average stress throughout the day. Not surprisingly (since they have lower average stress), they have a higher HRV as well. Most of the results we obtain from this between-groups comparison corroborates our findings from the earlier section. We learn that there are significant differences not just within the behavior of the promoted workers but also between the promoted and non-promoted participants. In section 5.4, we leverage our findings from these two analyses to train a machine learning model that can learn patterns from these differences.

Table 3. Differences in patterns of non-promoted/promoted workers. The table below lists the differences in behavior and physiology of non-promoted/promoted workers (after they get promoted). There are 141 promoted workers and 609 non-promoted workers in our analysis.

Group	**Differences in patterns**
Promoted	Time spent in vehicle is lower during overall day. Physically active duration is lower throughout the entire day. More time spent alone and at work desk while in the workplace. Higher rest stress duration and lower average stress. Higher sleep duration. Average heart rate variability while at desk and also while not at work is higher. Total number of phone unlocks in the overall day is higher.
Non-promoted	More distance travelled throughout the day. Travel further away from home than promoted participants in the entire day. Lower REM sleep duration. More number of steps throughout the day. Average heart rate at work is lower.
	Wilcoxon rank-sum test; p-value < 0.10 (FDR-BH)

5.2 Subgroup Analysis: Gender and Job Performance

Analysis of promotion on the entire cohort might mask important effects that occur for certain subgroup of participants. In order to investigate the differences within subgroups, we perform further analysis of behavior and physiology change considering gender and job performance.

5.2.1 Gender. Out of 141 participants, there are 45 females and 71 males who are promoted. The remaining 25 participants are blinded (that is, we do not know their gender). Table 4 lists the changes we find in participants based on gender.

Our gender based analysis shown in Table 4 clearly indicates that there are significant differences between females and males in terms of behavior and physiology changes after promotion. This is an interesting result. We find that male workers show an increase in phone usage during night/early morning whereas we do not see

Table 4. Changes in patterns after promotion when considering gender. We observe the change in behavior and physiology of promoted workers based on gender. There are 71 males and 45 females who are included in our analysis. Note, all the reported changes are significant with p-value of less than 0.10 after correcting for multiple comparisons using false discovery rate using the Benjamini Hochberg (FDR-BH) procedure [6].

Gender	**Changes in patterns after promotion**
Female	Stress is more variable during the entire day, in the working hours (9am - 6pm) and during the evening (6pm - 12am). Duration with low stress decreases during all periods (entire day, morning, working hours and evening). Medium stress duration increases during working hours (9am - 6pm). Average workplace interaction duration decreases across the entire day as well as during the evening (6pm - 12am). Total distance travelled increases for the entire day. Heart rate increases during evening (6pm - 12am). Heart rate variability decreases while not at work.
Male	Phone usage during night/early morning (12 am - 9 am) increases. Number of unique locations visited during working hours (9am - 6pm) decreases. Number of unlocks in the entire day, during night/early morning (12 am - 9 am) and during evening (6pm - 12am) increases. Average stress level in the entire day, during night/early morning (12 am - 9 am) and during evening (6pm - 12am) decreases. Total step count in the evening (6pm - 12am) decreases. REM sleep duration increases. Number of workplace interactions and unique participant decreases throughout the day (overall day, working hours and evening). Variation of heart rate while at the work desk increases.
	Wilcoxon paired test; p-value < 0.10 (FDR-BH)

similar changes for promoted female workers. Male workers show a decrease in the number of unique locations visited during working hours. For female workers, we find that the total distance travelled increases when we consider the entire day. The most significant changes for female workers are related to stress and heart rate variability. While we find that the average stress level decreases for male workers, females experience an increase in medium stress duration during working hours. Nyberg et al. [58] report that female workers' self-reported health deteriorates at a higher rate in comparison to male workers after getting promoted. Prior studies report that female workers with same position in a company as male workers have higher demands [80], less financially rewarding jobs [8, 57], larger sacrifices in private life [56] and lack of social rewards [31] which ultimately leads them to suffer more stress in jobs with higher occupational status in comparison to males. We also find that the female workers' HRV while not at work decreases. We do not find this for males. Previous work shows that female workers are more exposed to work-family conflicts and 'double burden' [17, 57, 80], which might be the reason why they experience a decrease in HRV while they are not at work. Interestingly, the number of workplace interaction seems to decrease for both genders, possibly reflecting that both males and females experience increased engagement in the workplace, or, workers experience a change in communication, regardless of gender.

5.2.2 Job Performance. In this subgroup analysis, we first make an attempt to uncover the association (if any) in our dataset with respect to job performance and promotion. We then analyze the behavioral and physiological changes.

As discussed in Section 3, the Tesserae study [51] collects self-reports along with mobile sensing data from N = 750 participants. One of the many self-reports collected from the participants includes job performance metrics every 3 days for the first 60 days of the year long study. We want to see if the self-reported performance metric is an indicator of future promotion. The participants self-report their performance using four metrics/surveys: *counterproductive work behavior (CWB)* [64, 65], *organizational citizenship behavior (OCB)* [14, 64], *in-role behavior (IRB)* [9, 87] and *individual task performance (ITP)* [9, 30]. For our analysis, we take an average of the participant's self-reported performance scores for these four metrics.

Table 5. Changes in patterns of higher/lower job performers. The table below lists the changes in behavior and physiology of higher/lower job performers after they get promoted. There are 72 higher performers and 44 low performers in our analysis.

Performance Class	**Changes in patterns after promotion**
High Performers	Number of unique locations visited during working hours (9am - 6pm) and overall day decreases. Stress has more variability during overall day, working hours (9am - 6pm) and during evening (6pm - 12am). Average interaction duration in evening (6pm - 12am) decreases. Average heart rate in overall day, during working hours (9am - 6pm) and during evening (6pm - 12 am) decreases. Average HRV increases in overall day. Number of unique participants in interactions with during the overall day increases, while it decreases during working hours (9am - 6pm) and evening (6pm - 12am).
Low Performers	Stress has more variability during working hours (9am - 6pm) and during evening (6pm - 12am). Stress duration decreases in overall. Number of unique participants goes down in the entire day, working hours (9am - 6pm) and evening (6pm - 12am). Heart rate increases during night/early morning (12am - 9am).
	Wilcoxon paired test; p-value < 0.10 (FDR-BH)

We cluster the participants into higher and lower performers based on the four self-reported performance metrics proposed in [52]. In this approach [52], multiple iterations of K-means clustering is run with initial centroid set to the maximum values of OCB, IRB and ITP survey scores whereas the minimum value of CWB survey score is used. The idea here being, OCB, IRB and ITP are positive performance metrics, whereas CWB is an indicator of negative workplace behavior (i.e., the higher the aggregated value of CWB response, the lower the job performance). Hence, K-means clustering with K=2 and initial centroid set to maximum value of positive indicators of job performance and minimum value of negative indicators of job performance should ideally lead to a division of participants into two classes: "high performers" (those that have high OCB, IRB and ITP but low CWB) and "low performers" (high CWB, but low OCB, IRB and ITP). After we cluster the workers into these two classes (viz. high performer and low performer), we investigate whether there is any correlation between the two performance groups and promotion. The question that drives this exploration is *"Is prior job performance an* *indicator of future promotion?"*. We find that the Mathew's Correlation Coefficient (MCC) [50] on performance class (higher/lower performers) and being promoted is -0.06. This indicates that there is no relationship between performance class and being promoted. For what it's worth, 72 of our 141 participants fall in high performer's class while 44 are in the low performer's class. The remaining 25 participants are blinded; we do not know their performance metrics. In order to investigate the influence of individual performance metrics on promotion (as opposed to the broad higher/lower performance class based on clustering), we fit a logit model with the performance metrics and performance class as independent variables and being promoted (or not) as dependent variable. However, again we do not find any statistically significant relationship. All this leads us to conclude that self-reported performance is not reflective of future promotion in our dataset.

Finally, we investigate the within-person changes of high performers and low performers after getting promoted. Table 5 shows the significant changes we observe. Perhaps what is more surprising about our findings is that higher performers have an increase in number of unique participants that they interact with in the overall day. These higher performers show an improvement in their HRV and heart rate. Stress reduces for both the groups after promotion but there is an increase in stress variability. For low performers, we find that heart rate increases during night/early morning, after leaving work for the day.

5.3 Changes in Individual Streams

We extract numerous features from our sensing dataset sourced from workers' phones, wearables and Bluetooth beacons. Before we discuss our predictive classification model, we want to explore if individual sensing streams change around the time of promotion. We use offline changepoint detection to study this. Changepoint detection [2] is an approach used to identify whether there is a change in a sequence of observations over time. Aptly named, a changepoint is a point at which there is an abrupt variation in the time series data. In offline changepoint detection [79], changes are identified by taking into account the entire dataset in a retrospective manner. Online changepoint detection [2] is used to detect changes in real time settings. Because we are exploring the streams after the collection of the dataset, we use offline changepoint detection to look back in time to recognize where the change happened. Note, that this is a search based technique; it is not using machine learning.

We first aggregate all the available sensing streams by week to create a weekly time series for each stream and then we pass each of the time series streams through binary segmentation search method [25]. Binary segmentation based search method works by trying to identify a single changepoint in the entire dataset in the first instance and then by breaking the time series up into further smaller splits to come up with the most significant changepoints. We use weekly summarised time series of each sensing streams for up to 20 weeks including the month of the promotion. We then investigate the composition of streams that change the most during the month of the promotion (i.e, the streams that have changepoints during that period). In Figure 2, we show three examples of changepoint detected on three separate streams (viz. light sleep duration, number of 15 minute breaks and stress duration at work). We observe that there is a significant change in each of the streams following the week where changepoint is identified. The vertical dashed line in the time series figures denotes the changepoint – the blue and red shaded regions represent the period before and after the changepoint, respectively. These sensing time series figures clearly show that there is a significant variation in the light sleep duration (i.e., decrease), number of 15 minute breaks (i.e., increase) and stress duration at work (i.e., increase) following the changepoint. Note, the changepoint time series selected are represented on many other examples and are just shown here to illustrate the visible change in physiological and behavioral patterns. For the entire cohort, we report the result by categorizing all the sensing streams (i.e., feature categories) into 5 groups: heart rate and heart rate variability (HRV), stress, sleep, workplace behavior (includes workplace activity and interactions) and activity (includes everything else: distance, activity, phone usage). In Figure 4, we show which of these groups change the most during promotion. As can be observed in the figure, stress based streams change the most (i.e., on average, 22% of the changepoints are identified through stress based streams). Activity (21%), sleep (20%) and HR/HRV (20%) based streams also change during the period of promotion. Streams related to workplace interactions and desk activity (17%) do not have as many changepoints in comparison to these more significant streams. In Table 6, we list a number of the individual sensing streams within each groups that change the most during promotion.

Table 6. Streams that change the most. In this table, we list some of the individual streams that changed the most during the period of promotion.

Category	**Features**
Activity	Total step count Calories burned Highly active duration
Workplace behavior	Time spent at work Average break duration Departure time from work
Stress	Duration of high stress Duration of medium stress Duration of low stress
Sleep	REM sleep duration Daily sleep debt Total sleep duration
Heart Rate	Average HR at desk Average HR not at work Average HRV at desk

5.4 Predictive Classification Model

A key goal of our study is to train a machine learning model such that it can detect the promotion period (whether promoted or non-promoted) that the sensing data is associated with. The exploratory analysis that we performed in the earlier section hinted towards the fact that passive sensing streams change during or after the period of promotion. With a predictive machine learning model, we investigate if using all these changes and variations of multiple streams together can give us a good performance. As mentioned in Section 4, we treat the problem as a time series classification task. Basically, for each promoted and non-promoted period, we first group the data into a 7 days time series. We extract features from the 7 days time series and feed it into the model. Once we have the predictions for 4 weeks, we perform a majority voting over all the 4 individual predictions to say whether the corresponding month belongs to a promoted period or a non-promoted period. In case of a tie (i.e., 2 weeks being predicted as non-promoted period and 2 weeks as promoted period), we break it by predicting class 0 (i.e, non-promotion), which is to say that we always require a majority in order to detect promotion. As we may not be able to detect promotions if we use a time series generated over a larger period of time because the important patterns may then be masked by having data aggregated over a longer period, we decided on using weekly time series instead of monthly and do majority voting over them for a month during testing. Weekly time series provides us a better granularity than a monthly time series allowing us to extract patterns that provide a more “cleaner” picture of promotion related events.

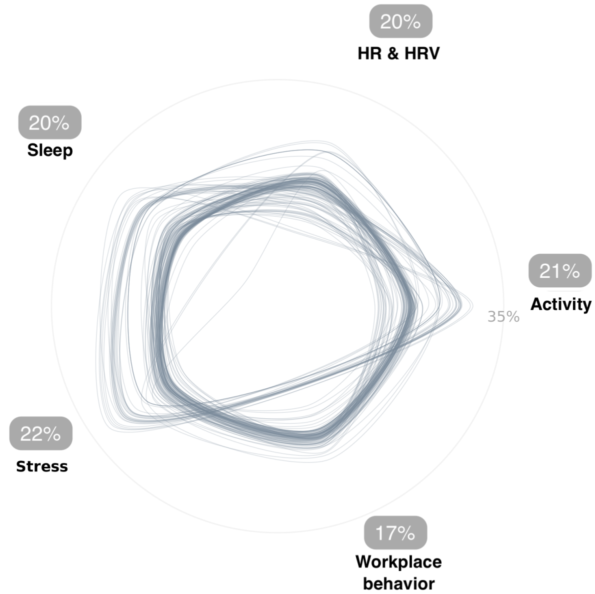

Fig. 4. Composition of different streams. The figure shows the composition of different groupings of streams towards changepoint detection. Each blue colored radial in the image represents a participant in our study. The closer the radials are towards a particular group, the more changepoints were detected by that group of streams for that particular person. We find that in average, 22% of changepoints are detected on Stress related streams, 21% on Activity related streams, followed by 20% on Sleep, 20% on HR & HRV and 17% on Workplace behavior related streams.

The process begins by first extracting features using 250 random kernels. The length of the kernels we use is randomly selected from {3, 5, 7} with equal probability. Other parameters are initialized as mentioned by Dempster et al. [16]; weights are sampled from normal distribution, bias is sampled from a uniform distribution b $\sim \mathcal{U}(-1, 1)$, dilation is sampled on an exponential scale $d = \lfloor 2^x \rfloor, x \sim \mathcal{U}(0, A)$ where $A = \log_2 \frac{l_{\text{input}} - 1}{l_{\text{kernel}} - 1}$ and padding is applied randomly with equal probability. Stride is always set to one. We implement ROCKET with the help of the PyTorch library [61] by performing a 1D convolution, where the number of channels is equal to the number of features, and the sequence length for the time series is equal to 7 (since we are extracting features for a week at a time). From the resultant feature maps, ROCKET [16] then computes two aggregated features – the maximum value (which is equivalent to global max pooling [60]) and the proportion of positive values. Overall, there are 500 features at the end. Let us provide an example of the process. For instance, we start with a 3D tensor, say, of size (432, 133, 7). Here, 7 refers to the number of columns representative of each day of week. Since we are predicting for a week at a time, we create a time series of 7 days so that we can extract features based on each week's data. The row size in the tensor refers to the number of features which we suppose to be 133 in this case. Let us also assume that we have 432 number of weeks available in total of all the participants. Thereafter, we generate 250 random kernels of size 3, 5 or 7 where each have equal probability of occurrence. The weights, bias, padding and dilation are also randomly generated based on the settings mentioned earlier. Once we generate 250 kernels, we then perform 1D convolution over all these 432 different arrays of size (133, 7), using the randomly generated kernels. After performing convolution operations between each kernel and the weekly arrays, we obtain feature maps. From each feature map, we generate two aggregated features – the maximum

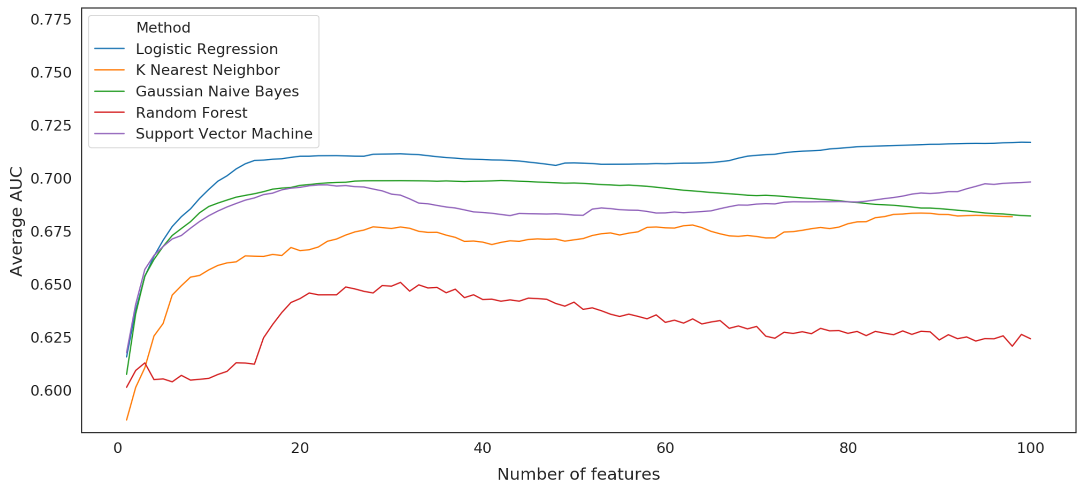

Fig. 5. Performance of models based on feature selection. The figure shows comparison of 5 different Machine Learning models with sequential forward selection. The x-axis is the number of features and the y-axis is the average AUC obtained as a result of leave subjects out cross validation. We can see that all the algorithms show an improvement in their performance as the number of features increase. The figure above shows a particular snapshot of feature selection when number of features is equal to 100. Eventually the models reach saturation point and there is a gradual drop in performance. The highest performance is obtained by Logistic regression.

value and the proportion of positive values. So basically, after performing 1D convolution using every kernel, we extract two features per kernel. At the end we have an array of size (432, 500), where 432 is the number of rows representative of each week of the participants and 500 is the number of columns representative of the features that we generated from the feature map (since we started with 250 kernels, we have 500 features each for every week – twice the number of kernels.) From an input tensor of size (432, 133, 7), we end up with an array of size (432, 500) after feature extraction. We then perform feature selection over these 500 features.

As suggested in [16], we normalize the data per sample before performing convolution operations. Participants were instructed to have 80 percent compliancy in order to qualify for monetary remuneration. However we still have numerous missing values in the dataset – some participants with more and some with less [48]. We perform imputation in order to handle these missing values; we use backward filling and forward filling based on the day of the week and iterative imputation based on Random Forest (also known as MissForest [74]). Then Sequential Forward Selection (SFS) approach is used to select features. Figure 5 shows the performance of different models as SFS proceeds with feature selection. Note that the figure shows a snapshot of one of the iterations limited to only first 100 features. Among the five different models that we try (viz. logistic regression, support vector machine, gaussian naive bayes, random forest and k nearest neighbor), we obtain the best performance with logistic regression. Logistic regression leads to an average AUC of 0.72 with 100 features. In Figure 6, we show a comparison of the performance of these different models when used with and without features generated by ROCKET.

To avoid overfitting and data leakage issues, we perform Leave 10 Subjects Out Cross Validation, which means that we do not use the same subject's records split in train and test set at the same time. Each participant has several months of data (as mentioned earlier, for promoted workers, we do not use the month of the promotion in the analysis), some of them belong to class 0 (referring to non-promoted period) and the remaining belong to class 1 (referring to promoted period). In case of the non-promoted workers, all their data points will be labelled as 0, while the promoted workers will have a mix of class 0 and class 1 labels, depending on when they were promoted. As we have a larger number (609) of non-promoted workers than we have the promoted ones (141), we divide the non-promoted workers into folds each consisting of 100 non-promoted participants. We then repeat the predictive analysis for each fold. Note that we have 241 participants in each iteration of analysis – 100 varying non-promoted participants and 141 fixed promoted ones. This division allows us to handle two things: firstly, it allows us to deal with data imbalance issues because with this division we have similar number of promoted and non-promoted workers as opposed to having large majority of non-promoted workers (i.e., 609) if we consider all of them. Secondly, using different groupings of participants makes the analysis more robust as we are not just relying on a certain group of participants. In addition, we perform another set of predictions to further validate the generalizability of our model. We divide the dataset into stratified K folds where K = 75. Here each fold has 10 participants consisting of a combination of promoted and non-promoted workers. As a result of using stratified K fold, the distribution of promoted and non-promoted workers in each fold is similar to that of the entire dataset. At every iteration, we train on 74 folds and perform testing on the remaining fold. In this manner, we evaluate our model in such situations where we may have a large number of non-promoted workers but a fewer number of promoted workers. In essence, we make predictions on an imbalanced dataset because the held-out stratified fold has similar class distribution as that of the entire dataset which is heavily imbalanced. The result is Leave 10 Subjects Out Stratified Cross Validation. With Logistic Regression, we obtain an average AUC of 0.68 with this procedure, which is not too far off of the result we obtain from our earlier analysis (i.e., AUC of 0.72).

We use ROCKET generated features in our predictive models. As discussed previously, ROCKET has only one hyperparameter: the number of random kernels. The number of random kernels essentially defines the number of features extracted from the given signal. ROCKET extracts two features per feature map. Therefore, there will be twice the number of features as there are the number of random kernels. We extract 500 features using 250 kernels. In figure 7, we show the distribution of learned weights when we do not use SFS for feature selection and also the distribution when we select top 100 features based on SFS. We can clearly see that SFS leads to coefficients that are non-zero. The coefficients, as shown in the figure, are obtained when we train a Logistic Regression model based on features selected by SFS. This shows that feature selection is promising and it helps isolate the key features that can improve the performance of the model.

6 DISCUSSION

In this section, we broadly discuss our approach and the implications and opportunities that it hints towards. We examine the effect of promotion on employee behavior and physiology using mobile sensing consisting of phones, wearables and Bluetooth beacons. We focus on short term effects that are brought about in physical activity, workplace behavior, wellbeing (stress, sleep), physiology (heart rate) and miscellaneous other features as a result of promotion. Subgroup analysis allowed us to consider gender differences and job performances (viz. high and low performers). Results show that at least within our cohort, promotion leads to reduced stress, more engagement at work and decrease in physical activity. We perform a between-groups analysis to investigate the difference in behavior between promoted participants (after promotion) and non-promoted participants. We also find that some of the effects of promotion are different for males and females. We report the findings with regards to higher and lower job performers and changes in behavior and physiology after promotion.

We believe that our findings present implications and opportunities for future work and applications. First, we show that it is possible to use passive sensing data from mobile devices to dive deeper into workplace promotion and its effects in an objective manner. Although we only concern ourselves with promotion based changes, passive sensing can help us better understand how people react to changes which might not just be limited to promotion (e.g., job changes and role transitions). Understanding the effect of certain job changes or role

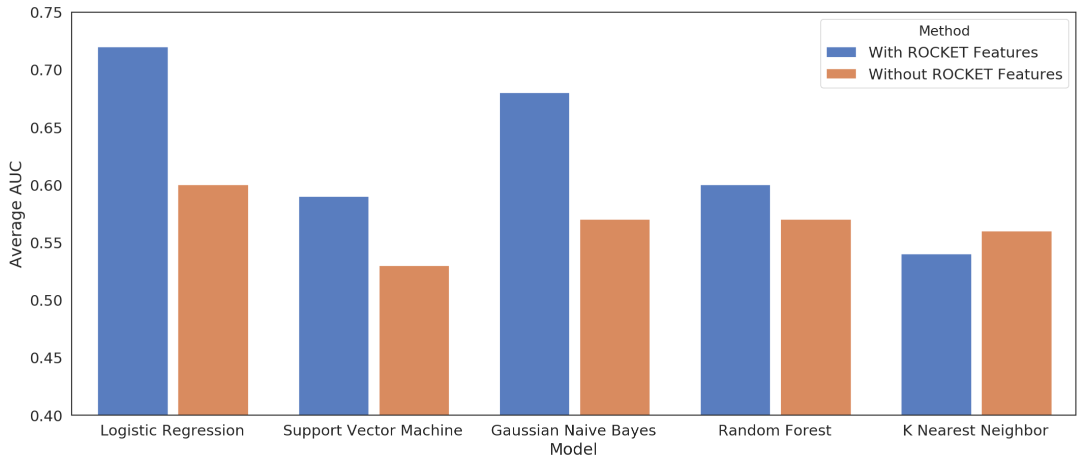

Fig. 6. With and without the use of ROCKET. The figure shows the performance of machine learning models when they're used with and without the features generated by ROCKET. "Without ROCKET features" refers to the original features. We see that features generated by ROCKET improve the performance compared to original features except in the case of K Nearest Neighbor classifier. Logistic regression outperforms other models in terms of performance.

transitions can help in improving workplace wellbeing. This could be achieved through interventions from management or by simply helping workers keep track of changes in their behavior and life as a result of role change. In addition, learning how a particular person reacts to role changes may help the company in personnel management because then they can take into account the person's specific situation or behavior in consideration. Management can understand the differences in individual worker's behavior, physiology and how to help them better deal with things such as changes in communication, getting used to new routines among other things. These are all associated with reacting to change in the job. In contrast, workplace sensing can also offer the ability to study how workers initiate the change themselves. An interesting avenue that this opens up is with regards to job crafting [89]. Job crafting refers to the proactive steps or actions taken by employees utilizing opportunities available to them in order to customize their jobs. For instance, workers can change their tasks (e.g., changing the scope or how they perform it) and interaction with peers (e.g., nature or extent of interactions). Prior research shows that job crafting has positive effects on employee's psychological wellbeing [7], work engagement and performance [78]. With the help of workplace sensing, employers can understand how people with promotion focus craft their job. Equipped with such knowledge, organizations can create appropriate policies to help foster a similar approach for other employees. Prior researchers [11, 43, 77] show that individuals who are focused on being promoted are more likely to craft their job, compared to individuals who are not, as it can lead to status gain, personal growth and success [11, 43, 77]. Ultimately, all this information allows management to retain and motivate workers by offering them an environment to grow while at the same time recognizing that there are individual differences in people.

Next, we explore whether we can detect promotion using passive sensing data. We do this via a machine learning based approach. Our pipeline consisting of ROCKET (for feature extraction), Sequential Forward Selection (for feature selection) and Logistic regression (predictive classification model) is able to achieve an AUC of 0.72. However, as a result of performing convolution operation on the time series signals, we loose interpretability of the model. Needless to say, it is complicated trying to interpret how multivariate time signals interplay with each other in order to result in the given performance based on the features obtained after performing the convolution

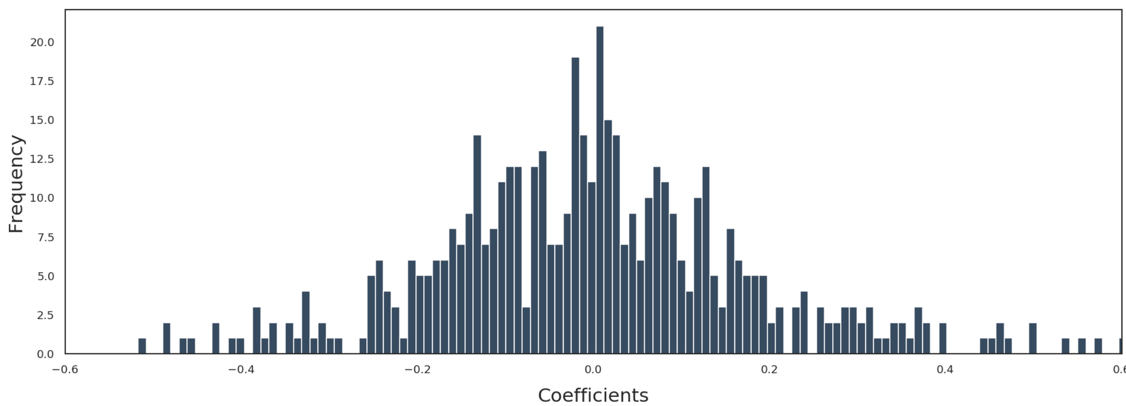

(a) Without Sequential Forward Selection

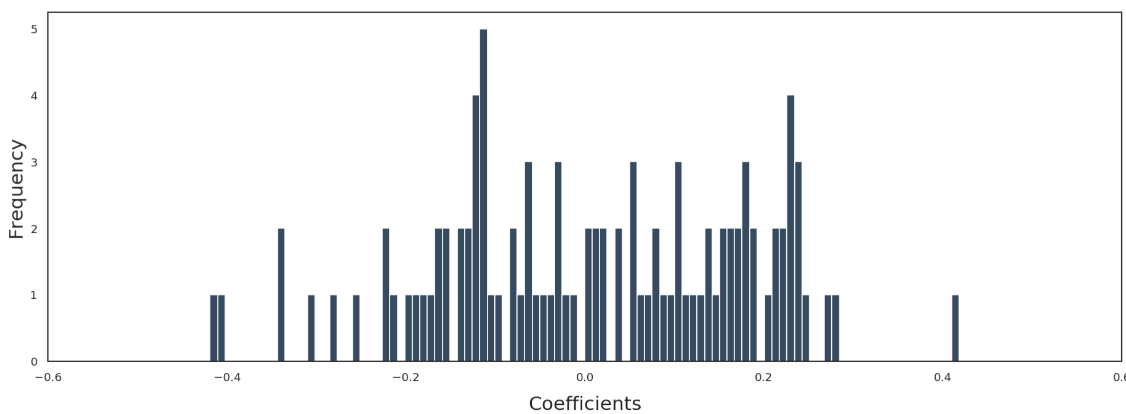

(b) With Sequential Forward Selection

Fig. 7. The figures above show the distribution of coefficients of the Logistic Regression model when trained on features obtained from ROCKET. (a) is the distribution of the coefficients when SFS is not performed. We can see that there are quite a number of values that are close to zero. (b) is the distribution of coefficients after performing SFS. The learned coefficients are non-zero, as is expected for important features in the case of logistic regression model.

operations. Also, with our classification model we used 100 features as obtained from SFS. But as we can see in figure 5, the performance somewhat plateaus for a while after about 20 features until it finally increases again once we hit 80 features mark. It might be possible to use some early stoppage criteria during training so that we can stop the training early to reduce the complexity of the model and still end up with similar performance.

Our findings show that passive sensing from mobile devices is capable of detecting promotion based changes in the workplace. In terms of our predictive classification model, it shows that a model trained on passively sensed data of a group of workers can, with some acceptable level of error, detect promotion events for other workers as well. From a practical implementation perspective, we could have, for example, a model pre-trained on a group of worker's non-promoted and promoted period's data and use it to detect promotion events for a new worker. Once we have enough data available from a new worker (in our study we use 7 days long time series), we could run it through the model, and it would let us know whether the employee was promoted during that period or not.

This is just one of the use-cases of the predictive classification model. We could also envision an online machine learning model that is learning and updating itself in real time settings based on the passive sensing data.

7 PRIVACY FOR FUTURE OF WORK

Privacy and security issues need to be considered with the utmost importance as more and more sensing technology is developed for the workplace. New agency programs on future of work [4] are funding groups to study privacy in the context of new technology in the workplace. Information workers use enterprise applications (e.g., a firm's social media, texting, communications and email systems), video conferences and swipe cards to enter and leave buildings – these are all examples of enterprise technologies that are widely deployed in industries today. As a result, there are a large number of "information breadcrumbs" captured as workers go about their day to day tasks in the enterprise.

The introduction of new technology in the workplace is typically under the remit of private companies that do not disclose how that data is collected, stored and analysed; or, specifically, how data is use in the case of promotion. An important question is how advanced human sensing technology that observes human behavior (and in our case physiology signals) is going to be safeguarded against abuse and bias of employees. This is a critical question that moves beyond privacy in to workers' rights. How are workers' rights to be protected? As ubiquitous computing researchers, we need to consider the implications of advanced sensing technology we help to develop. It is our opinion that workers will only adopt new and potentially invasive sensing technology if they feel they have complete control over their data (e.g., controlling who has access to it) and get something out of the new technology (e.g., health and wellbeing data, performance data to help improve their productivity). These types of sensor driven interventions might be valued by workers as long as they have control of what data is shared and what is private.

How these emerging future of work technologies impact workers' rights is an open and important area of research with little guidance right now. Many times only employees of the company are aware of what technology is used in the firm. For example, many companies offer health apps with insurance discounts to workers as an option. We understand that our study of using mobile sensing for information workers adds to the burden and pressure on privacy of workers. On the more narrower issue of conducting an ethical study, the participant's data in our study is kept at a secure, central server of the research institution. The study has an IRB and protocols. There is no personally identifiable information available to the researchers. In addition, workers in our study are notified that their mobile sensing data and surveys responses will not be shared with their employers. While these protections are important for a study like ours, they are small in consideration to the broader issue of protecting workers data in the era of future of work. In this section, we do not propose specific solutions. Rather, we raise the issue as being critically important to advances in sensing technology in the workplace that must start from the viewpoint of protecting workers' rights first.

8 LIMITATIONS

While our research findings are interesting, there are a number of limitations associated with our work. The analysis is based on a small sample of information workers spread throughout different companies in the USA. Therefore, it is difficult to determine the degree to which our results generalize to other contexts, occupations and industries. As such, application of specific findings should be made with caution. In addition, it is possible that the participants in our study had the physiological and behavioral changes because of situations unrelated to promotion. Next, in order to maximize the number of N, we take into account upto 60 days before and after promotion, which is a short duration to examine the effect of life events such as promotion. The results might be entirely different when consideration is put on longer term effects. However, prior studies based on self-reports show that the effect can be long lasting [3, 10, 36]. Future studies will consider using objective measures such as passive sensing to investigate if that is still the case when analysis is done on objective data. The proxy for interaction as well as workplace activity that we use in our study is based on the proximity of Bluetooth beacons, which is prone to error. If person A's phone is detecting person B's Bluetooth with strong signal strength (and vice versa) then we infer that as an interaction. But, we cannot know for certain whether it was a face to face interaction or simply collocation. Future researchers on this topic could make use of RFID tags which are more reliable in inferring face to face interactions. In addition, the job performance metrics as well as the promotion ground truths are collected via self-reports which are known to be subject to bias such as recall and social desirability. Furthermore, we do not know if workers report a major or minor promotion. Its entirely up to the workers to decide what they consider to be a promotion and we do not know the specifics of what the promotion entails. Different individuals may interpret promotion differently, although given that we have workers from different companies, a unified way to measure promotion objectively may not be possible. We achieved a good AUC for our predictive classification model utilizing random convolutions. Logistic regression trained on ROCKET based features appears to outperform all the other models. Although using ROCKET based features improve the performance of the model significantly, we end up losing interpretability.

9 CONCLUSION

The impact of promotion on an individual has been discussed in prior work. However, so far, self-reported responses have been used to draw insights related to promotion. In this paper, we presented a passive sensing based approach to collect objective data from phones, wearables and Bluetooth beacons in order to explore behavioral and physiological changes of employees after promotion. We utilized up to 60 days of data prior to promotion and 60 days of data after promotion and reported a number of insightful findings. We trained a machine learning model that can identify whether the sensing data belongs to a promoted period or a non-promoted period with an AUC of 0.72. Our work represents the first time mobile sensing has been used to understand the behavioral impact of job promotion on information workers. We believe our findings pave the way for further research in the future to understand the effects of promotion and other job changes on individuals in the workplace. Understanding employees and how they react to job changes might be useful in order to improve the general wellbeing of the workforce as well as to further future-of-work based research and applications. Finally, we also raised the important issue of securing workers' rights as new technologies for future of work accelerate.

ACKNOWLEDGMENTS

This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IARPA Contract No. 2017-17042800007. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

REFERENCES

[1] Vernon L. Allen and Evert van de Vliert (Eds.). 1984. *Role Transitions.* Springer US. https://doi.org/10.1007/978-1-4613-2697-7

[2] Samaneh Aminikhanghahi and Diane J. Cook. 2016. A survey of methods for time series change point detection. *Knowledge and Information Systems* 51, 2 (Sept. 2016), 339–367. https://doi.org/10.1007/s10115-016-0987-z

[3] Michael Anderson and Michael Marmot. 2011. The Effects of Promotions on Heart Disease: Evidence from Whitehall. *The Economic Journal* 122, 561 (Oct. 2011), 555–589. https://doi.org/10.1111/j.1468-0297.2011.02472.x

[4] Thereza Balliester, Adam Elsheikhi, et al. 2018. The future of work: A literature review. (2018).

[5] K. Barton and R. B. Cattell. 1972. Personality factors related to job promotion and turnover. *Journal of Counseling Psychology* 19, 5 (1972), 430–435. https://doi.org/10.1037/h0033250

[6] Yoav Benjamini and Yosef Hochberg. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)* 57, 1 (1995), 289–300.

[7] Justin M Berg, Adam M Grant, and Victoria Johnson. 2010. When callings are calling: Crafting work and leisure in pursuit of unanswered occupational callings. *Organization science* 21, 5 (2010), 973–994.

[8] Erik Berntson, Linda Wallin, and Annika Härenstam. 2012. Typical Situations for Managers in the Swedish Public Sector: Cluster Analysis of Working Conditions Using the Job Demands-Resources Model. *International Public Management Journal* 15, 1 (Jan. 2012), 100–130. https://doi.org/10.1080/10967494.2012.684026

[9] Walter C Borman and SM Motowidlo. 1993. Expanding the criterion domain to include elements of contextual performance. (1993).

[10] Christopher J. Boyce and Andrew J. Oswald. 2011. Do people become healthier after being promoted? *Health Economics* 21, 5 (April 2011), 580–596. https://doi.org/10.1002/hec.1734

[11] Veerle Brenninkmeijer and Marleen Hekkert-Koning. 2015. To craft or not to craft. *Career Development International* 20, 2 (May 2015), 147–162. https://doi.org/10.1108/cdi-12-2014-0162

[12] Jeanne M Brett. 1982. Job transfer and well-being. *Journal of Applied Psychology* 67, 4 (1982), 450.

[13] Luca Canzian and Mirco Musolesi. 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*. 1293–1304.

[14] Jose M Cortina and Joseph N Luchman. 2012. Personnel selection and employee performance. *Handbook of Psychology, Second Edition* 12 (2012).

[15] Annet H. de Lange, Hans De Witte, and Guy Notelaers. 2008. Should I stay or should I go? Examining longitudinal relations among job resources and work engagement for stayers versus movers. *Work & Stress* 22, 3 (July 2008), 201–223. https://doi.org/10.1080/02678370802390132

[16] Angus Dempster, François Petitjean, and Geoffrey I. Webb. 2020. ROCKET: exceptionally fast and accurate time series classification using random convolutional kernels. *Data Mining and Knowledge Discovery* (July 2020). https://doi.org/10.1007/s10618-020-00701-z

[17] Fabrizio D'Ovidio, Angelo d'Errico, Cecilia Scarinzi, and Giuseppe Costa. 2015. Increased incidence of coronary heart disease associated with "double burden" in a cohort of Italian women. *Social Science & Medicine* 135 (June 2015), 40–46. https://doi.org/10.1016/j.socscimed.2015.04.015

[18] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. 2019. Deep learning for time series classification: a review. *Data Mining and Knowledge Discovery* 33, 4 (March 2019), 917–963. https://doi.org/10.1007/s10618-019-00619-1

[19] Hassan Ismail Fawaz, Benjamin Lucas, Germain Forestier, Charlotte Pelletier, Daniel F Schmidt, Jonathan Weber, Geoffrey I Webb, Lhassane Idoumghar, Pierre-Alain Muller, and François Petitjean. 2019. InceptionTime: Finding AlexNet for Time Series Classification. *arXiv preprint arXiv:1909.04939* (2019).

[20] FJ Ferri, Pavel Pudil, Mohamad Hatef, and Josef Kittler. 1994. Comparative study of techniques for large-scale feature selection. In *Machine Intelligence and Pattern Recognition*. Vol. 16. Elsevier, 403–413.

[21] Jane Elizabeth Ferrie, Martin John Shipley, Stephen Alfred Stansfeld, and Michael Gideon Marmot. 2002. Effects of chronic job insecurity and change in job security on self reported health, minor psychiatric morbidity, physiological measures, and health related behaviours in British civil servants: the Whitehall II study. *Journal of Epidemiology & Community Health* 56, 6 (2002), 450–454.

[22] Gary Fields. 2002. Predicting Potential For Promotion: How The Data In Human Resource Information Systems Can Be Used To Help Organizations Gain Competitive Advantage. *CAHRS Working Paper Series* (01 2002).

[23] Marco Francesconi. 2001. Determinants and consequences of promotions in Britain. *Oxford Bulletin of Economics and statistics* 63, 3 (2001), 279–279.

[24] Michael Frese. 1984. Transitions in Jobs, Occupational Socialization and Strain. In *Role Transitions*. Springer US, 239–252. https://doi.org/10.1007/978-1-4613-2697-7_17

[25] Piotr Fryzlewicz. 2014. Wild binary segmentation for multiple change-point detection. *The Annals of Statistics* 42, 6 (Dec. 2014), 2243–2281. https://doi.org/10.1214/14-aos1245

[26] Ross B Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. 2013. Rich feature hierarchies for accurate object detection and semantic segmentation. CoRR abs/1311.2524 (2013). *arXiv preprint arXiv:1311.2524* (2013).

[27] Fred H. Goldner. 1970. Success Vs. Failure: Prior Managerial Perspectives. *Industrial Relations* 9, 4 (Oct. 1970), 453–474. https://doi.org/10.1111/j.1468-232x.1970.tb00527.x

[28] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press. http://www.deeplearningbook.org.

[29] Joan M Griffin, Rebecca Fuhrer, Stephen A Stansfeld, and Michael Marmot. 2002. The importance of low control at work and home on depression and anxiety: do these effects vary by gender and social class? *Social science & medicine* 54, 5 (2002), 783–798.

[30] Mark A Griffin, Andrew Neal, and Sharon K Parker. 2007. A new model of work role performance: Positive behavior in uncertain and interdependent contexts. *Academy of management journal* 50, 2 (2007), 327–347.

[31] Madeline E Heilman. 2001. Description and prescription: How gender stereotypes prevent women's ascent up the organizational ladder. *Journal of social issues* 57, 4 (2001), 657–674.

[32] Yu Huang, Haoyi Xiong, Kevin Leach, Yuyan Zhang, Philip Chow, Karl Fua, Bethany A Teachman, and Laura E Barnes. 2016. Assessing social anxiety using GPS trajectories and point-of-interest data. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 898–903.

[33] Jeremy F Huckins, Alex W DaSilva, Weichen Wang, Elin Hedlund, Courtney Rogers, Subigya K Nepal, Jialing Wu, Mikio Obuchi, Eilis I Murphy, Meghan L Meyer, et al. 2020. Mental Health and Behavior of College Students During the Early Phases of the COVID-19 Pandemic: Longitudinal Smartphone and Ecological Momentary Assessment Study. *Journal of medical Internet research* 22, 6 (2020), e20185.

[34] Kevin Jarrett, Koray Kavukcuoglu, Marc'Aurelio Ranzato, and Yann LeCun. 2009. What is the best multi-stage architecture for object recognition?. In *2009 IEEE 12th international conference on computer vision*. IEEE, 2146–2153.

[35] Abelino Jiménez and Bhiksha Raj. 2019. Time signal classification using random convolutional features. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 3592–3596.

[36] David W. Johnston and Wang-Sheng Lee. 2013. Extra Status and Extra Stress: Are Promotions Good for Us? *ILR Review* 66, 1 (Jan. 2013), 32–54. https://doi.org/10.1177/001979391306600102

[37] Timothy A Judge and Shinichiro Watanabe. 1994. Individual differences in the nature of the relationship between job and life satisfaction. *Journal of occupational and organizational psychology* 67, 2 (1994), 101–107.

[38] Robert A Karasek Jr. 1979. Job demands, job decision latitude, and mental strain: Implications for job redesign. *Administrative science quarterly* (1979), 285–308.

[39] Hye-Geum Kim, Eun-Jin Cheon, Dai-Seg Bai, Young Hwan Lee, and Bon-Hoon Koo. 2018. Stress and Heart Rate Variability: A Meta-Analysis and Review of the Literature. *Psychiatry Investigation* 15, 3 (March 2018), 235–245. https://doi.org/10.30773/pi.2017.08.17

[40] Michael W. Kramer and Timothy L. Noland. 1999. Communication during job promotions: A case of ongoing assimilation. *Journal of Applied Communication Research* 27, 4 (Nov. 1999), 335–355. https://doi.org/10.1080/00909889909365544

[41] Edgar Krau. 1983. The attitudes toward work in career transitions. *Journal of Vocational Behavior* 23, 3 (1983), 270–285.

[42] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521, 7553 (May 2015), 436–444. https://doi.org/10.1038/nature14539

[43] Nira Liberman, Lorraine Chen Idson, Christopher J Camacho, and E Tory Higgins. 1999. Promotion and prevention choices between stability and change. *Journal of personality and social psychology* 77, 6 (1999), 1135.

[44] Yuxi Long, Jiamin Liu, Ming Fang, Tao Wang, and Wei Jiang. 2018. Prediction of Employee Promotion Based on Personal Basic Features and Post Features. In *Proceedings of the International Conference on Data Processing and Applications - ICDPA 2018*. ACM Press. https://doi.org/10.1145/3224207.3224210

[45] Garmin Ltd. 2019. Activity Tracking. https://buy.garmin.com/en-US/US/p/567813. [Online Access].

[46] M. Malik, J. T. Bigger, A. J. Camm, R. E. Kleiger, A. Malliani, A. J. Moss, and P. J. Schwartz. 1996. Heart rate variability: Standards of measurement, physiological interpretation, and clinical use. *European Heart Journal* 17, 3 (March 1996), 354–381. https://doi.org/10.1093/oxfordjournals.eurheartj.a014868

[47] H. B. Mann and D. R. Whitney. 1947. On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other. *The Annals of Mathematical Statistics* 18, 1 (March 1947), 50–60. https://doi.org/10.1214/aoms/1177730491

[48] Gonzalo J Martinez, Stephen Mattingly, Shayan Mirjafari, Subigya Nepal, Andrew Campbell, Anind Dey, and Aaron D Striegel. 2020. On the Quality of Real-world Wearable Data in a Longitudinal Study of Information Workers. In *WristSense 2020: 6th Workshop on Sensing Systems and Applications using Wrist Worn Smart Devices (WristSense 2020)*. Austin, USA.

[49] Gonzalo J. Martinez, Stephen Mattingly, Jessica Young, Louis Faust, Anind Dey, Andrew Campbell, Munmun Choudhury, Shayan Mirjafari, Subigya Nepal, Pablo Robles-Granda, Koustuv Saha, and Aaron D. Striegel. 2020. Improved Sleep Detection Through the Fusion of Phone Agent and Wearable Data Streams. In *WristSense 2020: 6th Workshop on Sensing Systems and Applications using Wrist Worn Smart Devices (WristSense 2020)*. Austin, USA.

[50] B.W. Matthews. 1975. Comparison of the predicted and observed secondary structure of T4 phage lysozyme. *Biochimica et Biophysica Acta (BBA) - Protein Structure* 405, 2 (Oct. 1975), 442–451. https://doi.org/10.1016/0005-2795(75)90109-9

[51] Stephen M Mattingly, Julie M Gregg, Pino Audia, Ayse Elvan Bayraktaroglu, Andrew T Campbell, Nitesh V Chawla, Vedant Das Swain, Munmun De Choudhury, Sidney K D'Mello, Anind K Dey, et al. 2019. The Tesserae project: Large-scale, longitudinal, in situ, multimodal sensing of information workers. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–8.

[52] Shayan Mirjafari, Anind K. Dey, Sidney K. D'Mello, Ge Gao, Julie M. Gregg, Krithika Jagannath, Kaifeng Jiang, Suwen Lin, Qiang Liu, Gloria Mark, Gonzalo J. Martinez, Kizito Masaba, Stephen M. Mattingly, Edward Moskal, Raghu Mulukutla, Subigya Nepal, Kari Nies, Manikanta D. Reddy, Pablo Robles-Granda, Koustuv Saha, Anusha Sirigiri, Aaron Striegel, Ted Grover, Weichen Wang, Pino Audia, Andrew T. Campbell, Nitesh V. Chawla, Vedant Das Swain, and Munmun De Choudhury. 2019. Differentiating Higher and Lower Job Performers in the Workplace Using Mobile Sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 2 (June 2019), 1–24. https://doi.org/10.1145/3328908

[53] Bjarke Mønsted, Anders Mollgaard, and Joachim Mathiesen. 2018. Phone-based metric as a predictor for basic personality traits. *Journal of Research in Personality* 74 (June 2018), 16–22. https://doi.org/10.1016/j.jrp.2017.12.004

[54] Skanda Muralidhar, Marianne Schmid Mast, and Daniel Gatica-Perez. 2018. A Tale of Two Interactions. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 3 (Sept. 2018), 1–24. https://doi.org/10.1145/3264939

[55] Nigel Nicholson. 1984. A theory of work role transitions. *Administrative science quarterly* (1984), 172–191.

[56] Anna Nyberg, Linda L. Magnusson Hanson, Constanze Leineweber, and Gunn Johansson. 2015. Do Predictors of Career Success Differ between Swedish Women and Men? Data from the Swedish Longitudinal Occupational Survey of Health (SLOSH). *PLOS ONE* 10, 10 (Oct. 2015), e0140516. https://doi.org/10.1371/journal.pone.0140516

[57] Anna Nyberg, Constanze Leineweber, and Linda Magnusson Hanson. 2015. Gender differences in psychosocial work factors, work–personal life interface, and well-being among Swedish managers and non-managers. *International Archives of Occupational and Environmental Health* 88, 8 (March 2015), 1149–1164. https://doi.org/10.1007/s00420-015-1043-0

[58] Anna Nyberg, Paraskevi Peristera, Hugo Westerlund, Gunn Johansson, and Linda L Magnusson Hanson. 2016. Does job promotion affect men's and women's health differently? Dynamic panel models with fixed effects. *International Journal of Epidemiology* (Dec. 2016), dyw310. https://doi.org/10.1093/ije/dyw310

[59] Mikio Obuchi, Jeremy F. Huckins, Weichen Wang, Alex daSilva, Courtney Rogers, Eilis Murphy, Elin Hedlund, Paul Holtzheimer, Shayan Mirjafari, and Andrew Campbell. 2020. Predicting Brain Functional Connectivity Using Mobile Sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 1 (March 2020), 1–22. https://doi.org/10.1145/3381001

[60] Maxime Oquab, Léon Bottou, Ivan Laptev, and Josef Sivic. 2015. Is object localization for free?-weakly-supervised learning with convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 685–694.

[61] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., 8024–8035. http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf

[62] Ali Rahimi and Benjamin Recht. 2008. Random features for large-scale kernel machines. In *Advances in neural information processing systems*. 1177–1184.

[63] Abdul Razak, Sarpan Sarpan, Ramlan Ramlan, et al. 2018. Influence of promotion and job satisfaction on employee performance. *Journal of Accounting, Business and Finance Research* 3, 1 (2018), 18–27.

[64] Reeshad s. Dalal, Holly Lam, Howard M. Weiss, Eric R. Welch, and Charles L. Hulin. 2009. A Within-Person Approach to Work Behavior and Performance: Concurrent and Lagged Citizenship-Counterproductivity Associations, and Dynamic Relationships with Affect and Overall Job Performance. *Academy of Management Journal* 52, 5 (Oct. 2009), 1051–1066. https://doi.org/10.5465/amj.2009.44636148

[65] Paul R Sackett. 2002. The structure of counterproductive work behaviors: Dimensionality and relationships with facets of job performance. *International journal of selection and assessment* 10, 1-2 (2002), 5–11.

[66] Sohrab Saeb, Emily G Lattie, Konrad P Kording, and David C Mohr. 2017. Mobile phone detection of semantic location and its relationship to depression and anxiety. *JMIR mHealth and uHealth* 5, 8 (2017), e112.

[67] Saharuddin and Sulaiman. 2016. The Effect Of Promotion And Compensation Toward Working Productivity Through Job Satisfaction And Working Motivation Of Employees In The Department Of Water And Mineral Resources Energy North Aceh.

[68] Andrew Saxe, Pang Koh, Zhenghao Chen, Maneesh Bhand, Bipin Suresh, and Andrew Ng. 2011. On Random Weights and Unsupervised Feature Learning. *Proceedings of the 28th International Conference on Machine Learning*, 1089–1096.

[69] Emiliano Sironi. 2019. Job satisfaction as a determinant of employees' optimal well-being in an instrumental variable approach. *Quality & Quantity* 53, 4 (Jan. 2019), 1721–1742. https://doi.org/10.1007/s11135-019-00835-3

[70] M. Sornam, Kavitha Muthusubash, and V. Vanitha. 2017. A Survey on Image Classification and Activity Recognition using Deep Convolutional Neural Network Architecture. In *2017 Ninth International Conference on Advanced Computing (ICoAC)*. IEEE. https://doi.org/10.1109/icoac.2017.8441512

[71] Seymour Spilerman and Tormod Lunde. 1991. Features of Educational Attainment and Job Promotion Prospects. *Amer. J. Sociology* 97, 3 (Nov. 1991), 689–720. https://doi.org/10.1086/229817

[72] Stephen Alfred Stansfeld, Hans Bosma, Harry Hemingway, and Michael G Marmot. 1998. Psychosocial work characteristics and social support as predictors of SF-36 health functioning: the Whitehall II study. *Psychosomatic medicine* 60, 3 (1998), 247–255.

[73] Stephen A Stansfeld, Rebecca Fuhrer, Martin J Shipley, and Michael G Marmot. 1999. Work characteristics predict psychiatric disorder: prospective results from the Whitehall II Study. *Occupational and environmental medicine* 56, 5 (1999), 302–307.

[74] Daniel J Stekhoven and Peter Bühlmann. 2012. MissForest—non-parametric missing value imputation for mixed-type data. *Bioinformatics* 28, 1 (2012), 112–118.

[75] Vedant Das Swain, Koustuv Saha, Hemang Rajvanshy, Anusha Sirigiri, Julie M. Gregg, Suwen Lin, Gonzalo J. Martinez, Stephen M. Mattingly, Shayan Mirjafari, Raghu Mulukutla, Subigya Nepal, Kari Nies, Manikanta D. Reddy, Pablo Robles-Granda, Andrew T. Campbell, Nitesh V. Chawla, Sidney D'Mello, Anind K. Dey, Kaifeng Jiang, Qiang Liu, Gloria Mark, Edward Moskal, Aaron Striegel, Louis Tay, Gregory D. Abowd, and Munmun De Choudhury. 2019. A Multisensor Person-Centered Approach to Understand the Role of Daily

Activities in Job Performance with Organizational Personas. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 4 (Dec. 2019), 1–27. https://doi.org/10.1145/3369828

[76] Firstbeat Technologies. 2014. Stress and Recovery Analysis Method Based on 24-hour Heart Rate Variability. https://assets.firstbeat.com/firstbeat/uploads/2015/10/Stress-and-recovery_white-paper_20145.pdf

[77] Maria Tims and Arnold B. Bakker. 2010. Job crafting: Towards a new model of individual job redesign. *SA Journal of Industrial Psychology* 36, 2 (Dec. 2010). https://doi.org/10.4102/sajip.v36i2.841

[78] Maria Tims, Arnold B Bakker, and Daantje Derks. 2012. Development and validation of the job crafting scale. *Journal of vocational behavior* 80, 1 (2012), 173–186.

[79] Charles Truong, Laurent Oudre, and Nicolas Vayatis. 2019. Selective review of offline change point detection methods. *Signal Processing* (2019), 107299.

[80] Ari Väänänen, May V. Kevin, Leena Ala-Mursula, Jaana Pentti, Mika Kivimäki, and Jussi Vahtera. 2005. The Double Burden of and Negative Spillover Between Paid and Domestic Work: Associations with Health Among Men and Women. *Women & Health* 40, 3 (April 2005), 1–18. https://doi.org/10.1300/j013v40n03_01

[81] Rui Wang, Min SH Aung, Saeed Abdullah, Rachel Brian, Andrew T Campbell, Tanzeem Choudhury, Marta Hauser, John Kane, Michael Merrill, Emily A Scherer, et al. 2016. CrossCheck: toward passive sensing and detection of mental health changes in people with schizophrenia. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 886–897.

[82] Rui Wang, Fanglin Chen, Zhenyu Chen, Tianxing Li, Gabriella Harari, Stefanie Tignor, Xia Zhou, Dror Ben-Zeev, and Andrew T Campbell. 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*. 3–14.

[83] Rui Wang, Weichen Wang, Alex daSilva, Jeremy F Huckins, William M Kelley, Todd F Heatherton, and Andrew T Campbell. 2018. Tracking depression dynamics in college students using mobile phone and wearable sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 1 (2018), 43.

[84] Weichen Wang, Gabriella M. Harari, Rui Wang, Sandrine R. Müller, Shayan Mirjafari, Kizito Masaba, and Andrew T. Campbell. 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, Article 141 (Sept. 2018), 21 pages. https://doi.org/10.1145/3264951

[85] Weichen Wang, Shayan Mirjafari, Gabriella Harari, Dror Ben-Zeev, Rachel Brian, Tanzeem Choudhury, Marta Hauser, John Kane, Kizito Masaba, Subigya Nepal, et al. 2020. Social Sensing: Assessing Social Functioning of Patients Living with Schizophrenia using Mobile Phone Sensing. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–15.

[86] Frank Wilcoxon. 1945. Individual Comparisons by Ranking Methods. *Biometrics Bulletin* 1, 6 (Dec. 1945), 80. https://doi.org/10.2307/3001968

[87] Larry J Williams and Stella E Anderson. 1991. Job satisfaction and organizational commitment as predictors of organizational citizenship and in-role behaviors. *Journal of management* 17, 3 (1991), 601–617.

[88] David A Wise. 1975. Personal attributes, job performance, and probability of promotion. *Econometrica: Journal of the Econometric Society* (1975), 913–931.

[89] Amy Wrzesniewski, Nicholas LoBuglio, Jane E Dutton, and Justin M Berg. 2013. Job crafting and cultivating positive meaning and identity in work. In *Advances in positive organizational psychology*. Emerald Group Publishing Limited.

[90] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. 2014. How transferable are features in deep neural networks?. In *Advances in neural information processing systems*. 3320–3328.

[91] Jia Yuan, Qian-Ming Zhang, Jian Gao, Linyan Zhang, Xue-Song Wan, Xiao-Jun Yu, and Tao Zhou. 2016. Promotion and resignation in employee networks. *Physica A: Statistical Mechanics and its Applications* 444 (Feb. 2016), 442–447. https://doi.org/10.1016/j.physa.2015.10.039