

FIELD

5 The present invention relates to a method and system to detect stress outcome or depression outcome or both in a monitored subject.

BACKGROUND

10 Work-related stress is an increasingly common affliction in all walks of life. When left unmanaged, work stress can inhibit productivity, or even cause psychological distress such as depression.

 There exist initiatives to assess and manage stress. However, some initiatives require participants to install mobile or custom apps or provide regular survey data, making these solutions
15 hard to scale. These apps may carry higher privacy risks as a high amount of personal data is collected for real-time analysis. Other initiatives require sensors to be attached to the subject that is monitored to collect data, which is resource intensive.

 An object of the present invention is to provide an approach to detect the probability of stress or depression or both that addresses the shortcomings of the outlined existing initiatives.

20

SUMMARY OF THE INVENTION

 According to a first aspect of the present invention, there is provided a computer-implemented method of detecting stress outcome or depression outcome or both in a subject, the
25 method comprising: receiving, as input, location data providing information indicative of venues visited by the subject, the location data being provided by an access point situated at the venues; processing the location data against a library of activities, stored in a database, to deduce one or more activities done at each of the visited venues, the deduction being based on an event log recorded at the visited venue; constructing a current movement pattern to capture a sequence of an activity from the
30 deduced one or more activities being done; calculating change values of one or more features by comparing the current movement pattern against a past movement pattern of a sequence for the same activity being done; analysing the change values using a classification algorithm that models a relationship between data from one or more movement patterns and their indicative stress outcome or depression outcome; and providing, as output, a stress outcome or depression outcome or both in the
35 subject.

According to a second aspect of the invention, there is provided a computer-implemented system for detecting stress outcome or depression outcome or both in a subject, the system comprising: at least one microprocessor; at least one non-volatile storage device; at least one computer-readable memory device operatively associated with the microprocessor; and an input/output interface operatively associated with the microprocessor, wherein the computer-readable memory device contains computer-executable instruction code which, when executed via the microprocessor, configures the microprocessor to: receive, via the input/output interface, location data providing information indicative of venues visited by the subject, the location data being provided by an access point situated at the venues; process the location data against a library of activities, stored in the non-volatile storage device, to deduce one or more activities done at each of the visited venues, the deduction being based on an event log recorded at the visited venue; construct a current movement pattern to capture a sequence of an activity from the deduced one or more activities being done; calculate change values of one or more features by comparing the current movement pattern against a past movement pattern of a sequence for the same activity being done; analyse the change values using a classification algorithm that models a relationship between data from one or more movement patterns and their indicative stress outcome or depression outcome; and provide, as output, a stress outcome or depression outcome or both in the subject.

BRIEF DESCRIPTION OF THE DRAWINGS

Representative embodiments of the present invention are herein described, by way of example only, with reference to the accompanying drawings, wherein:

Figure 1 shows a block diagram of a system in accordance with one embodiment of the present invention.

Figure 2 summarises the demographic data used to train a machine learning engine of the system of Figure 1 to predict an outcome of severe stress or depression in a university student.

Figure 3 shows a screen capture of an online portal used to obtain survey data that serves as ground truth to the results of the machine learning engine.

Figure 4 lists the distribution of stress and depression labels for all participants contributing the training and test data of Figure 2.

Figures 5(a) to 5(c) plot different mobility features between normal stress and severe stress groups.

Figure 6 lists details of three different machine learning models for stress and depression cross validated on different student samples.

Figures 7(a) to 7(d) summarise results of cross validating a machine learning engine to detect severe stress trained with the data of Figure 2.

Figures 8 summarises the validation results of the generic stress model of Figure 6.

Figure 9 shows cross validation results for best settings of a depression model built in accordance with one embodiment of the present invention.

Figure 10(a) and 10(b) show results for optimising and validating the depression model.

Figure 11 is a schematic of a server used to host a machine learning engine of the system of Figure 1.

10 DETAILED DESCRIPTION

In the following description, various embodiments are described with reference to the drawings, where like reference characters generally refer to the same parts throughout the different views.

15 The present application, in a broad overview, relates to a computer-implemented method that can detect whether a subject is experiencing a more serious condition of stress or depression. This method may entirely be executed on a single computer server or across several computer servers, with each computer server responsible for executing a segment of the instructions that constitute the method.

20 The underlying operating principle of the method is its ability to quantify a change in a movement pattern and interaction habits of a subject, i.e. obtain a computational value in a format that can be processed by one or more algorithms, and use the computed value to determine whether this change is indicative of the subject being under severe stress or suffering from depression. The input that the method uses includes, but is not limited to, is location data providing information indicative of venues visited by the subject, which provides details on the subject's presence and movement between these venues. The term "location data" is thus a record of the subject's movement and is also interchangeably referred to as "movement data" or "location information" throughout this disclosure. The one or more computer servers may obtain this location data from one or more access points, where there may be overlap between the coverage of each access point so that an access point is not confined to only monitoring a single venue.

30 The location data is processed against a library that contains an archive of stored activities, to deduce one or more activities that are done at the visited venues. A movement pattern is then created from a sequence of these activities. This movement pattern which is derived from the input location data is called a "current movement pattern" because it provides the behavioural data which is being analysed to determine whether the subject is experiencing an increased level of stress or the more serious condition of depression.

The computer-implemented method uses a classification algorithm that models a relationship between data from one or more movement patterns and their indicative stress outcome or depression outcome to analyse whether a subject is under stress or suffering from depression. In one implementation, the stress outcome or the depression outcome has a binary value, i.e. the monitored subject is either under stress or not; or is heading towards depression or not. The model is learnt by training the classification algorithm with data from past movement patterns providing information of one or more activities done at the visited venues and an indicative stress outcome or depression outcome associated with each of these past movement patterns. In one implementation, a first set of data providing movement patterns on various activities over a time frame is obtained; along with its accompanying indicators of stress or depression, gathered (for example) through surveys completed by participants on which the first set of data is based. A second data set is then obtained by calculating change in movement patterns by comparing the current movement against its prior activity over the same time frame; along with its accompanying indicators of stress or depression, gathered (for example) through surveys completed by participants on which both the first and second set of data are now based. The second data set is used to train the classification algorithm, which then is configured to model a relationship between change in one or more movement patterns and their corresponding indicative stress outcome or depression outcome. Multiple second data sets may be used in this training, with each of the second data sets obtained from comparing different instances of the first and second sets of data.

The classification algorithm determines whether a subject is under stress or suffering from depression by analysing change values of one or more features calculated from comparing the current movement pattern against a past movement pattern. "Features" refer to input datasets that the classification algorithm is configured to classify, analogous to classification of input into a computational neural network. These "features" may also be referred to as "mobility features" in this disclosure.

It was found that performing the analysis on change values of features calculated from comparing the current movement pattern against a past movement pattern improved the prediction of whether a subject is experiencing a more serious condition of stress or depression. Using change values of features translates to using change in behavioural data as basis to perform the analysis, in contrast to the approach of using a trained classification algorithm to analyse features directly from the current movement pattern (which translates to current behavioural data).

Figure 1 shows a block diagram of a system 100 in accordance with one embodiment of the present invention. The system 100 has four blocks: a tracking sub-system 102, an activity mapper 103; a feature extractor 104 and a machine learning engine 106. The four blocks, discussed in detail below, may be implemented in a single computer server or each may be implemented in a separate computer server.

The tracking sub-system 102 collects location and group information directly from venue infrastructure, such as one or more access points 108 located in a venue monitored by the system 100. The access points 108 obtain such location and group information through communication with an electronic device 110 that is registered to a monitored subject. For example, the tracking sub-system 102 tracks all electronic devices 110, when connected to a wireless fidelity (WiFi) network every 5 seconds from WiFi signal strengths measured by WiFi access points.

In the implementation shown in Figure 1, the tracking sub-system 102 obtains location data from wireless fidelity (WiFi) signals monitored by the one or more access points 108 through inferring the position from the strongest received signal strength of all electronic devices 110 within range and finding the venue that best matches the received signal from a radio map. The system 100 has one or more access points 108 that may be situated over one or more venues in a monitored area. Each of these venues may be labelled to be associated with a certain activity, e.g. laboratory work in the context of a university.

The tracking sub-system 102 has a WiFi location system 112, which uses real-time location services (RTLS) to extract Receiver Signal Strength Information (RSSI) that can be used to determine for presence of the electronic devices 110. Detection indicates that the party to whom the electronic device 110 is registered has visited a particular venue. The RSSI is the signal strength of each electronic device 110 connected to the access point 108 as measured by the access point 108. These signal strengths decay as the electronic device 110 moves further away from the access point 108. By using RSSI observed by multiple access points 108, the position of each WiFi connected electronic device 110 can be determined using reverse triangulation with heuristics and filters. This approach uses data collected solely from infrastructure (i.e., from each WiFi access point 108) and thus can work across any mobile device platform (e.g. iOS, Android) and does not require installing any client software into the electronic device 110. The subject being monitored is traceable from identifying the electronic device 110 registered to the monitored subject through a unique attribute, such as the MAC address of the electronic device 110.

Server-side location data provides the ability not only to infer individual behaviours but to group interaction patterns non-intrusively. Adopting WiFi location tracking as the sensing medium is advantageous for being a resource already present or installable in buildings, which can be readily tapped, while WiFi capability is already a default feature in mobile devices.

Besides WiFi, other possible mediums to monitor venues visited by a subject include one or more of global positioning system (GPS); other radio-based mediums such as infrared (IR), radio frequency identification (RFID), ultra-wideband (UWB) and Bluetooth low energy (BLE); and non-radio based mediums such as camera, sound or magnetic fields. In addition, one or more of these mediums may be tapped (and even combined with WiFi) to obtain the necessary information of the venues visited by the subject. For instance, the short coming of GPS being possibly less accurate in

indoor settings may be supplemented by one or more of IR, RFID, UWB and BLE based techniques to track the movement of the electronic device 110 registered to the monitored subject. As for BLE, its higher beacon density provides accurate positioning than WiFi. However, continuous tracking of a fixed Bluetooth address over different periods may be difficult due to the Bluetooth address of a device being replaced with random values at specific time intervals.

WiFi is advantageous for allowing continuous identification of the monitored subject since WiFi uses MAC address, a globally unique ID for the electronic device 110. One way to maintain anonymity is to use a one-way hash function that is implemented at the WiFi location system 112, which enables monitoring of all tracked electronic devices 110 at a community-level without exposing their identities.

The tracking sub-system 102 is a component that identifies groups by looking at correlations in electronic devices 110 movement trajectories. The tracking sub-system 102 also has a group detector system 114 which uses one or more algorithms (e.g. Markov Cluster algorithm, MCL) that can process the location information compiled by the WiFi location system 112 to infer conclusions as to whether group data may be present. Such group data refers to whether there is location data with cluster characteristics, namely whether there is location data showing electronic devices, which include the electronic device 110 registered to a monitored subject, belonging to a cluster in the vicinity of monitored venues. The detection of such a cluster of electronic devices may be indicative of interaction of a monitored subject with the group to whom the other electronic devices are registered, providing an option to also utilise location data attributable to the monitored subject during such interaction in the determination as to whether the monitored subject is under severe amounts of stress or suffering from depression. Such a cluster of electronic devices is identified, for example, by detecting whether these electronic devices move together with the electronic device 110 registered to the monitored subject across monitored venues. The location data which shows the monitored subject interacting with a group may be further broken down into subsets, such as a subset for each activity that the monitored subject does with the group; the activity may be deduced from one or more of the duration, time at which the location data is logged and the venue at which the location data is logged. One or more of these subsets may be used as input in the determination as to whether the monitored subject is under severe stress or suffering from depression.

The activity mapper 103 is a component that applies in-situ heuristics to location and group data to map activities that are commonly engaged by a studied population and take places in designated spaces in a studied environment. The activity mapper 103 has a database 116 which stores a library of activities, which is populated with activities that are associated with each of the venues monitored in the system 100. For instance, programming is assumed to be done in a computer laboratory, project discussion is assumed to be carried out in a meeting room, while meals are assumed to be consumed in a dining facility. The activity mapper 103 receives location data 120 from

the tracking sub-system 102 block to determine, from the library of activities in the database 116, the most likely activities that occurred within the span of the location data based on one or more parameters present in each data set, such as a particular location of a detected venue, time threshold of routines calculated from population mean and regular measurements of activity in 15 minute windows.

5 An event log is recorded at each of the monitored venues in the system 100, the event log being a record of actions captured by each of the access points 108, e.g. a time stamp of when an electronic device 110 is detected to have entered or exited a monitored zone of the access points 108. These time entries may be used to provide duration and time as to when an electronic device 110 visits a venue. The venue can be mapped to different activities depending on a specific day and time
10 at which the venue is determined to have been visited from the location data. Broadly, the assignment of various types of activities to deduce one or more activities done at each of the visited venues in the location data 120 received by the activity mapper 103 uses two different conditions: 1) the activity most associated with each location (for that day and time. e.g. seminar room is marked as “Work” during office hours); and 2) if the monitored subject was at that location for at least 15 minutes.
15 Otherwise the activity is marked as in transition.

Activity heuristics 118 were developed in the following steps. First, common predetermined activities were manually assigned to each monitored venue (e.g. {"location": "<building name>_<level number>", "activities": ["seminar", "study", "transit"]}). Then, an activity is decided from a list based on a simple decision-making statement of time and day (e.g. in the context of a university, lectures are
20 conducted on fixed time slots on weekdays). Second, the activity is verified against the monitored subject's schedule. In one approach, the monitored subject may be asked to provide their schedule, or their schedule may be obtained from peers to which the monitored subject belongs. In an automated approach, the monitored subject's schedule may be obtained through synchronising their calendar. A further verification to confirm the deduced activities uses time thresholds of the monitored subject's
25 daily activities averaged over a sampled population to which the monitored subject belongs. For example, an instructor consultation is capped at 1 hour, gym at 2 hours, transition between places at 15 minutes, eating at 30 minutes. In addition, an activity is only assigned to a nomadic electronic device 110 switching connections between different access points 108 if a connection lasts for at least 15 minutes; otherwise labelled as “in transition”.

30 In the deduction of the one or more activities done at each of the visited venues in the location data 120 received, the activity mapper 103 may perform a pre-processing step 122 to obtain missing values, i.e. missing data required to deduce an activity done at a visited venue (which in turn may prevent formulating a sequence of one or more activities to construct a current movement pattern of the monitored subject), using data interpolation, such as AKIMA interpolation.

35 Missing data occurs from situations such as the electronic devices 110 not always being connected to the access point 108, thus constitute to missing values. Such missing values are treated

with AKIMA interpolation. AKIMA spline affects only the curve of neighbouring data points, minimising the error of estimated values used as a substitute to the missing data. The pre-processing noise step 122 may also be used to address fluctuating signals from situations such as detection of an electronic device 110 bouncing from one access point 108 to another during transit. To resolve the issue of fluctuating signals, the final location may be determined by the mode of a 5-minute sliding window.

Features, useful in the determination as to whether there is an increased level of stress or the more serious condition of depression in a monitored subject, are identified from the activities assigned to all data sets within the location data 120. These features are based on activity types; with most features being generalisable to a population pool to which the monitored subject belongs (e.g. activities for a student population would include: studying, attending a lecture, eating), while other features may be specific to the monitored subject's circumstances (e.g. if the monitored subject is a student enrolled in a Software Engineering (SE) course). Such domain-specific features may be verified against a project team schedule, where each schedule entry captures location, date, duration, types of task and attendees. Contradicting entries, such as a detected location which does not match a logged location (for a particular time of the day), are identified as 'unique' task features.

The feature extractor 104 is a component that generates statistical and temporal changes from mapped location traces that are representative of routines and group interactions as mobility features. When applied to analyse location data 120 based on campus activities, the feature extractor 104 of the system 100 identified four main types of activities, primarily "Work", "Workgroup", "Non-work", and "Group", which are common across all students. "Work" refers to events in locations designated for studying, including seminar rooms for lectures, and group study rooms for workgroup-related events. "Workgroup" is thus a subset of 'Work'. Workgroup-related events were also verified with project team schedules. "Non-work" refers to events in recreational spaces such as the campus gym, dance studios, and food courts. Finally, "Group" events occur in locations with group data activity, such as where multiple electronic devices 110 were found to belong to a cluster (using the logic of the group detector system 114). These activities are categorised by the feature extractor 104 into 'General' dataset 124 and 'Domain-specific' dataset 126, as summarised in table 1 below. Features from the domain-specific dataset 126 are specific to a particular project, while features from the general dataset 124 are applicable to everyone in the population pool.

Table 1: Features grouped in sets

Type	Set	Mobility Feature
General	Non-work (Location data)	count_building count_campus, time_campus count_gym, time_gym count_food, time_food count_clinic, time_clinic count_transit, time_transit
	Work (Location data)	count_study, time_study count_seminar, time_seminar count_consultation, time_consultation
	Group (Group data)	count_group, time_group count_smallgroup, time_smallgroup count_mediumgroup, time_mediumgroup count_largegroup, time_largegroup
Domain-specific	Workgroup (Location data)	count_meetings, time_meetings count_pairprogramming, time_pairprogramming count_knowledgeshare, time_knowledgeshare
	Workgroup (Group data)	count_groupmembers, time_groupmembers

The feature extractor 104 extracts features from the location data 120 based on campus activities as described in the following list

1. Number of unique visits per day records the number of different buildings visited.
2. Total time spent on <activity type> & number of times engaged in <activity type> per day consist of the following activities with 15 minutes unit time per activity: on campus (sums up all activities), studying, attending lecture, group meetings, study consultation, transiting, eating, exercising, visiting the clinic. Domain-specific workgroup activities include the types of tasks declared in a project schedule such as pair-programming, knowledge sharing, application design, milestone preparation and unique events.

3. Total time spent being in <group type> & number of times being in <group type> per day Characteristics such as activity frequency, duration, averages and proportions of different activities were extracted as raw features 128.

- In order to process the temporal changes in features 130 and 132, current behaviour of the monitored subject needs to be established to provide a basis for comparison against their past behaviour. The current behaviour of the monitored subject is derived by constructing, from the location data 120, a current movement pattern that captures a sequence of an activity over a period of time for all deduced activities. An example of a mobility pattern (i.e. either a current movement pattern or a past movement pattern) comprises of a sequence of a single deduced activity over a period of time. A mobility pattern for a study activity may capture studying done at time 1, time 2 and time 3, while a mobility pattern for a seminar activity may capture attending a seminar at time 1, time 2

and time 3. A change in mobility pattern for the studying activity compares the samples between time 3 and time 2, time 2 and time 1 for the same studying activity. Similarly, a change in mobility pattern for the seminar activity compares the samples between time 3 and time 2, time 2 and time 1 for the same seminar activity.

5 In one implementation, the current movement is compared against a past movement, where the change in movement pattern of one feature from the current record and the past record is obtained. In the event that there is no current movement for a particular activity, like “time spent on study activities”, the above mentioned AKIMA interpolation produces an estimated value for this activity based on neighbouring data points. Change values are calculated the same way, i.e. a change in an
10 activity that is not present in the constructed current movement is obtained from interpolating neighbouring data points, and the estimated value is treated as baseline for future movement patterns.

The approach of using calculated change in the values of one or more features seeks to establish whether changes in a monitored subject’s movement patterns and interaction habits in reference to themselves are indicative of perceived stress or depression. This is advantageous for
15 providing an advance indication of whether the changes in the movement indicate that the monitored subject is struggling to reorient or adapt before serious consequences occur, such as depression. For example, while it may be intuitive that spending more time in a stressful activity should consequently increase stress, confirming whether this is indeed the case from a determination based on the change in the values of the one or more features against stress or depression indicators of past movement
20 patterns may instead show that the monitored subject is not experiencing a commensurate increase in stress or depression levels or both levels remain the same.

The change values of the one or more features may be calculated in one of several ways. In a first implementation, the change value for each feature is calculated from comparing a same feature from the movement pattern of the same subject $x_{i,j}^u$ over two different intervals, as absolute change
25 ($abs_{i,j}^u$, see Equation 1 below). These movement patterns may be constructed from activities done during each of the two intervals by the monitored subject. In a second implementation, the change value for each feature is calculated from comparing a same feature from both the monitored subject's movement pattern and a movement pattern of a peer population, $x_{i,j}^0$ over two different intervals, as relative change ($rel_{i,j}^u$, see Equation 2 below). This peer population is a population pool that contains
30 data from peers to which the monitored subject belongs. The movement pattern of this population pool may be constructed from the same activity done during the selected interval by the population pool. These change values may be calculated over different time windows which correspond to the frequency over which ground truth data is collected, e.g. in multiples of 3 days for ground truth data collected every 3 days).

$$x_{i,j}^0 = \sum_{v \in [1..N] \setminus u} x_{i,j}^v / N - 1$$

$$\hat{x}_{i,j}^* = \sum_{k=i}^{i+w} x_{i,j}^k \quad i \in [1..K-w], [w] := \{3, 6, 9, \dots, w\}$$

$$\mathbf{abs}_{i,j}^u = \hat{x}_{i+1,j}^u - \hat{x}_{i,j}^u \quad (1)$$

$$\mathbf{rel}_{i,j}^u = (\hat{x}_{i+1,j}^u - \hat{x}_{i+1,j}^o) - (\hat{x}_{i,j}^u - \hat{x}_{i,j}^o) \quad (2)$$

5

The calculated change values attributed to absolute change is shown using reference numeral 130 in Figure 1, while the calculated change values attributed to relative change is shown using reference numeral 132.

10 The change values 130 and 132 are then used as input to the machine learning engine 106. During initialisation of the machine learning engine 106, these change values 130 and 132 and their corresponding indicative stress outcome or depression outcome, serve as training data so that the machine learning engine 106 can model a relationship between change in one or more movement patterns and their corresponding indicative stress outcome or depression outcome. When these change values 130 and 132 are in respect of change values of a current movement pattern against a past movement pattern, they provide the machine learning engine 106 with the behavioural data which is being analysed to determine whether a monitored subject is experiencing stress or depression. Suitable choices for the machine learning engine 106 used by the system 100 include those using a classification algorithm with a machine learning pipeline of feature selection and classification.

20 The machine learning engine 106 is a component that systematically builds detection models for stress and depression and outputs predictions. The machine learning engine 106 has a classification algorithm that models a relationship between data from one or more movement patterns and their indicative stress outcome or depression outcome, i.e. the classification algorithm recognises a relationship between a movement pattern and its indicative stress outcome or depression outcome and is configured to receive data from movement patterns to compute a stress outcome or a depression outcome. In Figure 1, the relationship modelled by the machine learning engine 106 encompasses a change in one or more movement patterns and their corresponding indicative stress outcome or depression outcome. The classification algorithm of the machine learning engine 106 analyses the change values 130 and 132 and provides, as output, an indicative stress outcome 150 or an indicative depression outcome 152 or both in the monitored subject. In providing this output, the classification algorithm may also optionally analyse the raw values 128 besides the change values 130 and 132.

The machine learning engine 106 includes a recursive feature elimination (RFE) module 134 that is configured to perform optimisation by analysing the importance of one or more features. The RFE module 134 searches through the different features in the change values 130 and 132 and selects the most optimal ones for the classification algorithm to provide its output of stress outcome 150, the depression outcome 152 or both.

The parameter tuning block 136 allows for threshold settings of the machine learning engine 106 to be adjusted. For instance, it may be more important to build a classifier with high true positive rate (TPR) to prioritise the accuracy of detecting subjects with signs of severe stress or depression than true negative rate (TN), which detects subjects as being healthy. Default settings for certain machine learning algorithms have the threshold for positive/negative class set at 0.50. The parameter tuning block 136 allows the threshold for the machine learning engine 106 to be set at a lower level, such as 0.45, at the cost of a high false negative rate.

The classification algorithm of the machine learning engine 106 may use a protocol that is based on any one of the following algorithms: Logistic Regression (LR), Support Vector Machine (SVM), and Random Forest (RF). The model 142 used to provide an output of the stress outcome 150 is different from the model 144 used to provide an output of the depression outcome 152.

The classification algorithm may receive one or more of the following input:

i) change values of one or more features derived from location data 120 that comprises information indicative of venues visited by a group that is detected to have interacted with the monitored subject. Activities that the monitored subject does with the group at the visited venues is deduced by the group detector system 114, by confirming the current movement pattern of the monitored subject is the same as the sequence of deduced activities with the group. The calculated change values, derived from comparing this current movement pattern (that factors group interaction) against a past movement pattern of the monitored subject, thus include one or more activities the monitored subject is deduced to have done while interacting with the group.

ii) change values 130. These are change values of features obtained from comparing the current movement pattern of the monitored subject against a past movement pattern of the same subject, see $abs_{i,j}^u$ from Equation 1.

iii) change values 132. These are change values of features obtained from comparing the current movement pattern of the monitored subject against a past movement pattern of a population pool to which the monitored subject belongs, see $rel_{i,j}^u$ from Equation 2.

iv) raw values 128. These are features obtained from the movement pattern of the monitored subject over a period of time.

For the classification algorithm of the machine learning engine 106 to be able to model a relationship between data from one or more movement patterns and their stress outcome or depression outcome, the machine learning engine 106 is trained using a subset of participant data recording

movement patterns of one or more activities and an indicative stress outcome or depression outcome over a period of time. The following paragraphs describe how the machine learning engine 106 can be trained to be able to predict the stress outcome or depression outcome when the monitored subject is a university student. It will be appreciated that the machine learning engine 106 is not limited to being trained to only detect stress or depression outcomes for university students, but can also be used to detect stress or depression outcomes for another population pool, although a different set of training data would be required for the machine learning engine 106 to predict the stress outcome or the depression outcome for this other population pool.

Figure 2 summarises the demographics of the training data used to train the machine learning engine 106 to predict the stress outcome or depression outcome for a university student. It began with *Study_{SE}*, as the main study conducted for 81 days in Fall AY2017 (Academic Year 2017). All 76 participants (Information Systems major, Sophomore) were enrolled in a Software Engineering (SE) module that requires students to work in pre-assigned groups to build a semester-long project. SE groups are made up of 5 pre-assigned members, with each group balanced along multiple dimensions such as gender, GPA, nationality, and programming capability. Consequently, as the groups are pre-assigned and balanced, students find themselves in a situation where they must work with people they do not know and have varying capabilities, personalities, and work styles. Students get to know the project requirements and team grouping on the first day of class, and each team has the full semester to build a fully working software artifact (a fairly complicated web application deployed to a commercial cloud provider) while following the various processes required – these processes require team members to spend equal amounts of time coding, enforce pair programming for coding tasks, require rotation of programming pairs etc. We chose this course for our study as it is continuously cited as a substantially stressful course, mostly due to the pressures of having to work closely with people whom one does not know, and who may not share the same capabilities and aspirations.

Study_{ValidA} (13 participants) and *Study_{ValidB}* (51 participants) are validation studies conducted in Spring AY2017 (36 days) and Fall AY2018 (81 days), respectively. Students ranged from freshmen to seniors, and were enrolled in different majors and modules. In all these modules, students would form groups to work on either semester-long or short-term projects. None of these courses, unlike SE, required intensive juggling between technical and management practice. At the end of our study, 62 (*Study_{SE}*), 11 (*Study_{ValidA}*) and 35 (*Study_{ValidB}*) active students remained. These are the students who contributed at least 80% of all survey data (used for ground truth validation) and attended at least one interview session. We validated from our mid-semester and end-of-semester interview findings that none of our inactive participants had reported experiencing a highly stressful semester.

Each participant filled out a pre-study questionnaire outlining their personality traits using a Big-5 Inventory, their current GPA, and their normal campus routines (e.g. meal breaks, sports, extra-curricular activities). During the study, participants reported their perceived stress levels using a PSS-

4 survey every three days, and a retrospective assessment to measure depression using a PHQ-8 survey approximately every two weeks. Additionally, semi-structured interviews were arranged with participants at recess week and at the end of the semester to better understand their sources of stress, experiences of team-related stress, and how these stresses were managed. Students participated in face-to-face interviews for *Study_{SE}*, and online interviews for *Study_{ValidA}* and *Study_{ValidB}*. We used cross-validated survey and interview data as ground truth. Finally, *Study_{SE}* participants provided access to their SE project schedule (a graded document maintained by all teams to keep track of SE work progress and plans), while *Study_{ValidA}* and *Study_{ValidB}* students provided meeting logs every three days. These schedules and logs included information on meeting dates, duration, and location. To enhance our data collection process, all surveys were embedded and administered from a custom online portal. Our portal, as illustrated in Figure 3, served as the primary platform for survey access and facilitating reminders for participation.

The majority of our data consisted of location and group data, derived from connection logs between participants' mobile phones and the campus WiFi network. We used a WiFi location system that uses RTLS (real-time location services) to extract RSSI information about every connected WiFi device directly from every access point in the environment. These raw RSSI values are then processed, using reverse triangulation with many heuristics and filters, to produce the final location of every connected WiFi device.

For this study, we extracted three months' worth of WiFi localisation data for *Study_{SE}* and *Study_{ValidB}* and one-month data from *Study_{ValidA}*. This localisation data was then processed by a group detector to produce group information. For *Study_{SE}*, this resulted in 310,000 processed location data points and 89,280 processed group data points for all 62 participants each day, which is equivalent to an average of 7 hours' worth of location and 2 hours of group interaction data per participant (i.e., $310,000/62 * 5$ seconds and $89,280/62 * 5$ seconds). Participants were also detected to have visited, on average, 96 unique locations on campus each month. We extracted mobility features from location and group data points on a per day basis and mapped them to stress (PSS-4) and depression (PHQ-8) labels. In summary, the final dataset consisted of 5,022 data points ($62 * 81$ -days) from *Study_{SE}*, 2835 data points from *Study_{ValidB}*, and 468 data points from *Study_{ValidA}*.

Self-reports of PSS-4 and PHQ-8, coupled with verification of students' experiences from two interview findings were the main mechanisms for collecting ground truth. We now describe how we grouped the assessment scores.

PSS-4 is a well-established scale, ranging from 0 to 16. While a score close to 16 suggests high perceived stress, it is not designed as a diagnostic tool. Hence, we referenced work by Wartig et al. (Wartig, S. L., Forshaw, M. J., South, J., and White, A. K. New, normative, english-sample data for the short form perceived stress scale (pss-4). *Journal of health psychology*, 18(12):1617–1628, 2013) that provides norms for an English sample ($N > 1500$, with various ethnicity: White, Mixed,

Black African, and Asian) for PSS-4. In computing a 95% prediction interval, the upper bound ends up being $6.11 + 3.14(1.96) = 12.26$. This implies that a PSS-4 score over 12.26 would place someone above the 95th percentile of the population, supporting an interpretation of 12+ as severe stress. Moreover, the norms in the study place score of PSS-4 = 9+ in the moderate group that is ≈ 1 SD above the mean ($6.11 + 3.14$). A score of PSS-4 = 12+ places someone in the severe category, 2 SD above the mean. Accordingly, scores of 12 and above are grouped as **1: severe stress** (positive class), otherwise **0: normal stress**.

The use of PHQ-8 for a diagnostic tool has a scale with clear cut-offs – minimal: 0-4, mild: 5-9, moderate: 10-14, moderate-severe: 15-19, and severe: 20- 24. We are particularly concerned with students falling in the “moderate-severe” and “severe” depression categories and consider these occurrences as “critical”. Accordingly, PHQ-8 scores of 10 and above are grouped as **1: depressed** (positive class), otherwise **0: non-depressed**.

Figure 4 lists the distribution of stress and depression labels for all studies. The PSS-4 conversion resulted in a distribution of more than 90% normal stress labels for all studies. Prior work [Burnaev, E., Erofeev, P., and Papanov, A. Influence of resampling on accuracy of imbalanced classification. Eighth International Conference on Machine Vision (ICMV 2015), volume 9875, page 987521] suggests that this imbalance in labels, seen in our ground truth, is to be expected as individuals overwhelmed by stress tend to be outliers. Such skewed datasets can lead to poor prediction performance if not corrected. Thus, we addressed this problem by applying SMOTE [Batista, G. E., Prati, R. C., and Monard, M. C. A study of the behaviour of several methods for balancing machine learning training data. ACM Sigkdd Explorations Newsletter, 6(1):20–29, 2004] to synthetically oversample training set data in the severe stress and depression classes.

We explored the feasibility of detecting stress using only coarse-grained location data collected from the campus WiFi network. We developed a set of mobility-driven features and hypotheses based on interview studies with students. Hypothesis testing was performed to statistically validate features that can differentiate students experiencing severe stress. For the ground-truth labels, we placed the students in *Study_{SE}* into severe stress ($n=4$) and normal stress ($n=58$) categories, based on their self-reported average PSS-4 scores. Referring to Figures 5(a) to 5(c) (which plots different mobility features between normal stress and severe stress groups, where the dashed line is for the data for students with normal stress and the solid line is for students with severe stress). The analysis was conducted as follows: First, we visually examined the changes in mobility features over time, between the two groups, as shown in Figure 5(a) by averaging features every three days and plotting them over the study duration of 81 days. We define *Time Point*, T_x as a sample made every 3 days – i.e., $T_{24} = 24 * 3 = \text{day } 72$ of the study. Second, we performed one-way MANOVA to investigate the significance of the multivariate mean effects on different features and ran individual t-tests with Bonferroni correction to check for specific mean differences across time periods.

Overall, we observed that students with normal stress incrementally spent more time on campus, especially towards the second half of the semester. This is an expected trend because students typically spend more time working on projects and preparing for examinations. Yet, students with severe stress were found to spend significantly less time on campus ($p < 0.05$), specifically on T_{24} and T_{26} (see Figure 5(a)).

We observed significantly less involvement in work activities for students with severe stress ($p < 0.01$) than students with normal stress (see Figure 5(b)). While students with severe stress began with more work participation, this pattern did not persist. In addition, two large dips occurred around the recess week (T_{12} - T_{14}) and the end of the semester (T_{23} - T_{25}). These time points closely corresponded to important project milestones for the Software Engineering (SE) course and were significantly different between both groups.

From Figure 5(c), we found that students with severe stress spent significantly more time with small groups ($p < 0.05$), with a significant difference at T_7 . However, a noticeable dip happened during the recess week, with a significant difference at T_{17} . Additionally, we found that students with severe stress spent significantly more time with their groups ($p < 0.05$).

Overall, these results suggest that features generated by coarse-grained location data can be strong indicators of stress levels.

The above training data allowed the machine learning engine 106 to have three different model configurations for severe stress and depression: (1) $Model_{S_{SE}}$, a highly specific stress model which uses domain-specific features, (2) $Model_S$, is a generalised stress model which excludes domain-specific features and (3) $Model_{D++}$, is a depression model which uses generalised features while accounting for personality profile. Figure 6 lists details of how these three different models were set up and performed. Selection of various model settings were determined through a standard procedure of Group K-fold (5) cross validation. SMOTE was applied to correct the highly imbalanced training dataset. Correspondingly, samples from $Study_{ValidA}$ and $Study_{ValidB}$ were used to externally validate the final model. Performance metrics were defined to cater for class imbalance in our test/validation sets, mainly using *true positive rate* (TPR), *true negative rate* (TNR), *overall accuracy* (ACC), and *area under the ROC curve* (AUC). AUC score close to 1 indicates a perfect predictor.

Figures 7(a) to 7(d) summarise results of our cross validation (CV) of the trained machine learning engine 106 at different stages.

We first checked if all features (raw + rel + abs) make stronger predictors of stress. Using three different algorithms, we achieved the best results using Random Forest (RF) (see Figure 7(a)).

Continuing the use of different sets of features as input, we sought to investigate which feature set works best. As shown in Figure 7(b), while we achieved best performance using all sets, the addition of raw features did not lead to significantly better performance. Hence, we retained only the change sets to avoid developing an overfitted or computationally expensive model. As part of the

feature selection process, we used recursive feature elimination (RFE), with a backward elimination with a step size 1, on all change features. However, no change features were completely redundant, and the performance of the RF classifier peaked with all change features considered. Figure 7(c) lists the ten most important features (and their sets), in order of variable importance.

Finally, we determined the time window, in increments of 3 days (corresponding to the frequency of PSS-4 samples), that provides the best stress detection. From Figure 7(d), we observed that TPR remained high (>90%), for both cross validations, using 3- or 6-day intervals. We selected 6-day intervals, compared to 3, as the misclassification rate was significantly lower for 5-fold CV. While the overall accuracy increases over time, it is important to consider time as key factor of intervention. That is, while stress is regarded as an everyday experience and chronic stress evolves over a longer period of time, prolonging detection of severe stress by more than a week might result in students missing out on vital help, leading to depression. Since timeliness should be prioritised, we concluded the best setting to be 6-day interval. Overall, we used a Random Forest with change type features calculated at 6-day interval as our main stress model, *Model_{S,SE}*.

Figure 8 summarises our test and validation results of the variant *Model_S*. As mentioned above, the only difference between them is the exclusion of domain-specific features, which are highly tailored to SE-students in the *Study_{SE}* sample. Generalising the model to exclude domain-specific features also allows for validation on other samples.

Our solution yielded AUC score of 0.91 and 0.94 from external validation. Specifically, while TPR dropped to 66.67% for *Study_{ValidA}*, 1 out of 3 severe stress cases was misclassified. The same test on *Study_{ValidB}* successfully yielded 100% TPR as it correctly detected 1 severe stress instance. The overall accuracy remained above 80%.

One way to demonstrate scalability is to build an effective stress model using all student populations. That is, the model can be used effectively on students, regardless of the training samples. To test this, we built an all-population model, using data from *Study_{SE}*, *Study_{ValidA}* and *Study_{ValidB}*. As shown in Figure 8, our cross validation results achieved 96.01% TPR, where 4 out of 149 severe stress instances would go unnoticed.

Testing *Model_S* on all populations revealed at least 19 (out of 108) students who did not report any accounts of severe stress being misclassified multiple times during the study. Upon investigation, many of these students had reported feeling depressed from their PHQ-8 assessment despite not experiencing severe stress (PSS-4 scores were not high enough, over 12). This large amount of depression without a corresponding severe stress indication reaffirmed our decision to extend the problem assessment to also detect depression. We now show how we modified our stress model to also detect depression – *Model_{D+}*.

We first tried to detect depression using *Model_S*. However, this had a low TPR of 59%. Digging deeper, we found that students only reported depression (PHQ-8 score above 10) after 39

days on average, with the earliest being after 12 days. This matches clinical wisdom as, typically, people develop quicker responses to stressful experiences, whereas depressive episodes often intensify over longer time periods. Thus, we hypothesised that a longer time interval ($Model_S$ uses 6-day intervals) was needed to accurately detect depression. Subsequently, we repeated the time window experiment to determine an appropriate interval for detecting depression. We maintained a 3-day window size to match the original collection of PSS-4 labels, and used windows sizes ranging from 6 up to 18 days in length – as shown in Figure 9. We found that the Random Forest classification algorithm using change features and a 15-day interval yielded highest TPR on cross validation. We call this $Model_D$.

We validated $Model_D$ on two test sets, $Study_{ValidA}$ and $Study_{ValidB}$, and achieved AUC score of 0.69-0.70. We expanded our analysis to review the profiles of our students such as their gender, academic year, GPA, and Big-5 personality test filled out as part of the demographics survey. Our training set has an approximately fair distribution of freshmen, sophomores, and juniors and we noticed that a high percentage of students (90% of sample) had high score on Neuroticism (only 1 case of depression by a student who scored low on neuroticism, score ≤ 2.25 out of 5) on their Big-5.

We thus decided to add Big-5 factors as additional features in our depression model. As shown in Figure 10(a), the addition of various personality traits as features improved overall accuracy by approximately 10% (difference is significant at $p=0.05$ level). A brief examination on each dimension revealed Neuroticism as the strongest predictor (achieving 100% variable importance, while the next best- ranking trait, Openness, yielded a 10% importance). Using only Neuroticism resulted in a higher TPR of 90.21%, however, this difference is not significant. Considering the practicality of using Big-5 assessment, which is often deemed too lengthy (assessment comprises 44 questions), using only a subset of this information could increase model accuracy (Neuroticism has 8 questions). Hence, we concluded $Model_{D+}$ to include Big-5 Neuroticism score as an optimisation feature. Figure 10(b) shows test results on depression model, $Model_{D+}$.

The above thus describes a system that can detect severe stress and depression episodes in individuals by using, for example, coarse-grained location data collected passively from WiFi infrastructure, to calculate changes in an individual's routines and group interactions. A trained classification algorithm was discussed which uses changes in mobility features to detect stress at 6-day intervals, and depression at 15-day intervals. This classification algorithm, trained with three rounds of semester long user studies at an Asian university campus, achieved an ideal 96.01% TPR (81.76% overall accuracy) in detecting stress on all samples using a generic stress model and up to 91.21% TPR (75.06% overall accuracy) in detecting depression on all three student populations using a more generalised depression model with personality profile as additional features. Accordingly, the use of training data which excludes one or more activities that are unique to a domain to which a

monitored subject belongs would result in a classification algorithm that has more generic applications.

Figure 11 is a schematic of a server 1100 which is a component of the system 100 of Figure 1. The server 1100 hosts the machine learning engine 106 and one or more of the other three blocks of the tracking sub-system 102, the activity mapper 103 and the feature extractor 104. For the sake of simplicity, only selected electronic components of the server 1100 are shown in Figure 11.

The server 1100 has a microprocessor 1102 designed with circuitry that is programmable to perform instructions that result from executing computer code that is contained in memory 1108, with which the microprocessor 1102 is in operative association. These instructions, when executed, cause the microprocessor 1102 to receive, via a communication port 1124, location data providing information indicative of venues visited by a monitored subject, the location data being provided by an access point situated at the venues. The location data is processed against a library of activities stored in a non-volatile storage device 1116, to deduce one or more activities done at each of the visited venues, the deduction being based on an event log recorded at the visited venue. The microprocessor 1102 constructs a current movement pattern to capture a sequence of an activity from the deduced one or more activities being done and calculates change values of one or more features by comparing the current movement pattern against a past movement pattern of a sequence for the same activity being done. The microprocessor 1102 analyses the change values using a classification algorithm that models a relationship between data from one or more movement patterns and their indicative stress outcome or depression outcome; and provides, as output, a stress outcome or depression outcome or both in the monitored subject. The output of stress or depression may be sent to display 1104 or to the communication port 1124 for transmission to external devices (not shown) for further action. The communication port 1124 thus acts as an input/output interface for the server 1100.

In providing the output of the stress outcome or depression outcome, the microprocessor 1102 sends commands, over communication infrastructure 1106, with all components (such as the memory 1108, the storage device 1116, the communication port 1124 and the display 1104). The communication infrastructure 1106 refers to data communication channels such as a bus.

In the application, unless specified otherwise, the terms "comprising", "comprise", and grammatical variants thereof, intended to represent "open" or "inclusive" language such that they include recited elements but also permit inclusion of additional, non-explicitly recited elements.

While this invention has been described with reference to exemplary embodiments, it will be understood by those skilled in the art that various changes can be made and equivalents may be substituted for elements thereof, without departing from the spirit and scope of the invention. In addition, modification may be made to adapt the teachings of the invention to particular situations and materials, without departing from the essential scope of the invention. Thus, the invention is not

limited to the particular examples that are disclosed in this specification, but encompasses all embodiments falling within the scope of the appended claims.

CLAIMS

1. A computer-implemented method of detecting stress outcome or depression outcome or both in a subject, the method comprising:
 - 5 receiving, as input, location data providing information indicative of venues visited by the subject, the location data being provided by an access point situated at the venues;
processing the location data against a library of activities, stored in a database, to deduce one or more activities done at each of the visited venues, the deduction being based on an event log recorded at the visited venue;
 - 10 constructing a current movement pattern to capture a sequence of an activity from the deduced one or more activities being done;
calculating change values of one or more features by comparing the current movement pattern against a past movement pattern of a sequence for the same activity being done;
analysing the change values using a classification algorithm that models a relationship
15 between data from one or more movement patterns and their indicative stress outcome or depression outcome; and
providing, as output, a stress outcome or depression outcome or both in the subject.
2. The method of claim 1, wherein the stress outcome or the depression outcome has a binary
20 value.
3. The method of claim 1 or 2, wherein the past movement pattern is constructed from activities done over an earlier period of time by a population pool.
- 25 4. The method of claim 3, wherein the population pool contains data from peers to which the subject belongs.
5. The method of any one of the preceding claims, wherein the past movement pattern is constructed from activities done over an earlier period of time by the subject.
30
6. The method of any one of the preceding claims, wherein the location data comprises information indicative of venues visited by a group that is detected to have interacted with the subject, wherein the sequence of the deduced one or more activities includes activities the subject is deduced to have done with the group at the visited venues and wherein the calculated change values include
35 one or more features derived from the activities the subject is deduced to have done with the group at the visited venues.

7. The method of any one of the preceding claims, wherein the classification algorithm is configured to perform optimisation by analysing the change values of a selection from the one or more features.
- 5
8. The method of any one of the preceding claims, further comprising undertaking data interpolation to obtain missing data required to complete the sequence of one or more activities when constructing the current movement pattern.
- 10
9. The method of any one of the preceding claims, wherein the classification algorithm is trained using data of change in movement patterns providing information of one or more activities done at the visited venues and their corresponding indicative stress outcome or depression outcome.
10. The method of claim 9, wherein the data from the movement patterns include interpolated data used to complete a sequence of activities captured by the movement patterns.
- 15
11. The method of claim 9 or 10, wherein the training data excludes one or more activities that is unique to a domain to which the subject belongs.
- 20
12. The method of any one of the preceding claims, wherein the deduction of the one or more activities done at each of the visited venues is based on duration of the activity and time at which the activity was done.
13. The method of any one of the preceding claims, wherein the model used to provide an output of the stress outcome is different from the model used to provide an output of the depression outcome.
- 25
14. A computer-implemented system for detecting stress outcome or depression outcome or both in a subject, the system comprising:
- at least one microprocessor;
 - 30 at least one non-volatile storage device;
 - at least one computer-readable memory device operatively associated with the microprocessor;
 - and
 - an input/output interface operatively associated with the microprocessor, wherein the computer-readable memory device contains computer-executable instruction code which, when
 - 35 executed via the microprocessor, configures the microprocessor to:

receive, via the input/output interface, location data providing information indicative of venues visited by the subject, the location data being provided by an access point situated at the venues;

5 process the location data against a library of activities, stored in the non-volatile storage device, to deduce one or more activities done at each of the visited venues, the deduction being based on an event log recorded at the visited venue;

 construct a current movement pattern to capture a sequence of an activity from the deduced one or more activities being done;

10 calculate change values of one or more features by comparing the current movement pattern against a past movement pattern of a sequence for the same activity being done;

 analyse the change values using a classification algorithm that models a relationship between data from one or more movement patterns and their indicative stress outcome or depression outcome; and

15 provide, as output, a stress outcome or depression outcome or both in the subject.