EAI Endorsed Transactions
on Pervasive Health and Technology
Research Article **EAI.EU**

Mobile Stress Interventions: Mechanisms and Implications

Luis G Jaimes[1], Robert Steele[1]

[1]Florida Polytechnic University, 4700 Research Way, Lakeland, FL 33805, US

Abstract

According to the American Psychological Association, 49% of the U.S. population suffers from chronic, daily stress. Chronic stress also has significant long-term behavioral and physical health consequences, including an increased risk of cardiovascular disease, cancer, anxiety and depression. In this work, we examine how smartphones and mobile sensing can help address the short and long-term consequences of stress. First, we define a conceptual framework for thinking about the interaction between real-time pervasive devices and the real-time physiology of stress. Second, using this framework, we propose a set of guidelines or requirements for pervasive just-in-time intervention (JITI) systems. Third, based on these guidelines, we specify a three-layer software/hardware architecture to support just-in-time interventions for stress. Several themes emerge from this discussion, including the need for robust and accurate context-sensitive forecasting of future stress. Fourthly we describe our experiments and results demonstrating the feasibility of forecasting future stress from current measurements and the effectiveness of the intervention management approach. Finally we discuss the broader implications of mobile-based stress interventions. Whilst this work focuses on chronic stress, we believe the ideas presented are generalizable to other types of just-in-time pervasive interventions.

Received on 18 September 2016; accepted on 12 July 2017; published on 28 February 2018

Keywords: Time Series Forecasting,Mobile Stress Interventions, Ecological Momentary Intervention, Just-in-Time Intervention, Ubiquitous Sensing, Mobile Health

doi:10.4108/eai.28-2-2018.154343

1. Introduction

Advancements in mobile computing and sensing are rapidly changing healthcare. Under the status quo, the average healthy individual visits the doctor rarely, perhaps just once a year. The doctor assesses the patient and then may prescribe medications, recommend changes in behavior (reduce fat consumption, exercise more, etc.), and other forms of intervention. One year later, the patient returns and a similar process is repeated.

In the emerging new model of health care, the patient is augmented with a smartphone and other sensors that monitor personal health in real-time, as the patient goes about his/her normal daily life The smartphone and potentially cloud-based services can assess the monitored data at a higher frequency (on the order of minutes or seconds, where applicable). In turn, health interventions can also be prescribed frequently. More importantly, interventions are tailored to the known health state of the patient at a specific moment in time and are delivered in the context of real-life, precisely when and where needed. This vision of patient-centered intervention in the natural environment is sometimes called Ecological Momentary Intervention (EMI) In this paper, we investigate the characteristics of mobile device-based interventions and the implications for the technical design of such systems.

We examine these system design challenges in the context of delivering EMIs for chronic stress. Stress is a "silent killer", in that the negative impacts of stress on the body are not instantaneously noticeable. Rather, the effects of stress accumulate over time and lead to impacts on the cardiovascular system amongst other bodily systems. A well-designed EMI system could reduce this accumulation of negative effects by helping individuals reduce stress levels on a daily basis.

*Luis G Jaimes. Email: ljaimes@floridapoly.edu
†Robert Steele. Email: rsteele@floridapoly.edu

In this work, we go one step further, we propose a sensor-based system for preventative/just-in-time intervention for psychological stress and describe the implications of a stress intervention system. The proposed system is based on future stress prediction that permits an intervention before the impacts of stress occur. This type of intervention may help to avoid a stressful episode or reduce the impact of one that is already on its way.

The architecture is composed of three main modules: a continuous sensing and real time stress recognition module, a forecasting module or generator of stress future predictions, and a intervention management module. This work is presented in terms of this three layer architecture system, this metaphor allows us to go into detail in each layer, the connection between the layers and the flow of information between them.

The remainder of the paper is organized as follows: Section II surveys the state of the art in mobile EMI systems. Section III justifies and explains the reasons why future prediction may increase the power and effectiveness of EMI. Section IV describes the different components of the proposed architecture, with emphasis on the prediction and intervention layers. Section V presents the requirements for a system for preventive/just-in-time intervention. Section VI describes a set of three experiments to predict the future occurrence of stress one, two, and three minutes ahead and also experiments to evaluate the intervention management effectiveness. Section VI also presents the results of these experiments. Finally, Section VII describes the broader implications of such EMI-based mobile stress intervention systems.

2. Related Work

In this section, we describe three areas of related work: Stress detection in naturalistic environments, heath intervention management, and forecasting and prediction techniques for physiological data.

2.1. Stress Detection in Naturalistic Environments:

There has been substantial work on stress detection in supervised environments where stress variables can be continuously controlled [14, 30]. Under the hypothesis that psychological stress can be explained through the observation of physiological signals, several studies have focused on the observation of the arousal of the user physiology. In these studies patients are submitted to well known stressors and their physiological signals are measured and analysed by experts.

However, the real challenge consists in the detection of stress in natural environments. Here, unlike the supervised approach, there is not control of the user environment. Readings and recordings of the user physiology are carried out by unobtrusive sensors that do not require the user's intervention.

Some of the most commonly used signals for stress inference in natural settings include: Heart Rate (HR) [23], Respiratory Inductance Plethysmography (RIP) [7], Galvanic Skin Response (GSR) [9], and speech [6, 16].

Examples of systems for stress detection in natural environments include: AutoSense [16] a mobile platform for stress detection that carries out the sensing, processing, and stress inference in real time. The system uses an unobtrusive suite of sensors which via Bluetooth transmit sensor data to a smart phone where the whole process takes place.

Another representative example corresponds to StressSense [16]. This unobtrusive system uses the smart phone microphone to capture speech and then by the use of speech processing and machine learning techniques infer psychological stress.

2.2. Ecological Momentary Interventions

The emergence of ubiquitous technologies such as smart phones and wearable sensors have led to the rise in healthcare methodologies such as Ecological Momentary Intervention(EMI) [8]. The goal of EMI is to complement and reinforce the treatments that take place in the hospital by the use of interventions in natural environments. To accomplish this goal usually EMI is used in conjunction with Ecological Momentary Assessment (EMA).

EMA is used to gather user's feedback in real time (i.e.,written descriptions about feelings, cravings, and emotions) basically non-sensor data, and EMI is used to provide real time support given the feedback acquired through EMA.

EMI-EMA have been widely used as a treatment mechanism for a variety of conditions such as: smoking cessation, weight loss [10], anxiety reduction [19] as well as eating disorder reduction [24].

On the other hand, the creation of intervention strategies as well as their effectiveness evaluation constitutes what is called Dynamic Treatment Regimes and adaptive interventions. Here, factors such as: the decision about whether or not to intervene, the timing, and the intensity that maximize the intervention's effectiveness are the main components of this approach. The research question behind this approach is: given a set treatments or available interventions, when to apply or deliver an intervention, in what order, and which level of intensity to apply in order to maximize treatment efficacy. An appealing characteristic of this approach is its adaptive component which takes into account the differences among patients, which means that treatment (i.e., sequence of interventions in a period of time) is tailored according to characteristics of each patient. Research projects in this field include the work of Murphy and Chakraborty [17] in reinforcement learning as well as Rivera et al. [29] who model the problem using dynamic systems and control theory.

2.3. Prediction and Forecasting of Physiological Signals

An important component of this area has to do with the forecasting of physiological signals to predict stress in advance (i.e., some steps ahead). The goal of the prediction in the early stages is to prevent the user from experiencing a stressful event by the use of an appropriate intervention. This intervention might change the course of a user's actions which would prevent a stressful episode. In this work we use Heart Rate Variability (HRV) as a quantitative marker for stress prediction. Several researchers such as AJ Camm *et al.* [2] point out an increase in the HRV observed during mental stress and moderate exercise in healthy subjects. Thereby, we assume that by predicting HRV k steps-ahead, we can obtain a good approximation of a user's stress k steps-ahead.

A multi-step ahead time series forecast consists of predicting the next H values $[y_{N+1}, ..., y_{N+H}]$ of a time series $[y_1, ..., y_N]$, where $H \geq 1$ denotes the forecasting horizon. In terms of the number of terms predicted in every iteration, we can identify two main approaches: Single-Output, and Multi-Input Multi-Output strategies (MIMO). In the first case, just one term is predicted at a time, while in the second case a sequence of terms are predicted in one step.

Single-Output Strategies. One of the most common strategies corresponds to the *Recursive approach*. Here a single model f is trained to perform a one-step ahead forecast. In other words, $y_{t+1} = \{y_t, \dots, y_{t-d+1}\} + w$. Thereby, to forecast H steps ahead, we forecast the first step by applying the model. Subsequently, we use the value just forecasted in the previous step as part of the input variables for forecasting the next step. We continue in this way until we have forecast the entire horizon. Typically a Hidden Markov Model (HMM) forecasting approach follows this model, there, a time series segment of length d is used to train and create a probability transition matrix which is used to forecast the $n+1$ terms of the time series.

Unlike the recursive approach, with the *Direct strategy* forecast each horizon term independently from the other. In other words, H models f_h are learned from the time series $[y_1, \dots y_N]$. Here, $y_{t+h} = f_h(y_t, \dots, y_{t-d+1}) + w$, where $t \in \{d, \dots, N-H\}$ and $h \in \{1, \dots, H\}$. Because this strategy does not use approximate values or previous forecasting values to forecast, the predicted values are free of error accumulation. A clear disadvantage of this method is the high computational cost. Machine learning models such as neural networks, nearest neighbors, and decision tress have been widely used to implement this forecasting strategy.

Finally, a combination of two previous approaches is called the *DirRec strategy*. The *DirRec* strategy computes the forecast with a different model for every horizon (like the Direct strategy) and, at each time step, it enlarges the set of inputs by adding variables corresponding to the forecasts of the previous step (like the Recursive strategy). Usually this strategy outperforms the previous two.

MIMO Strategies. This strategy learns one Multiple-Output model F from the time series $[y_1, \dots, y_N]$ where $[y_{t+H}, \dots, y_{N-d+1}] = F[y_t, \dots, y_{t-d+1}] + w$, with $t \in \{d, \dots, N-H\}$, $F : R^d \rightarrow R^H$ is a vector-valued function. Here, the output of the forecasting in each step corresponds to a set of several terms. Typical models that follow this approach include a Gaussian regression multi-output process.

In real word applications often these methods are utilized using a sliding windows approach.

3. Problem Definition

Let P be a set of patients who experience regular episodes of psychological stress, and I a pervasive system of stress management and intervention which continuously monitors P. Thereby, the problem for pervasive stress management and intervention can be stated as follows: How to keep P free of stressful episodes for as long as possible; and at the same time create the psychological resilience that successfully helps patients to cope with stress and eventually reduce the episode's frequency?

To address this research question, we propose a framework for pervasive systems that includes the following three components: preventers, reducers, and reflectors. Preventers aim to prevent the occurrence of significant stress. Reducers aim to reduce a user's current level of stress. Lastly, reflectors aim to facilitate reflection on past stress, with the goal of preventing future stress.

The first two component of the proposed system are meant to keep the patients for as long as possible free of the experience of stress and reduce the negative consequences of it. Finally, the reflection component aims to help the user to reflect upon the causes and behavior that induce stress as well as to build the phycological resilience to face or avoid future stress episodes.

In this work we represent a stress episode by an increasing signal or wave that reaches a peak or point of inflexion at the time of maximum stress Figure 1 and then starts to decrease until reaching its previous state. This representation allows us to define stress interventions in relation to time, intervention effectiveness, and stress reduction as follows:

Definition 1. Once a stress event has been detected if the stress intervention is delivered before the user reaches the peak of his/her maximum stress levels, then we say that the intervention was delivered on time. See Figure 3

Definition 2. If as a result of a stress intervention the user never reaches the peak of the maximum stress levels that may be caused by a stressful event, or the intervention smoothes his/her stress curve, then we say that the intervention has been effective.

Definition 3. The difference in maximum between the point of inflexion of the potential stress peak (i.e., no intervention) and the new smooth curb (i.e., after the intervention) is called the stress reduction see Figure 3.

Figure 1. Scheme of mobile intervention for stress vs no intervention

3.1. Preventers

The idea here is to predict stress with some degree of probability in advance (i.e., some steps ahead) of the occurrence of a stressful episode. The goal of stress prediction is to prevent the user from experiencing a stressful event by the use of the appropriate intervention. This intervention may change the user's mind and change the course of his/her actions preventing stress. Figure 2 shows this case. Here, at time $t - \delta$ the system predicts stress at time t, and an intervention is delivered at time $t - \delta + \alpha$. As a result, the graph shows that the potential stressful episode represented by the solid curve was avoided and in its place a dash line illustrates that the patient continues in his/her non-stressful state. This ideal situation is possible thanks to the stress prediction occurring some steps ahead.

3.2. Reducers

This is the case, when the Preventer fails in predicting stress. Here, the stress is detected in its early stages, and the goal is to prevent the user reaching the peak of his/her maximum stress. Figure 3 shows an ideal situation. Here, the stress recognition system uses a

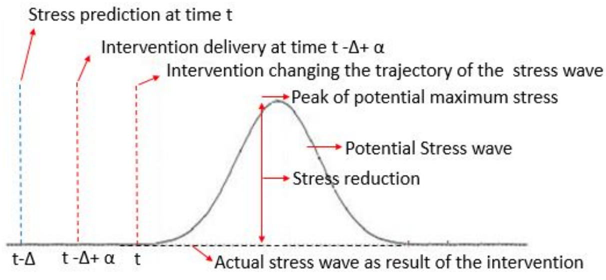

Figure 2. Stress prevention by using stress forecasting

time windows $(t - \alpha, t)$ of length α to compute the probability of finding stress in this time period. At the end of the interval, i.e., at time t the system determines whether or not there was stress in this time interval and delivers an intervention at time $t + \delta$. If the intervention is delivered on time, we can expect an effective result, i.e. a significant reduction of the user's stress levels.

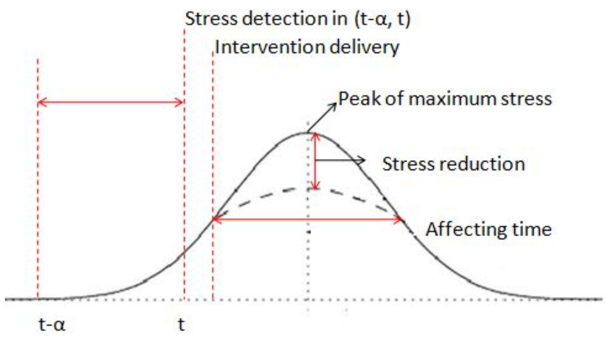

Figure 3. Effect of a mobile intervention on stress events

However, factors such as the duration in time of the stressful event or the fact that the inflexion point (i.e., peak of maximum stress) could be located in the interval of sampling (i.e., time window), may make it difficult to deliver a stress intervention on time. In the former case, Figure 4 illustrates the situation. Here, the stress recognition system monitors in real time the user's stress state. However, given the short duration of the stressful event its wave representation is located completely inside of the time window $(t - \alpha, t)$. Unfortunately, in time interval $(t - 2\alpha, t - \alpha)$ covered by the previous time windows there was not stress, thus, nothing was detected and just up to the time t stress and intervention delivered at time $t + \delta$. At this time the stressful episode has completely concluded and, thus the intervention results are not of use.

In the latter case, the stress event may last enough time to be detected. However, as Figure 5 shows the system splits the signal in chunks of length α, extracts the features of this interval and then using a machine learning algorithm makes an an inference with a time delay of α. The example in Figure 5 shows that even

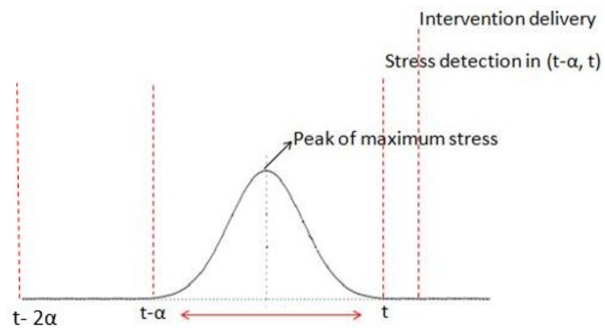

Figure 4. Unable to deliver a stress intervention on time due the short duration of the stress event

though the stressful event starts almost at the middle of $(t-2\alpha, t-\alpha)$ the system does not infer stress at that point because on average most of the time there was- not stress in this interval. However, in the second interval $(t-\alpha, t)$ the system infers stress, unfortunately at this point the user has already reached the peak of maximum stress.

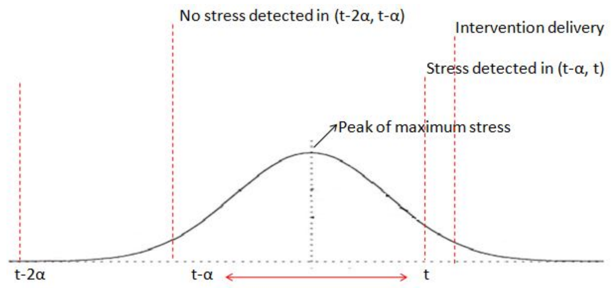

Figure 5. Stress nflexion point inside of the sampling interval

3.3. Reflectors

Figure 5 illustrates a stress intervention episode after the user has reached the point of maximum stress. Here the purpose of the intervention is not to avoid the stressful event, but to provoke a reflection process in order to minimize the stress frequency in the long term. Additionally, this late intervention is aimed at accelerating the user's recovery process or transition to his/her normal or non-stressful state.

4. Proposed architecture

4.1. System Hardware Architecture and Communication Protocols

A proposed hardware architecture for a pervasive system for stress management and just in time intervention consists of the following four main components.

Wearable Sensors: The task of these small and unobtrusive devices is to continuously measure the user's physiological signals and any other user contextual information. They can be implemented via chest bands, waist bands, and embedded in smart devices. The collected data is usually transmitted via Wireless Personal Area Network (WPAN) to a mobile device for aggregation and processing.

Mobile device: The role of this smart device is to receive and consolidate data from a variety of wearable sensors. Additionally, the smart device may add its own sensor data such as acceleration, video, or speech. The whole task of processing, inference (i.e. stress recognition), and user feedback (i.e. intervention management and delivery) can take place in the smart device (i.e. local model). Alternatively, the mobile device can be used as a bridge where data from the different sensors is consolidated and then transmitted to the cloud (i.e. remote model).

Communication: Depending of the system hardware architecture model, the system may use two different communication models. In the case of the local model, the system can use a WPAN usually based on Bluetooth communication protocols. On the other hand, if the system uses a remote model, a combination of WPAN and cellular/WiFi networks with connectivity to the Internet via TCP/IP protocols, could be the appropriate model.

Cloud: This optional component plays a role just if we use a remote architectural model. Here, the storage, processing, inference and intervention management take place in the cloud. The cloud transmits the intervention data to the user's mobile device which in turn displays the intervention to the user via their smart device (e.g., smart phone, smart glasses).

4.2. System Model Architecture

This section describes the three main components of the system architecture with emphasis on the second and third layers. The first layer corresponds to a system of real time stress recognition, the second corresponds to a module for stress forecasting, and finally, the third corresponds to a stress intervention management module.

Sensing and Stress Recognition Layer. This first layer includes hardware and software components. We use data from AutoSense a wearable sensor suite for stress recognition in real time developed for a consortium of universities led by the WiSeMANet lab of the University of Memphis. This platform has been widely documented in [5, 21, 22] and tested in several studies all of them in naturalistic environments. The AutoSense suite of sensors includes: two lead ECG, a respiratory inductive plethysmography band for measurement of respiration rate, skin conductance response (SCR) between two electrodes placed under the chestband, skin temperature with surface probe thermistor, and three-axis accelerometer for motion sensing. Physiological data collected by the sensors is sent via Bluetooth to a smart phone where a stress model based upon machine learning makes stress inferences every minute. The output of this layer corresponds to a set of time series which include low, medium ,and high level measurements. Low level time series corresponds to raw sensor measurement of physiological data. At the medium level we find the features computed during a time window's period. Finally, using as input the medium level data the stress models make inferences about the user's stress state (i.e., high level time series). All the details of this platform can be found in the references provided above.

Stress Forecasting Layer. The input of this layer corresponds to the output of the sensing and stress recognition layer. In order to deal with the high volatility of the input signals, it is transformed into the log domain. On the other hand, to take advantage of the prediction power of HMM, we discretized the HRV input signal using the Symbolic Aggregate approximation algorithm (SAX) [15]. SAX allows a time series of arbitrary length n to be reduced to a string of arbitrary length $w \leq n$ and uses an intermediate representation between the raw time series and the symbolic strings. In this work the SAX algorithm was applied on the log(HRV), transforming a continuous time series of length 4506 into a 4506-symbol string, with alphabet size 20. The whole procedure can be summarized as follows:

$$HRV \longrightarrow Log(HRV) \longrightarrow SAX(Log(HRV)).$$

The prediction process was carried out using a sliding window approach. This dynamic approach uses one thousand elements of the HRV time series to continuously create a prediction model, and using this model predict one, two, and three steps-ahead as is shown in Figure 6. The forecasting algorithm is based on a Poisson-HMM implementation and uses a likelihood prediction approach. Here, every time that the Sensing and Stress Recognition Layer provides a new and fresh value (i.e., HRV computation or stress inference), this value is introduced as a new element of the training set, and used to re-compute the whole prediction model (i.e. HMM transition probability matrix, emission matrix). Once created the model as per Algorithm 1 computes the likelihood for the potential candidates to be the next sequence element and picks the one with the highest likelihood as the prediction.

Intervention Layer. The goal of this layer is to deliver a set of stress relieving interventions. The main idea is to deliver the right intervention or the right combination of them, at the right time to maximize the effectiveness (i.e. reduce the frequency of the interventions) while also reducing the number of interventions. This way, the user can enjoy the benefits of stress reduction or prevention with the least amount of disturbance. In this section we explore a probabilistic mechanism based on reinforcement learning to compute an optimal policy (i.e. a set of rules which decide on the intervention to perform).

The layer's inputs correspond to output of either the sensing and stress recognition (red arrow) or the stress forecasting (black arrow) layers. In the former case, the type of intervention provided to the user corresponds to the Reducer and Reflector (red lines) treatments, as explained in section 3. In the latter case, if the stress prediction is accurate the intervention corresponds to a

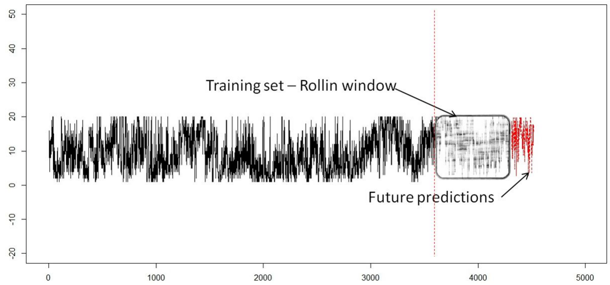

Figure 6. Rolling window as training set and prediction

Algorithm 1: HMM Forecaster

```
input  : O, m, λ, γ, δ, Orange, H
output: S, Matrix of Orange × m of probabilities
        predictions
begin
    n ← length(O)
    prob← outer(O,λ, dpois)
    foo ← δ× prob[1,]
    sumfoo ← sum(foo)
    lscale ← sum(foo)
    foo ← foo / sumfoo
    for i ← 2 to n do
        foo ← foo ×γ× prob[i,]
        sumfoo ← sum(foo)
        lscale ← lscale + log(summfoo)
        foo ← foo / sumfoo
    end
    temp ← matrix(NA, nrow =m, ncol=H)
    for i ← 1 to H do
        foo ← foo ×γ
        temp[,i] ← foo
    end
    prob ← outer(Orange, λ, dpois)
    S ← prob × tempMat[ , 1:H]
    return S
end
```

Preventer treatment (black lines), preventing the user from experiencing a stressful episode.

Reinforcement Learning (RL) [28] is identified to be an appealing modeling framework to optimize the effectiveness of the delivery of stress interventions, because it is designed to solve multi-stage delayed-result decision problems. In particular, we use an episodic Q-Learning(λ) (QL(λ)) algorithm.

In order to use the QL framework, we restate the problem using its standard concepts and notation.

We consider an episode of learning as the interval between the moment stress is detected (or forecasted) and the moment the patient shows no more sign of stress. Additionally, only one intervention of each type can be delivered in each episode. Thus, an episode is also considered terminated if no other intervaention is left to be delivered.

Let $\bar{A} = (a_1 \dots a_M)$ be the possible treatments that can be delivered to the patient during an episode of stress. Let $A_t = a_i$ the set of treatments given to the patient so far in the episode. Since an episode ending implies that the patient has been relieved from stress, those interventions are considered to have failed in curing the patient. Then, we define the QL state at time t as $S_t = A$. We interpret A as a set instead of an ordered list. The underlying assumption is that the order in which treatments are delivered is not important. This assumption greatly reduces the state space. Notice, however, that the number of possible states still grows exponentially with the number of interventions.

A positive reward r_{relief} is given whenever the system succeeds in treating the patient and a negative reward $r_{intervention}$ is given after each intervention. The problem is then to find the policy $\pi : S \times A \rightarrow A$ that maximizes the expected reward, which is equivalent to relieving the patient's stress while minimizing the number of interventions.

Eligibility traces were implemented to improve the algorithm learning rate and no discount was associated to the traces. Thus, the implemented algorithm is a QL($\lambda = 1$) algorithm.

Exploration and exploitation were balanced through the use of an ϵ-greedy strategy, where ϵ depended on the relation between the value of the best known choice and the sum of all values for that state.

Algorithm 2 summarizes one iteration of the simultaneous desicion-making and learning process of the QL(1) algorithm.

Just in Time Interventions (JITI) is a particular instance of the intervention approach described above. The only difference is that the algorithm should learn with new data (i.e., online-learning). In other words, it should incorporate new knowledge on the fly into its learning model with each new state-action interaction. One of the main problems in using an online learning algorithm is storage of old information (i.e. if the history is taken into account). An intermediate solution is to use a sliding window of a fixed length and dynamically generate new rules that take into account the new knowledge. However, sporadically, each t time an off-line algorithm that takes into account the whole history of the data could be used to adjust the policy generation model. Figure **??** shows the proposed architecture.

Algorithm 2: QL(1) intervention delivery system.

```
input : stress levels detected or forecasted
output: the set of interventions to deliver
begin
    forall stata action do
        value(state, action) = 1
    end
    while stress and there are more interventions to
      perform do
        i = chooseIntervention(treatment, value);
        stress = applyIntervention(treatment,
          intervention);
        r = reward(treatment, intervention, stress);
        maxVal = max(value([treatment +
          intervention]));
        updateETrace(treatment, intervention);
        applyQLLearningRule(treatment, intervention,
          reward, maxVal);
    end
end
```

5. Requirements

We motivate the discussion about the requirements or design issues for a sensor based system for just-in-time intervention for stress, focusing in two main components: a module for stress future predictions and an intervention management module.

5.1. Design Issues for Stress Forecasting Systems

Every t seconds the stress recognition layer provides a new fresh stress measurement. In other words, there is a continuous stream of data, and the prediction algorithm should be able to add this new knowledge on the fly to update the prediction model. One of the disadvantages of HMM is its off-line nature. The algorithm uses a portion of the time series as a training set to create the prediction model (i.e., transition probability matrix) and the remaining data as testing. The problem with this approach is that the new incoming data is not being taken into account to update the model; in other words the model becomes old. The way this can be tackled is to re-build the entire model as new samples show up, noting this pseudo-online way of applying HMM has a high computational cost.

The following questions address the design issues in relation to stress forecasting:

- Can the system capture any new information as it becomes available?
- Does the system represent the overall knowledge about the problem without memorizing a large amount of the representative dataset?
- Is the system able to update any knowledge in real time which is observed in the recent data set that was not previously observed during building the initial system and thereby be able to avoid rebuilding a new system when there is no change in the model?

A new generation of algorithms that partially address these problems corresponds to DENFIS [11] and FIS [18] which are able to update new knowledge in real time and thus adapt the parameters of the model only, while keeping the rest of the model unchanged. The downside of these smart prediction algorithms is their complex implementation as well as their computing performance.

5.2. Design Issues for Just in Time Intervention Systems

The question of whether intervention outcomes are independent of each other or not plays an important role in the design of JITI systems because it affects the difficulty of the problem. If the interventions' outcomes are independent of each other, statistical approaches can be used to find the set of most effective interventions for a given patient. Then, the system would always try to deliver those interventions in decreasing order of effectivenes. However, in the medical setting, there is consensus that previous interventions may influence the effectiveness of new ones. We argue in favor of this statement with a graphical model of the stochastic outcome of different interventions. Figure **??** depicts the model in which the stress state s depends on the combination of N independent interventions i_j and a set of M attributes a_k that define the patient.

The joint probability $p(s, a_1, .., a_M, i_1, .., i_N)$ is then split into different factors expressing the effectiveness of each intervention under the pattient attributes, according to Equation 1. Note that, in this case, intervention outcomes are indeed independent of each other.

$$p(s, a_1, .., a_M, i_1, .., i_N) = \prod_{j=1}^{j=N} \phi(s, a_1, .., a_M, i_j) \qquad (1)$$

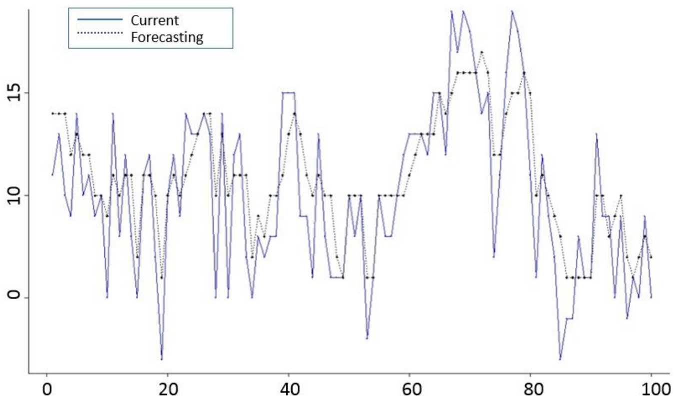

Figure 7. A step-ahead forecasting using a four hidden state Poisson HMM

As the complete set of attributes can never be known, i.e. the factors that affect the effectiveness of each intervention cannot be fully determined, we would like to work with the marginal distribution $p(s, i_1, ...i_N)$. Then, due to marginalization of the attribute variables, the graphical model turns into the one shown in Figure **??**. Under this assumptions, the effectiveness of each intervention becomes dependent of one another.

Consequently, the decision regarding the intervention to apply at a given moment must take into account the outcomes of previously given interventions.

The previous insight also suggests that the solution to this problem could be improved by introducing additional context. Namely, some of the known attributes for that patient. Firstly, in the case of JITIs for stress, context information such as the time of the day and the patient's location could give clues on the cause for the stress, which could be taken into account when selecting the interventions to apply. In addition, extra context such as whether physical activity is being performed by the patient may be important to distinguish a true or false event of stress (i.e., physical activity and stress have similar effects in the physiology). Finally, an a priori questionaire could gather information on the patient's preferences for possible interventions.

6. Experiments and Results

6.1. Stress Forecasting

In these set of experiments, HRV is used as an approximation of psychological stress. A set of three experiments were carried out to measure the prediction accuracy of HMM to predict stress. In all the experiments the training set corresponds to a sliding window of 1000 minutes of stress measurements which runs through the time series. Four, three, and two hidden states were used for experiment one, two, and three respectively.

Table 1. Initial Parameter.

Parameters	Experiment 1	Experiment 2	Experiment 3
Time series size	4506		
Distribution	Poisson		
Parameter estimation	EM Algorithm		
Forecasting range	100	100	100
Forecasting horizon (h)	1, 2, and 3 steps-ahead		
Hidden states (m)	4	3	2
Training set size	1000		
Type of training set	sliding window		
Transition-Prob matrix	$\begin{pmatrix} 0.7 & 0.1 & 0.1 & 0.1 \\ 0.1 & 0.7 & 0.1 & 0.1 \\ 0.1 & 0.1 & 0.7 & 0.1 \\ 0.1 & 0.1 & 0.1 & 0.7 \end{pmatrix}$	$\begin{pmatrix} 0.9 & 0.05 & 0.05 \\ 0.05 & 0.9 & 0.05 \\ 0.05 & 0.05 & 0.9 \end{pmatrix}$	$\begin{pmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \end{pmatrix}$
Delta	(0.25, 0.25, 0.25, 0.25)	(0.333, 0.333, 0.333)	(0.5, 0.5)
Lambda	kmeans(training-set, m)		

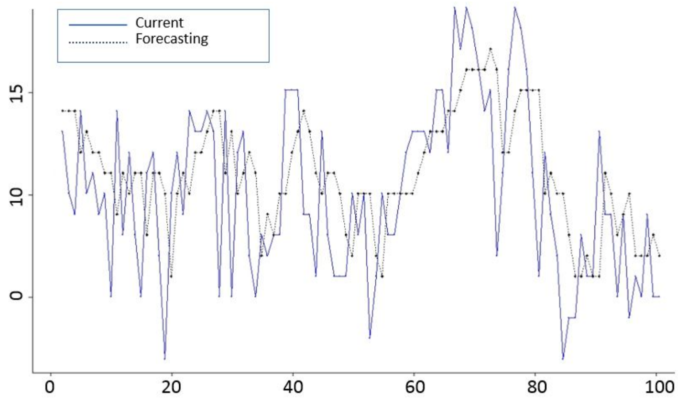

Figure 8. Two step-ahead forecasting using a four hidden state Poisson HMM

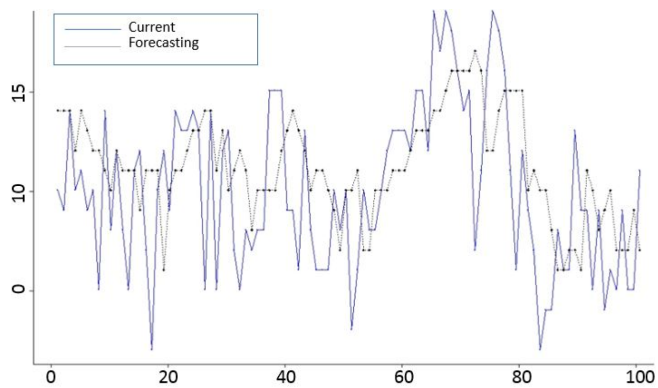

Figure 9. Three step-ahead forecasting using a four hidden state Poisson HMM

Besides, for every experiment one, two, and three steps-ahead predictions were computed as well as prediction accuracy measures in terms of the metrics described in Table 4. For all the experiments, we use a discrete Poisson-HMM implementation using maximum likelihood as a prediction approach. The pseudocode version of the prediction algorithm which corresponds to Algorithm 1 is presented in section 4. Table 1 and Table 2 summarizes the initial experiment's parameters and experiment's results respectively. Figure 7, Figure 8, and Figure 9 show the results of experiment one. Here, a Poisson HMM is used using four hidden states to forecast one, two, and three steps-ahead, which corresponds to the first, second, and third sub-columns of the column Experiment 1 in Table 2. The results of the Experiments 2, and 3 corresponds also to the sub-columns of the columns Experiment 2 and Experiment 3 respectively in Table 2. As can be observed in Table 2 there is not significant difference in terms of the Coefficient of determination (R^2) and Index of Agreement when the HHM predictor was used with four and three hidden states. However, in terms of computational cost, prediction using four states is far more costly than using just three states. Thus, using three states seems to offer a good trade-off between prediction accuracy and computational cost.

6.2. Stress Intervention

In order to performm tests over the intervention platform, we simulated the intervention system using a mathematical model of the patient's reaction to the interventions. Carrying out these tests before including actual subjects in the loop is important due to the online nature of the experiments. Namely, the system must be fully functional and optimized before performing costly experiments on real human beings. In addition, the

Table 2. Experiment results.

METRIC	EXPERIMENT 1			EXPERIMENT 2			EXPERIMENT 3		
ME	-0.72	-0.91	-1.15	-0.83	-1.22	-1.37	0.12	-0.04	-0.09
MAE	1.88	2.99	3.19	1.91	3.12	3.49	2.16	3.14	3.27
RMSE	2.36	3.66	3.94	2.35	4	4.3	2.66	3.95	4.13
NRMSE	83.2	135	150.4	71.7	135	149.9	69.8	104.2	109.3
PBIAS	-6.6	-8.3	-10.3	-7.6	-10.8	-12	1.2	-0.4	-0.9
RSR	0.83	1.35	1.5	0.72	1.35	1.5	0.7	1.04	1.09
rSD	1.43	1.5	1.55	1.24	1.38	1.42	1.06	1.08	1.08
NSE	0.3	-0.84	-1.28	0.48	-0.84	-1.27	0.51	-0.1	-0.21
mNSE	0.16	-0.41	-0.55	0.3	-0.35	-0.6	0.41	0.15	0.1
rNSE	0.15	-1.03	-1.45	0.38	-1.01	-1.51	0.3	-0.37	-0.52
d	0.87	0.67	0.60	0.89	0.64	0.56	0.87	0.72	0.69
md	0.67	0.47	0.43	0.69	0.47	0.4	0.69	0.56	0.53
rd	0.85	0.64	0.57	0.87	0.61	0.51	0.82	0.65	0.61
r	0.84	0.51	0.43	0.84	0.44	0.34	0.77	0.49	0.44
r^2	0.71	0.26	0.18	0.7	0.2	0.12	0.59	0.24	0.2
br^2	0.67	0.24	0.16	0.66	0.17	0.1	0.59	0.23	0.18
KGE	0.54	0.29	0.2	0.7	0.32	0.21	0.76	0.49	0.44
VE	0.83	0.73	0.71	0.83	0.72	0.69	0.78	0.69	0.68

Table 3. QL parameter search data.

Parameter	Min	Step	Max	Best
α	0.6	0.05	0.9	10
γ	0.7	0.05	1	-6
relief reward	10	10	100	.8
intervention cost	0	2	10	.8

developed testing platform provides us with a suitable testbed to try out new ideas that might improve the system perfomance in any way.

We modeled a patient by creating an artificial distribution $p(s|t_1, \ldots, t_N)$ corresponding to the factor graph shown in Figure **??**. The probability of not having stress under a set of treatments was sampled from a beta distribution with the parameters shown in Table **??**. Figure **??** shows the plot for this beta distribution. As one can see, most of the probability mass is concentrated near zero. Thus, most of the treatment combinations will be inefficient, while a few of them will have a greater chance of relieving the patient.

Then, an exhaustive search of the best set of parameters for the QL algorithm was made. Table 3 summarizes the parameters, their search space and the best result. Each set of parameters was evaluated by running the QL system over ten simulations with different random patient models. The mean number of intervention was the chosen metric to minimize. The obtained set of parameters was used for the remaining experiments.

After that, 100 different models were generated using this sampling technique. For each of the models, 1000 episodes were simulated. In each episode, the patient stress was set and the QL system was executed. For each intervention, the probability of stress was taken from the conditional probability of the patient model. Algorithm 3 summarizes this. In order to be able to compare, we implemented a simple system, which

Algorithm 3: Simulated experiments for the intervention system.

begin
  for *experiment in 1:100* **do**
    for *each combination c of interventions* **do**
      model(s|c) = sampleBeta(alpha=.1, beta=1);
    end
    for *episode in 1:1000* **do**
      treatment(:) = 0 **while** *stress and remaining interventions* **do**
        intervention = QL.pickIntervention(treatment) treatment(intervention) = 1 stress = sampleBinom(model(treatment)) QL.learn(treatment, intervention, stress)
      end
    end
  end
end

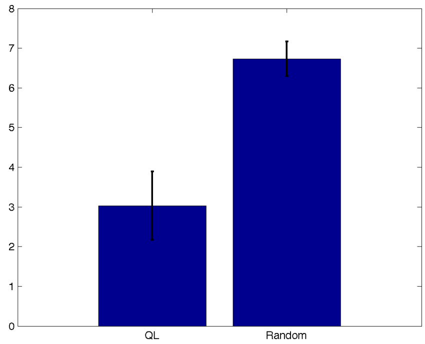

Figure 10. Number of interventions for each system

picked interventions at random. This system served as a way of measuring how difficult it was to relieve stress under the current models. Figure 10 shows the mean number of interventions per episode for each system across all simulations. Figure 11 shows the cumulative average of the number of interventions per episode of one of these experiments.

7. Implications

The above experiment results show that for the stress prediction techniques it is possible to predict stress with up to a 80% probability.

In the case of intervention management the results show that the reinforcement Q-Learning approach taken can optimize the delivery of interventions so as to maximize the effectiveness of the treatment in terms of the number of interventions required.

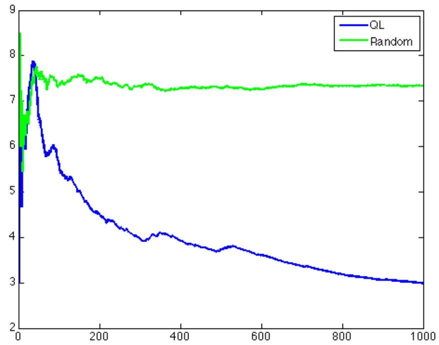

Figure 11. Cumulative average of the number of interventions per episode over one of the experiments.

Table 4. Evaluation Metrics Used in this paper

Metric	Description	Definition
ME	Mean Error	$\frac{1}{N}\sum_{i=1}^{N}(S_i - O_i)$
RMSE	Root Mean Square Error	$\sqrt{\frac{1}{N}\sum_{i=1}^{N}(S_i - O_i)^2}$
NRMSE	Normalized Root Mean Square Error	$\frac{\sqrt{\frac{1}{N}\sum_{i=1}^{N}(S_i - O_i)^2}}{O_{max} - O_{min}}$
PBIA	Percent Bias	$100\frac{\sum_{i=1}^{N}(S_i - O_i)}{\sum_{i=1}^{N}(O_i)}$
RSR	Ratio of RMSE to the standard deviation of the observations	$\frac{RMSE}{STDEV(O)}$
rSD	Ratio of standard deviations between predicted and observed values	$\frac{STDEV(S)}{STDEV(O)}$
rNSE	Relative Nash-Sutcliffe efficiency between predicted and observed values	$1 - \frac{\sum_{i=1}^{N}(\frac{S_i - O_i}{\hat{O}})^2}{\sum_{i=1}^{N}(O_i - \hat{O})}$
r	Pearson product-moment correlation coefficient	$\frac{E\left[(O-\hat{O})(S-\hat{S})\right]}{sd(O)sd(S)}$
r^2	Coefficient of determination	$1 - \frac{SS_{res}}{SS_{tot}}$
KGE	Kling-Gupta Efficiency	$1 - ED$

That is, the introduction of the mechanisms in this paper demonstrates the preliminary feasibility of just-in-time mobile stress interventions and here we consider a number of additional implications of such systems. These include: privacy, clinical relationship, health information system integration, scalability, public health and ethical considerations and in particular the complex inter-relationships between these.

7.1. Privacy

Implicit in such a system is the continuous sensing carried out to detect the onset of stress. As described, this captures HRV plus potentially additional physiological data as indicators of the onset of stress as noted in Section II and Section IV.

Such continuous monitoring of physiology could potentially raise concerns in relation to privacy [1]. It should also be noted that the privacy of a patient's data is protected under the Health Insurance Portability and Accountability Act(HIPAA) when in the possession of a health provider or other covered entity but not when transmitted between individuals or other non-covered entities [20]. HIPAA also does not cover data stored on an individual's mobile device [20] However there are various ameliorating factors, both technical and non-technical that may modify any privacy concerns.

- it is architecturally possible for the processing of the collected data to occur entirely locally on the mobile device and not be sent across the network or to a third-party. As described in Section IV, the local model system architecture allows for processing, stress recognition and intervention management to occur on the device. This decreases both the real and perceived risk of third-party observation of or access to physiologic data

- implicit in the digitized capture and processing of HRV or other data is also the storage of such data. While short-term storage is necessary for the operation of the stress forecasting algorithm, it is longer-term storage of physiological data that will be more pertinent to privacy considerations. Three factors have significant implications for this: (1) control via user preferences; (2) the local device architecture; and (3) the relationship to personal health records

- as per (1) above, allowing user control of whether data is captured, the periods of data capture, where data is transmitted and how data is stored will decrease privacy concerns for the user [25].

- as per (2) above, storage of data on the local device potentially provides greater privacy control and assurance for the end user.

- whether such patient generated data should or must be stored in the person's electronic medical record (EMR) or in particular in their personal electronic health record (PHR) has further privacy implications and would be an important implementational and regulatory aspect of such a system. There are various models of PHR including tethered, where the PHR is integrated with the health care provider's EMR, or standalone, including portable, where the PHR is not conected to other systems[27]. The latter model will more strictly enforce patient privacy but may also not support various clinical and integration functionality.

7.2. Clinical Relationship

Inter-related with each of the other implications is that of clinical relationship: that is, will the use of a mobile stress intervention system occur under the clinical guidance of a healthcare practitioner or can it be used as a lifestyle-aid, not requiring clinical prescription or consultation. It is also conceivable that different versions of mobile stress intervention system might fall into one of these categories or the other.

In a potential case, where the mobile stress intervention system does not require clinician prescribing, then there will be less requirement for the transmission of data beyond the device, potentially less need for integration with clinical and other information systems and more personal responsibility and discretion in how such a system is utilized.

In the case where the system is used under clinician prescription and clinician guidance, there will be a greater likelihood that patient-generated data from the system will be integrated with digital storage such as with the patient's PHR. It could however be the case that the collected data is still not stored in a record, or does not become shared or communicated to the clinician. That is, there are various contemporary digital health devices, that whilst they capture patient-generated health data, these readings are not also added to a PHR [13].

Clinician guidance will enable the usage of such a system in a more supervised way, but of course as an emerging technological area of development, clinical usage guidelines do not as yet exist, and the regulatory and reimbursement environment is still to evolve.

7.3. Health Information System Integration

In considering the hardware architecture we have in the above sections only described the immediate architectural components as they relate to the local model or the remote model (Section IV).

There are three areas of significant interest in relation to integration with broader health information systems:

Clinical Information Systems (CIS): as described there are various circumstances where it is desirable that mobile stress intervention systems may have some level of integration with CIS and with PHRs in particular. Currently, detailed physiological data is only kept longer term in some instances of intensive care unit (ICU) monitoring and for some fitness wearables and sensors, but often not in PHRs. PHRs are being extended to store a greater variety and new types of data. Patient-generated data from such systems as mobile stress monitoring systems may potentially be stored in a patient's PHR. As is the case currently, often patient-generated can be optionally added at the choice of the patient. To address any related privacy concerns, such systems could allow for user-preferences governing such longer-term storage. A positive of such integration, is that more detailed data pertinent to a patient will be kept potentially allowing analysis for early detection of other potential health conditions.

Public Health Information Systems (PHIS): it is possible to integrate public health information systems with mobile stress intervention systems. Future PHIS that allow the aggregation of anonymized mobile device-sensed data from populations or sub-populations, provide a mechanism to signficantly extend current public health information systems and public health capabilities both in terms of population health data capture and also potentially public health intervention capabilities.

Research Information Systems (RIS): whether utilizing aggregated data from CIS, or as separate systems, there is the potential to analyze the data from groups or populations to realize new research results in stress onset, clinical guidelines and effective stress interventions [12].

7.4. Scalability and Affordability

Notably, the introduced mobile stress intervention system, involving a software/ hardware-based system is highly scalable. The proposed system does involve specific sensing hardware such as a chest strap but other sensing approaches described and emerging sensors may make the needed personal sensing capability even cheaper and more ubiquitous. As such, it is relatively inexpensive to deploy such a system to a wide number of people with the cost of initial development also further amortized as a larger number of users utilize the system. In comparison with medication-based interventions for stress, whose cost grows in proportion to the number of users and also in proportion to the duration of patients' treatments, the cost of mobile stress intervention systems does not grow in this exponential way and this makes them a potentially preferable and affordable option.

For either architectural model, local or remote, the costs per additional user will be low and amortized. This also implies the capacity to improve the stress levels of a higher numbers of people via such a system.

An implication of this is the feasibility of application to each of the broader health information systems identified above, namely CIS, PHIS and RIS. It should also be noted that such pervasive health systems can offer comparative advantages for care in rural and regional areas. That is, they can potentially be an improvement over the cost scalability of medications and other care options, their application will not pose bandwidth constraints even in under-developed regions particularly in the case of the local model architecture, and they also help to overcome a challenge often faced in rural and regional areas of a shortage of trained clinicians and specialists resident in those areas [26].

The scalability of a mobile stress intervention system also adds greater import to the consideration of the ethical considerations identified later in this section.

7.5. Public Health

New modes and practices of public health and new PHISs are a potential implication of the development of such systems. In the first instance such systems, if implemented as a remote model, will support previously unavailable population health data collection capabilities.Whilst traditional population health data capture measures still include such means as surveys and interviews, a mobile stress intervention system will allow much more detailed data to be captured and aggregated. Once anonymized, this can support public health data goals [3].

In the area of stress management, this can also support new modes of smartphone-based public health intervention [4]. By the interventions being adaptive, there is also the potential for the traditional boundaries between individual clinical care and public health intervention not being as defined, but rather that a continuum might exist between these two aspects of healthcare.

7.6. Ethical Considerations

Beyond privacy considerations, there are various additional ethical considerations implied by mobile stress intervention systems. Such systems provide the possibility to augment human experience to produce a less stressful living experience, which is recognized to have various health benefits. However there are also other impacts to consider.

Autonomy: such systems can act as personal stress coaches, decreasing the stress of an individual, but the system hence implicitly entails affecting the psychological state of individuals potentially on a frequent basis. This can affect such cognitive aspects as emotional state and hence even decision making processes. This poses the interesting question of the impact on volition and autonomy, and the effects of the system on these areas. This is an area warranting further consideration and research.

Emergency situations: in some cases, stress is a response to immediate, even life-threatening situations. In such cases, stress provides a valid and valuable function of heightened awareness and reaction, and hence in those cases it may not be desirable to attempt to decrease stress. There may be a need to differentiate chronic and non-emergency stress situations from situations where acute reaction or action is required, and for which a heightened level of stress is justified or even important.

Dependency: ideally over time, individuals would improve in their ability to manage their own stress without a supporting system. An aspect of the proposed system, Reflectors, addresses this area.

Population effects: as described, the scalability of such systems poses important implications. They could potentially have the effect of decreasing the stress and improving the health of large numbers of individuals. This effect on a large number of individuals also increases the need for careful consideration of the ethical implications of the system.

8. Conclusion

In this paper, we have introduced an architecture for pervasive stress management and just in time intervention, including describing experimentation and evaluation. We focused on the stress prediction and intervention management layers. In the first case we used real data from a system of stress recognition in real time, to show that it is possible to predict stress with up to 80% probability. In the second case, we showed a probabilistic model to optimize the intervention application in order to maximize the effectiveness in terns of number of interventions. The proposed system has the capacity not only to alleviate or reduce the impact of stress on the patients, but also the possibility to predict the stressful episodes and avoid the user experiencing such stressful periods.

We also address the broader implications of such a pervasive stress management system. Whilst this work concentrates on the mechanisms for addressing chronic stress, we believe many of the insights are also applicable to other just-in-time pervasive health interventions.

References

[1] Louis Atallah, Benny Lo, and Guang-Zhong Yang. Can pervasive sensing address current challenges in global healthcare? *Journal of epidemiology and global health*, 2(1):1–13, 2012.

[2] AJ Camm, M Malik, JT Bigger, G Breithardt, S Cerutti, RJ Cohen, P Coumel, EL Fallen, HL Kennedy, RE Kleiger, et al. Heart rate variability: standards of measurement, physiological interpretation and clinical use. task force of the european society of cardiology and the north american society of pacing and electrophysiology. *Circulation*, 93(5):1043–1065, 1996.

[3] Andrew Clarke and Robert Steele. Health participatory sensing networks. *Mobile Information Systems*, 10(3):229–242, 2014.

[4] Andrew Clarke and Robert Steele. Smartphone-based public health information systems: Anonymity, privacy and intervention. *Journal of the Association for Information Science and Technology*, 66(12):2596–2608, 2015.

[5] Emre Ertin, Nathan Stohs, Santosh Kumar, Andrew Raij, Mustafa al'Absi, and Siddharth Shah. Autosense: unobtrusively wearable sensor suite for inferring the onset, causality, and consequences of stress in the field.

In *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems*, pages 274–287. ACM, 2011.
[6] Ling He, Margaret Lech, Namunu Chinthaka Maddage, and Nicholas Allen. Stress detection using speech spectrograms and sigma-pi neuron units. In *Natural Computation, 2009. ICNC'09. Fifth International Conference on*, volume 2, pages 260–264. IEEE, 2009.
[7] Jennifer A Healey and Rosalind W Picard. Detecting stress during real-world driving tasks using physiological sensors. *Intelligent Transportation Systems, IEEE Transactions on*, 6(2):156–166, 2005.
[8] Kristin E Heron and Joshua M Smyth. Ecological momentary interventions: incorporating mobile technology into psychosocial and health behaviour treatments. *British journal of health psychology*, 15(1):1–39, 2010.
[9] Frank Horvath. An experimental comparison of the psychological stress evaluator and the galvanic skin response in detection of deception. *Journal of Applied Psychology*, 63(3):338, 1978.
[10] Nam-Seok Joo and Bom-Taeck Kim. Mobile phone short message service messaging for behaviour modification in a community-based weight control programme in korea. *Journal of Telemedicine and Telecare*, 13(8):416–420, 2007.
[11] Nikola K Kasabov and Qun Song. Denfis: dynamic evolving neural-fuzzy inference system and its application for time-series prediction. *Fuzzy Systems, IEEE Transactions on*, 10(2):144–154, 2002.
[12] Santosh Kumar, Wendy Nilsen, Misha Pavel, and Mani Srivastava. Mobile health: Revolutionizing healthcare through transdisciplinary research. *Computer*, 46(1):28–35, 2013.
[13] Jeon Lee and Joseph Finkelstein. Activity trackers: a critical review. In *25th European Medical Informatics Conference, MIE 2014*. IOS Press, 2014.
[14] Wenhui Liao, Weihong Zhang, Zhiwei Zhu, and Qiang Ji. A decision theoretic model for stress recognition and user assistance. In *AAAI*, volume 2005, pages 529–534, 2005.
[15] Jessica Lin, Eamonn J. Keogh, Stefano Lonardi, and Bill Yuan chi Chiu. A symbolic representation of time series, with implications for streaming algorithms. In *DMKD*, pages 2–11, 2003.
[16] Hong Lu, Denise Frauendorfer, Mashfiqui Rabbi, Marianne Schmid Mast, Gokul T Chittaranjan, Andrew T Campbell, Daniel Gatica-Perez, and Tanzeem Choudhury. Stresssense: Detecting stress in unconstrained acoustic environments using smartphones. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, pages 351–360. ACM, 2012.
[17] Susan A Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):331–355, 2003.
[18] PC Nayak, KP Sudheer, DM Rangan, and KS Ramasastri. Short-term flood forecasting with a neurofuzzy model. *Water Resources Research*, 41(4), 2005.
[19] Michelle G Newman, Andres Consoli, and C Barr Taylor. Computers in assessment and cognitive behavioral treatment of clinical disorders: Anxiety as a case in point. *Behavior Therapy*, 28(2):211–235, 1997.
[20] C Petersen and P DeMuro. Legal and regulatory considerations associated with use of patient-generated health data from social media and mobile health (mhealth) devices. *Applied clinical informatics*, 6(1):16, 2015.
[21] Kurt Plarre, Andrew Raij, Syed Monowar Hossain, Amin Ahsan Ali, Motohiro Nakajima, Mustafa Al'absi, Emre Ertin, Thomas Kamarck, Santosh Kumar, Marcia Scott, et al. Continuous inference of psychological stress from sensory measurements collected in the natural environment. In *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*, pages 97–108. IEEE, 2011.
[22] Kurt Plarre, Andrew Raij, Syed Monowar Hossain, Amin Ahsan Ali, Motohiro Nakajima, Mustafa al'Absi, Emre Ertin, Thomas Kamarck, Santosh Kumar, Marcia Scott, Daniel P. Siewiorek, Asim Smailagic, and Lorentz E. Wittmers. Continuous inference of psychological stress from sensory measurements collected in the natural environment. In *IPSN*, pages 97–108, 2011.
[23] Pramila Rani, Jared Sims, Robert Brackin, and Nilanjan Sarkar. Online stress detection using psychophysiological signals for implicit human-robot cooperation. *Robotica*, 20(06):673–685, 2002.
[24] Sarah Robinson, Sarah Perkins, Stephanie Bauer, Neil Hammond, Janet Treasure, and Ulrike Schmidt. Aftercare intervention through text messaging in the treatment of bulimia nervosa?feasibility pilot. *International Journal of Eating Disorders*, 39(8):633–638, 2006.
[25] Christoph Stach and Bernhard Mitschang. Privacy management for mobile platforms–a review of concepts and approaches. In *Mobile Data Management (MDM), 2013 IEEE 14th International Conference on*, volume 1, pages 305–313. IEEE, 2013.
[26] Robert Steele and Amanda Lo. Telehealth and ubiquitous computing for bandwidth-constrained rural and remote areas. *Personal and ubiquitous computing*, 17(3):533–543, 2013.
[27] Robert Steele, Kyongho Min, and Amanda Lo. Personal health record architectures: technology infrastructure implications and dependencies. *Journal of the Association for Information Science and Technology*, 63(6):1079–1091, 2012.
[28] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*, volume 1. Cambridge Univ Press, 1998.
[29] Kevin P Timms, Daniel E Rivera, Linda M Collins, and Megan E Piper. A dynamical systems approach to understand self-regulation in smoking cessation behavior change. *Nicotine & Tobacco Research*, page ntt149, 2013.
[30] Jing Zhai and Armando Barreto. Stress recognition using non-invasive technology. In *FLAIRS Conference*, pages 395–401, 2006.