A Stress-Free Life: Just-in-Time Interventions for Stress via Real-Time Forecasting and Intervention Adaptation

Luis G Jaimes
University of South Florida
ljaimes@mail.usf.edu

Martin Llofriu
University of South Florida
mllofriualon@mail.usf.edu

Andrew Raij
University of South Florida
raij@usf.edu

ABSTRACT

Chronic stress has significant long-term behavioral and physical health consequences, including an increased risk of cardiovascular disease, cancer, anxiety and depression. This paper conducts post-hoc experiments and simulations to demonstrate feasibility of both real-time stress forecasting and stress intervention adaptation and optimization. Using physiological data collected by ten individuals in the natural environment for one week, we first show that simple Hidden Markov Models can be used to forecast heart rate variability - a proxy for stress - up to 3 minutes in advance. Second, we expand Q-Learning (QL), a reinforcement learning methodology previously used to adapt non-pervasive interventions, to take advantage of the always-available nature of pervasive technology. Using eligibility traces, we demonstrate how QL could be used by a pervasive health system to adapt and deliver any number and type of interventions for a given health event. Our hope is that this work will take us one step closer to using pervasive devices to assist in the daily management of chronic stress and other health challenges.

Keywords: Time Series Forecasting; Just-in-Time Adaptive Intervention; Ubiquitous Sensing; Mobile Health.

1. INTRODUCTION

Advancements in pervasive computing are rapidly changing preventative 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 and recommend behavior changes (reduce fat consumption, exercise more, etc.). One year later, the patient returns and this process is repeated. In the emerging new model of health care, the patient carries sensors that monitor health in real-time, as the patient goes about normal daily life [3]. A smartphone and cloud-based services assess monitored data at a much higher frequency (on the order of minutes or seconds, if needed), allowing health interventions to be prescribed and delivered more frequently and in the natural environment.

BODYNETS 2014, September 29-October 01, London, Great Britain

DOI 10.4108/icst.bodynets.2014.258237

This vision of intervention in the natural environment is sometimes called Ecological Momentary Intervention (EMI) [3]. An important variation on EMI are just-in-time adaptive interventions (JITAIs) [1]. JITAIs leverage real-time and historical information about the user to maximize the success of the intervention. The JIT component of JITAI refers to the idea that the intervention is delivered precisely when needed (neither too early nor too late; just in time). The AI component of JITAI refers to algorithms that aim to maximize intervention success via real-time adaption of interventions to the user and their context. JITAIs search the space of intervention parameters (i.e., intervention type, timing, and dose) to select the intervention that would best address the user's health at the current moment in time. Murphy and Chakraborty examined the use of reinforcement learning [9] for this optimization. Rivera [14] models the problem using dynamic systems and control theory. These optimization approaches have been used as a treatment mechanism for variety of disorders such as: Smoking cessation, weight loss, anxiety reduction, and eating disorder reduction [12].

Building on previous work on the timing constraints imposed by JITAI [4], we identify a need for dedicated intervention management as well as the ability to forecast health state in advance. The latter enables intervention before a sudden increase in stress. We then propose a three layer architecture for JITAIs: a continuous sensing layer, a real time stress recognition and forecasting layer, and an adaptive intervention management module. Several mHealth systems include continuous sensing and health state recognition [2,8]. However, to our knowledge, forecasting and real-time adaptive intervention components have not yet been investigated in the literature. We investigate these two new components of the architecture more closely with post-hoc simulations on real-world data.

We ran three experiments that show it is feasible to forecast physiological variables associated with stress up to 3 minutes in advance. We used a Hidden Markov Model with two, three, and four hidden states. Our experiments show that it is possible to obtain stress predictions with a r^2 of 71% and a coefficient of agreement of 89%, given a large enough training set (i.e., one thousand minutes of time series data).

Another key innovation of this paper is the use of Q-Learning to deliver any number of interventions at any number of time points. In previous work, Nahum-Shani et al [10] introduced the notion of using q-learning for adaptive intervention. They used Q-Learning with linear regression to choose from two interventions at two points in time. This intervention frequency and quantity was chosen because of the nature of medical practice, in which health professionals and patient interact infrequently, and there are little opportunities to deliver interventions and try new interventions. With the expansion of pervasive technology into our daily lives, it is now possible to deliver health interventions at far more frequent intervals and to try a much larger set of interventions over a longer period of time. We adapt Q-Learning to enable selection and delivery of interventions via pervasive, always-on, always-available technologies by introducing eligibility traces. The adaptations we describe here are an important step toward taking advantage of the characteristics of pervasive systems to improve health care.

2. RELATED WORK

2.1 Physiological Stress Detection

Under the hypothesis that psychological stress triggers physiological changes, much previous work on stress measurement has focused on measuring stress by observing the arousal of user physiology. Some of the most commonly used signals for stress inference include: Heart Rate (HR), Respiratory Respiration Rate (RR), Galvanic Skin Response (GSR), and speech [8].

In lab settings, stress is studied by exposing individuals to well known stressors and then measuring the physiological response [6]. Collection of stress in the lab has the advantage of enabling careful control over the experiment, and thus enabling stress researchers to minimize confounding factors in their studies. However, data collection in the lab does not capture the way people interact with stressors in the real-world, and thus is less valid in a the real-world's mobile, pervasive context.

Recent advances in pervasive sensing overcome these data validity challenges by enabling passive capture of stress data in the natural environment. AutoSense [2] is a mobile platform for stress detection that carries out the sensing, processing, and inference of stress in real time. The system uses an unobtrusive suite of sensors to capture heart rate, respiration and other physiological signals, which are then transmitted to a smartphone for storage and processing. StressSense [8] assesses stress unobtrusively through analysis of speech data. A smartphone microphone captures speech signals. Speech processing and machine learning are then used to identify periods of stress.

2.2 Modeling for JITAIs

In this section, we briefly review the most common approaches to model the delivery of heath interventions. Two main modeling approaches have been suggested to address the problem of management and delivery for heath interventions [1], the Markovian approach (Markov Decision Process or Partial Observable Markov Decision Processes) and the reinforcement learning approach, which is usually using a form of q-learning [13]. Both cases, represent two different modeling frameworks whose application depends on the assumptions made about the problem.

The Markovian approach supposes a basic knowledge of the system dynamics, i.e., knowledge about the effect of applying each intervention on each user state. This knowledge is represented in the form of a Transition Probability Matrix (TPM), and the initial parameters of each state. The reinforcement learning case is initially model-free. Through trial and error, the system learns the underlying model.

2.2.1 Markov Decision Process

MDP is a model-based approach that corresponds to an interaction process between an agent and an environment. At every decision point t the agent observes the environment at state s, and based on the observation takes an action a from a set of actions A. As result, the environment makes a transition to a new state $s^{'}$ with transition probability $P(s^{'}|s,a)$ and returns a reward $R(a,s)$ [13]. The main goal of this approach is to obtain a policy that maximizes the sum of the reward.

2.2.2 Q-Learning

Q-Learning is a model-free reinforcement approach that, unlike MDP, assumes no knowledge about the state transition probabilities. Again, this model is made up of four elements: agent, environment states, set of actions, and rewards.

Here, the agent observes the state of the environment and then takes action. At the beginning, the agent chooses a random action in response to a stimulus, because it does not yet have information about which action will yield the greatest reward. Over time, the best actions over each state are learned through travel over the path of state transitions. Once the agent learns the actions that produce the highest reward, they may choose that action (i.e., exploitation) or choose a random one (i.e., exploration) [13]. In the context of JITAI, balancing exploitation and exploration is critical, as people change with time. Occasional exploration is needed to re-adapt intervention policies to the user.

3. ARCHITECTURE OF A JITAI SYSTEM

We propose a JITAI three layer architecture for stress in concordance with our concept of a stressful event. We represent stress as a continuous variable, and a stress event as an increasing wave that reaches a peak or point of inflection at the time the user experiences maximum stress. The overall goal of a JITAI system for stress is to maintain stress at an acceptable minima. We note that this representation of stress is a simplification, but is similar to previous physiological proxies for stress [2,4,15].

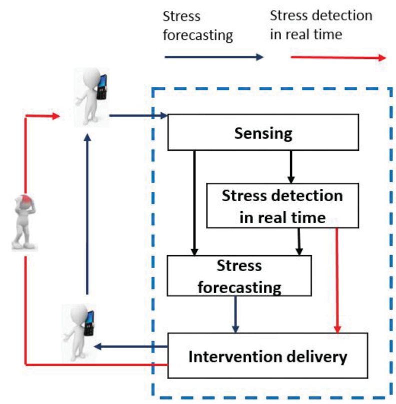

Figure 1: Flow of the Top-down stress three layer architecture

The first layer continuously senses human physiology. The second layer has two modules, one for stress inference in real time and the another one for stress forecasting. The third layer examines the outputs from the second layer and makes decisions as to how and when to intervene. Figure 1 outlines the proposed architecture emphasizing the advantages of using predicted stress values as input into intervention layer. Here, black lines represent the flow and connection between the architecture layers when forecasted values are used. The goal of anticipate an stressful event is to avoid the user pass through a traumatic experience.

3.1 Sensing Layer

As in other context-aware systems, the first layer is where sensing of user behavior and health state occurs. It consists of sensing hardware and software to pass sensed samples on to the next layer. For example, in the experimental work described in later sections, the sensing layer consists of a two-lead ECG and a Bluetooth connection to transmit captured samples to a smart phone.

3.2 Recognition and Forecasting Layer

The second layer has two components, a real-time recognizer and a real-time forecaster. The recognizer processes samples from the sensing layer to determine the health state of the user. The forecaster uses the current user's health state (as determined by the recognizer) and other contextual information to predict what the user's health state will be in the near future. Thus the forecaster has a temporal dependency on the recognizer. First, the user's health state at time t is determined. Then, the forecaster can predict what this health state will at time $t+\delta$. This information is then passed on to the intervention layer, where the decision to deploy the intervention can be made.

For example, in the experimental work described in later sections, raw ECG data is passed from the continuous sensing layer to the stress recognition layer. There, the ECG data is processed into heart rate variability (HRV), a proxy for continuous stress. At the same time, HRV data is passed to a stress forecaster. For every minute t, the forecaster produces a prediction of stress level at $t+1$, $t+2$,and $t+3$ minutes in the future.

3.3 Intervention Layer

The goal of this layer is to deliver a set of stress relieving interventions. The system must choose from a set of interventions, the right intervention or the right combination of them, and deploy at the right time and dosage to maximize the effectiveness while also reducing the number of interventions.

In the remainder of this paper, we discuss, simulate and test algorithms to forecast future stress (Section 4) and to choose the best intervention (Section 5).

4. FORECASTING FUTURE STRESS

The stress forecasting algorithm is based on a first-order Hidden Markov Model (HMMs) and uses the criterium of the maximum likelihood as the prediction method. Algorithm 1 depicts the proposed approach.

HMMs are a probabilistic graphical modeling approach used widely to model physical, behavioral, and social phenomena. They are useful to model situations which involve time and processes whose states are not evident, but whose manifestations can be observed through sensor readings [11].

A HMM is characterized by the following elements: A set of hidden process states C; a set of *observable or sensor readings* X; the number of states m; a *State Transition Probability Matrix* Γ, $P(C_t = j|C_{t-1} = i)$; a set of *Initial State Probabilities* δ; and an *Emission Probability Matrix* Λ, $P(X_t = j|C_t = i)$. Here, the sum of the elements of matrices' rows Γ, Λ, and δ should be equal to 1. Algorithm 1 operationalizes the forecasting process.

Algorithm 1: Forecasting Discrete Poisson-HMM

input : Observations X, Number of states m, Emission Probability Matrix Λ, Transition Probability Matrix Γ, Initial states probability δ, prediction horizon H, X range $Xrange$
output: $statepred$, Matrix of $X \times m$ of probabilities predictions
begin
- $xrange \leftarrow qpois(0.001, min(\Lambda)) : qpois(0.999, max(\Lambda))$
- $n \leftarrow length(X)$
- $la \leftarrow forw(X, m, \Lambda, \Gamma, \delta)$
- $c \leftarrow max(la[, n])$
- $llk \leftarrow log(sum(exp(la[, n] - c)))$
- $statepred \leftarrow zeros(m, H)$
- $R1 = exp(la(, n) - llk)$
- $R2 \leftarrow diag(m)$
- **for** $i \leftarrow 1$ **to** H **do**
  - $R2 \leftarrow R2 * \Gamma$
  - $statepred(:, i) \leftarrow R1 * R2$
- **end**
- $prob \leftarrow outer(X, \Lambda, dpois)$
- $dist \leftarrow$prob%*%$statepred[, 1 : H]$
- $print(Xrange$,dist) **return** S

end

The HMM forecasting algorithm takes the form of Algorithm 1. The state prediction starts at Line 3 by the use of the forward algorithm whose output is a matrix of probability weights α. The column vector α_t corresponds to the probability of occurrences of state C_i, at time t. In order to avoid overflow, in Line 5 the value of likelihood obtain at Line 4 is re-scale in the log domain. Then, in Lines 9, 10, and 11 the $for-loop$ fills the $h \times m$ matrix *statepred* with $P(C_{T+h} = i|X_{1:T}) = \alpha_T \Gamma^h(, i)/L_T$ which is the conditional probability that the hidden process at time C_{T+h} is in state i . Finally, at line 14 the $h \times m$ matrix *dist* stores $P(X_{T+h} = i|X_{1:T})$ which corresponds to the likelihoods of the potential candidates to be the next sequenceâĂŹs element in the time series.

In summary, when the current stress level is determined, it is added as a new element of the training set. Then, the Expectation Maximization Algorithm (EM) is used to optimize the initial parameters and the whole prediction model is recomputed using the new training set. Once the model is created, Algorithm 1 computes the likelihood for the potential candidates to be the next sequence's element and pick the one with the highest likelihood as prediction.

5. MODELING JITAI FOR STRESS

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). Reinforcement Learning (RL) [13] is 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.

5.1 Defintions

We define a stress episode as the interval between the moment stress is detected or forecasted above a minimum acceptable level of stress to the moment stress is once again below the accepted minimum level of stress (set as zero here, without loss of generality).

Let $A = (a_1 \dots a_M)$ be the possible interventions that can be delivered to the patient during an episode of stress. Let $A_t = \{a_i\} \in$

A be interventions given to the patient so far in the episode, and we define the Q-Learning (QL) state (different from the user state) at time t as $s_t = A_t$. Note that A_t is an (unordered) set, meaning that 1) each intervention can only be used once in any one stress episode, and 2) the order in which the interventions are delivered does not affect the outcome. In reality, reusing an intervention in an episode and/or varying the order in which the interventions are delivered can affect the outcome of the intervention. Nonetheless, this simplification 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 $\pi : S \rightarrow A$, where S is the power set of 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, were ϵ depended on the relation between the value of the best known choice and the sum of all values for that state. Thus, exploration was performed with a probability as stated in Equation 1.

$$p(exploration, s) = (1 - \frac{Q(s, a^*)}{\sum_a Q(s, a))}) \quad (1)$$

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

Algorithm 2: QL(1) intervention delivery system.

```
input  : stress levels detected or forecasted
output: the set of interventions to deliver
begin
    forall the 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
```

6. EXPERIMENTS AND RESULTS

In this section, we describe experiments and results for forecasting and adaptive intervention.

6.1 Data

We use a real-world continuous physiological dataset in our experiments described below. As part of a larger study on the inferring stress from a suite of physiological measures, the data was collected by 20 students at the University of Memphis. Participants wore the AutoSense [2] system underneath their clothes for seven days as they went about their normal daily lives. AutoSense includes a range of sensors selected to measure stress. In this paper, we only examine the ECG data and use Heart Rate Variability (HRV) as a proxy for stress. HRV measurements are reported as averages over consecutive one minute windows.

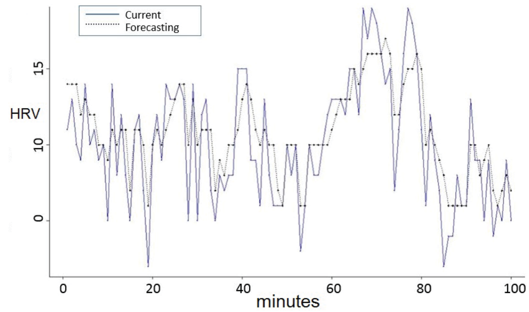

Figure 2: A step-ahead forecasting using a four hidden states Poisson HMM

Table 1: Initial parameters common to all experiments

Parameters	Value
Number of experiments	3
Number of participats	10
Number of repetitions per experiment	15
Time series size	4506 minutes of HRV
Distribution	Poisson
Parameter estimation	EM Algorithm
Forecasting range	100
Forecasting horizon (h)	1-3
Training set size	1000
Type of training set	Sliding Window
Lambda	kmeans(training-set, m)

In order to use the forecasting advantages of HMM, the HRV time series was discretized. The discretization process was carried out by using the Symbolic Aggregate Approximation Algorithm (SAX) [7]. We transformed a 4506 length continuous time series (i.e., 4506 minutes) into a 4506-symbol string time series, with alphabet size 20.

6.2 Stress Forecasting

A set of three experiments were carried out to measure the capacity of HMMs to predict stress. In all the experiment, the training set corresponds to a sliding window of a thousand minutes of HRV which runs through a time series.

Every minute, the sensing and stress recognition layer provides a new HRV value x_{t+1} (i.e., HRV computation) which is added as a new element of training set. At the same time, the first element x_k of the training set is removed (i.e., sliding window movement) resulting in the new training set $x_{k+1} \ldots x_{t+1}$ or X_{t+1}.

Thereby, every minute, a training set X_i along with the initial parameters described in Table 2 are used as input of the Expectation Maximization algorithm (EM) [11] to optimize the initial parameters and re-compute the HMM model.

Then, this new model is used as input of Algorithm 1 which computes the likelihood for the potential candidates to be the next sequence element in the time series and pick the one with the maxi-

Table 2: Experiment specific initial parameters

Parameters	Experiment 1	Experiment 2	Experiment 3
Hidden states (m)	4	3	2
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	$(0.25, 0.25, 0.25, 0.25)$	$(0.333, 0.333, 0.333)$	$(0.5, 0.5)$

mum likelihood as prediction.

Each experiment was repeated fifteen times to reach statistical significance, and for each repetition, we used data from ten participants. An instance or repetition consists of forecasting a hundred minutes (e.g., using the one, two, or three minutes-ahead forecasting approach) of HRV using the parameters specified in Table 2.

For each repetition, we computed the metrics values: ME, MAE, RMSE, NRMSE, d, and r^2 described in Table 5 for each participant. At the end of the repetition, the metric values of the ten participants are averaged. Finally, at the end of the fifteen repetitions, we computed the average of the average values obtained for each of the repetitions.

We carried out a set of three experiments, and for each experiment, we used the one, two, and three minutes-ahead forecasting approach. We repeated experiments one, two, and three with four, three, and two hidden states. Table 1 summarizes the initial parameters common to all experiments, while Table 2 shows the specific set of parameters for each experiment.

Table 3 summarizes the results for the three experiments. The first, second, third, and four table rows show the experiment number, the forecasting approach (i.e., one, two, and three minutes-ahead), the number of hidden states used in that experiment, and the metric values computed for each experiment.

For instance, the last row under experiment 1 shows that the r^2 values obtained from forecasting the HRV time series for one, two, and three minutes-ahead using a HMM with four hidden states were 0.71 0.26 0.18 respectively.

Figure 2 shows a forecasting instance for participant 1. Here, we used a HMM with four hidden states to forecast a minute-ahead during a period of one hundred minutes.

The type of forecasting used in this project is not recursive , where a forecasted value x_t is used to predict the next element x_{t+1}. Instead, the Sensing and Stress Recognition layer provide us a new x_t value every minute, which in turns is used to forecast the x_{t+1}, x_{t+2}, x_{t+3} values at the same time.

As can be observed in Table 3 there is little difference in terms of the Coefficient of determination (r^2) and Index of Agreement (d) when the HMM forecaster was used with four and three hidden states. However, in terms of computational cost, prediction using four states is more costly. Thus, using three states offers a good trade-off between prediction accuracy and computational cost.

Table 4: QL parameter search data

Parameter	Min Value	Step	Max Value	Best value
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

6.3 Stress Intervention

We simulated the intervention system using a mathematical model of the patient reaction to the interventions. Carrying out simulations before testing with actual subjects 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. The platform also provides a suitable testbed to try out new ideas before deployment with real users.

We modeled a patient by creating an artificial distribution $p(s|t_1, \ldots, t_N)$ corresponding to the factor graph shown in Figure 4b. The probability of not having stress under a set of interventions was sampled from a beta distribution with the parameters $\alpha = 0.1$ and $\beta = 1$. Here, most of the probability mass is concentrated near zero. Thus, most of the intervention combinations will be inefficient, while a few of them will have a greater chance of relieving the patient.

Then, an exhaustive search of the QL algorithm parameters was made. Table 4 summarizes the parameters, their search space and the best result. Each set of parameters was evaluated over ten simulations with different random patient models. The mean number of interventions was the chosen metric to minimize. The obtained set of parameters were 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. We also implemented a simple system that picked interventions at random. This system served as a way of measuring how difficult it was to relieve stress under the current models.

In the case of QL, the mean number of interventions per episode for each system across all simulations was around 3 versus 8 in the random case. Figure 3 shows a the cumulative average of the number of interventions per episode of one of these experiments.

7. DISCUSSION

7.1 Stress Forecasting

Table 3: Experiment results

	Experiment 1			Experiment 2			Experiment 3		
Predict approach (minutes-ahead)	1	2	3	1	2	3	1	2	3
Hidden States		4			3			2	
MAE	1.88	2.99	3.19	1.91	3.12	3.49	2.16	3.14	3.27
d	0.87	0.67	0.60	0.89	0.64	0.56	0.87	0.72	0.69
r^2	0.71	0.26	0.18	0.7	0.2	0.12	0.59	0.24	0.2

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
```

Table 5: Evaluation Metrics Used in this paper

Metric	Description	Definition
MAE	Mean Absolute Error	$\frac{1}{N}\sum_{i=1}^{N}\left\|(S_i - O_i)^2\right\|$
d	Index of Agreement	$1 - \frac{\sum_{i=1}^{N}(O_i - S_i)^2}{\sum_{i=1}^{N}(\|S_i - \hat{O}\| + \|O_i - \hat{O}\|)^2}$
r^2	Coefficient of determination	$1 - \frac{SS_{res}}{SS_{tot}}$

A disadvantage of HMMs is its off-line nature. Our forecasting algorithm uses a portion of a HRV time series as a training set to create the prediction model (i.e., transition probability matrix) and the remaining data as testing. This approach does not take into account new incoming measurements to update the model. We addressed this problem by re-building the entire model as new samples arrived. However, this pseudo-online application of HMM has a high computational cost. Thus, an immediate challenge to address is how to update the transition probabilities faster. A new generation of algorithms that partially address this problem are DENFIS [5]and FIS. These methods 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.

7.2 Just in Time Intervention Systems

The question of whether an intervention's outcomes are independent of each other or not plays an important role in the design of JITAI 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 effectiveness.

However, the medical consensus is that previous interventions may influence the effectiveness of new ones. Thus, we use a graphical model that integrates the stochastic outcome of different interventions. Figure 4a 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 factors expressing the effectiveness of each intervention under the patient attributes, according to Equation 2. Note that, in this case, intervention outcomes are indeed independent of each other.

$$p(s, a_1, .., a_M, i_1, .., i_N) = \prod_{j=1}^{j=N} \phi(s, a_1, .., a_M, i_j) \qquad (2)$$

As the full set of attributes can never be known, i.e. the factors that affect the effectiveness of each intervention cannot be fully determined, we focus on the marginal distribution $p(s, i_1, ...i_N)$. Due to marginalization of the attribute variables, the graphical model becomes that shown in Figure 4b. Here, the effectiveness of each intervention becomes dependent on one another. Consequently, deciding which intervention to apply at a given moment must take into account the outcomes of previous interventions.

The previous insight also suggests that additional context could improve this solution. For example, context information such as the time of the day and the patient's location could give clues as to the cause of 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 on physiology). Finally, an a priori questionnaire could be used to gather information on the patient's preferences of possible interventions.

8. CONCLUSION AND FUTURE RESEARCH

In this paper, we describe a three layer architecture for just-in-time adaptive intervention, in which there are two components that are different than what has been presented previously. The first is a forecaster that can predict future health states. The second is an intervention layer that chooses the best intervention at a given time. We then presented algorithms that drive these two components and demonstrated their feasibility in simulation.

We tested a HMM-based forecaster using ECG data collected from multiple users in the natural environment. We show it is possible to forecast HRV, a proxy for stress, with up to 89% of coefficient of agrement and r^2 of 71%. We also tested the intervention layer via simulation of a stressed patient. We use Q-Learning with intelligibility traces to select the best intervention from a set of many.

A key assumption of this work is in the shape and nature of the stress signal. We assume stress can be viewed as a continuous signal and that stress events are waves or bumps in the signal. As future work, we would like to better characterize the stress signal's shape. A better understanding of the stress' signal behavior may help us understand when, and for how long to intervene, as well as better assess the effectiveness of the interventions.

We also want to address the challenges in terms of online learning in both the prediction and the intervention layers. The introduction of a priori knowledge of the patient preferences for treatments and context information is important to accelerate the intervention layer's learning process and increase each intervention success rate.

Lastly, we would like to implement a full system and run studies in a mobile environment to explore the acceptability of JITAIs. We want to explore the burden imposed by JITAIs (i.e., what is the level of user tolerance to false positives) as well as how incentives can encourage user participation.

9. ACKNOWLEDGMENTS

We would like to thank Dr. Santosh Kumar and the WiSeMANET lab at the University of Memphis for the stress dataset used in this work. The present work was partially supported by NSF grant #1117303 and by the "Agencia Nacional de Investigacion e Innovacion, Uruguay".

10. REFERENCES

[1] Bibhas Chakraborty and Susan A Murphy. Dynamic treatment regimes. *Statistics*, 1, 2014.

[2] 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 *SenSys*, pages 274–287, 2011.

[3] 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.

[4] Luis G Jaimes and Andrew Raij. Forecasting biorhythms for preventative stress interventions. In *Proceedings of the ACM Workshop of Biological rhythms and technology, at CHI 2014*. ACM, 2014.

[5] 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.

[6] 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.

[7] 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.

[8] Hong Lu, Denise Frauendorfer, et al. Stresssense: Detecting stress in unconstrained acoustic environments using smartphones. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, pages 351–360. ACM, 2012.

[9] Susan A Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):331–355, 2003.

[10] Inbal Nahum-Shani, Min Qian, et al. Q-learning: A data analysis method for constructing adaptive interventions. *Psychological methods*, 17(4):478, 2012.

[11] Lawrence Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.

[12] Sarah Robinson, Sarah Perkins, Stephanie Bauer, Neil Hammond, Janet Treasure, and Ulrike Schmidt. Aftercare intervention through text messaging in the treatment of bulimia nervosaŮfeasibility pilot. *International Journal of Eating Disorders*, 39(8):633–638, 2006.

[13] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*, volume 1. Cambridge Univ Press, 1998.

[14] 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.

[15] Tanja GM Vrijkotte, Lorenz JP van Doornen, and Eco JC de Geus. Effects of work stress on ambulatory blood pressure, heart rate, and heart rate variability. *Hypertension*, 35(4):880–886, 2000.

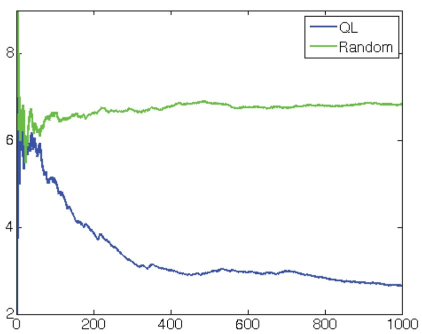

Figure 3: Cumulative average of the number of interventions per episode over one of the experiments

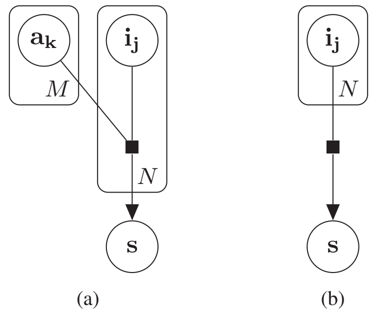

Figure 4: Graphical models of the patient's stress. Model (a) depicts the whole joint distribution while (b) represents the model after a_k has been marginalized out.