Adaptive Symptom Monitoring using Hidden Markov — an Application in Ecological Momentary Assessment

DOI:

[10.1109/JBHI.2020.3031263](https://doi.org/10.1109/JBHI.2020.3031263)

Document Version

Accepted author manuscript

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):

Hulme, W. J., Martin, G. P., Sperrin, M., Casson, A. J., Bucci, S., Lewis, S., & Peek, N. (2021). Adaptive Symptom Monitoring using Hidden Markov — an Application in Ecological Momentary Assessment. *IEEE Journal of Biomedical and Health Informatics*, 25(5), 1770-1780. [9226077]. <https://doi.org/10.1109/JBHI.2020.3031263>

Published in:

IEEE Journal of Biomedical and Health Informatics

Citing this paper

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights

Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.

Adaptive Symptom Monitoring using Hidden Markov Models — an application in Ecological Momentary Assessment

William J Hulme, Glen P Martin, Matthew Sperrin, Alexander J Casson, Sandra Bucci, Shôn Lewis, and Niels Peek

Abstract—Wearable and mobile technology provides new opportunities to manage health conditions remotely and unobtrusively. For example, healthcare providers can repeatedly sample a person’s condition to monitor progression of symptoms and intervene if necessary. There is usually a utility-tolerability trade-off between collecting information at sufficient frequencies and quantities to be useful, and over-burdening the user or the underlying technology, particularly when active input is required from the user. Selecting the next sampling time adaptively using previous responses, so that people are only sampled at high frequency when necessary, can help to manage this trade-off.

We present a novel approach to adaptive sampling using clustered continuous-time hidden Markov models. The model predicts, at any given sampling time, the probability of moving to an ‘alert’ state, and the next sample time is scheduled when this probability has exceeded a given threshold. The clusters, each representing a distinct sub-model, allow heterogeneity in states and state transitions.

The work is illustrated using longitudinal mental-health symptom data in 49 people collected using ClinTouch, a mobile app designed to monitor people with a diagnosis of schizophrenia. Using these data, we show how the adaptive sampling scheme behaves under different model parameters and risk thresholds, and how the average sampling can be substantially reduced whilst maintaining a high sampling frequency during high-risk periods.

Index Terms—mHealth; Ecological Momentary Assessment; Ecological Momentary Intervention; Digital Phenotyping; Hidden Markov Models; Risk Prediction; Schizophrenia; Adaptive Sampling.

I. INTRODUCTION

THE increasing availability of wearable and mobile technology provides opportunities to monitor, predict, and manage long-term health problems in real-time, in everyday settings, and with minimal disruption. A common framework for this is Ecological Momentary Assessment

(EMA), defined as the “repeated sampling of people’s current thoughts, emotions, behaviour, physiological states, and context, in their natural environment, typically (but not necessarily) via electronic wearable devices” [1]. EMA offers distinct advantages over retrospective interview- or questionnaire-based assessment methods, as the sampling occurs in situ and at higher frequency. This increases ecological validity, reduces recall bias, and is often cheaper and less intrusive. EMA, or Ecological Momentary Intervention (EMI) as it is called if there is an additional interventional component [2], is primarily used in psychology with sampling frequency typically ranging from hourly to weekly. However, EMA has a range of applications and we consider EMA here to include any mobile assessment tool collecting physiological, psychological, or behavioural data via either passive or active monitoring devices. Recent examples include investigating the relationship between sleep quality and suicidal ideation [3], pollution and allergies [4], and the weather and arthritis [5].

To ensure that EMA systems are deployed effectively, there is a utility-tolerability trade-off between collecting information with sufficient frequency to be useful, and over-burdening the user or the underlying technology. Tolerability may be reduced if the sampling frequency is too high, for instance due to app-fatigue leading to disengagement, or by battery limits which prevent the device being used as frequently as intended. Conversely, utility may be reduced if the sampling frequency is too low, as responses are not abundant or timely enough for appropriate monitoring and care.

In many EMA settings where sampling times are pre-set or requested at fixed intervals, balancing utility and tolerability is a crude exercise in choosing the fixed sampling frequency that represents the best compromise. By contrast, if sampling times were chosen adaptively in real-time, with high frequency sampling during “interesting periods” and low frequency sampling during “uninteresting periods”, this trade-off might be managed more effectively. Indeed, the overall sampling frequency could be reduced while increasing the utility of the collected data as samples are taken at the most informative and relevant times.

We propose the use of continuous-time Hidden Markov Models (CTHMMs) as a means to implement a generic representation of real-time, adaptive sampling in EMA-type settings. A CTHMM is a type of state-switching

W. J. Hulme, G. P. Martin, M. Sperrin, and N. Peek are with the Centre for Health Informatics, Division of Informatics, Imaging and Data Science, School of Health Sciences, University of Manchester, United Kingdom.

A. J. Casson is with the Division of Materials, Devices and Systems, School of Electrical and Electronic Engineering, University of Manchester, United Kingdom.

S. Bucci and S. Lewis are with the Division of Psychology and Mental Health, School of Health Sciences, University of Manchester, United Kingdom.

model where unobserved states are inferred from irregularly-observed response sequences [6]. In this paper we use the term “alert state” to denote a state that is of particular interest, and requires close monitoring or attention. We fit clustered CTHMMs to describe the underlying state-switching process and the distribution of observed responses in each state and use this model to predict in real-time the probability of visiting a state at any time in future using the sampled responses so far (or a subset thereof). The state predictions are used as a basis for selecting the next sampling time, where the sampling time is sooner if the probability of visiting an alert state is high. What constitutes an alert state is context dependent, and requires expert input. For instance, it could be when a person is experiencing distressing or fluctuating symptoms, or is physically active or travelling. It may also differ from person to person if there are phenotypes of responses and state-switching patterns. If they exist, these phenotypes can be exploited to improve the relevancy of the model for each person, hence our use of clustered CTHMMs where each cluster represents the distinct set of states through which individuals of the same phenotypes can move.

Adaptive sampling schemes have been considered previously for reducing energy consumption in passive monitoring scenarios such as accelerometry data [7], [8] or geolocation data [9], and this has sometimes included the use of Markov or hidden Markov chains. For example, [10] a discrete-time hidden Markov model is used to infer the state of the user but, unlike in our proposal, the classifier that determines whether to increase or decrease the sampling rate (based on *reward-inaction learning*) is independent of the HMM. Another example is [8], where the approach was to use a discrete-time Markov chain to describe user activities over time. Activities were classified without error at the sampling times, which were selected to maximise the probability of correctly classifying the user’s activity between sampling times, subject to a limited budget (e.g., battery life or maximum samples). HMMs have also been employed with EMA-type data to model psychological processes retrospectively [11], [12]. However, while many works have looked at different aspects of mobile sampling in a dynamic or non-constant way, there are none for EMA-type settings and the use of CTHMMs for the purposes of adaptively sampling in real-time has not been considered.

The aim of this paper is to introduce a CTHMM approach for adaptive sampling in EMA-type systems, motivated by and demonstrated using multivariate psychotic symptom data collected using the mobile app Clin-Touch [13]. We investigate how the burden of continuous active symptom monitoring could be reduced to improve engagement and allow better harmony with the flow of everyday life. Finally, we consider a number of potential extensions to this approach for personalised monitoring and interventions.

II. CTHMMs AND ADAPTIVE SAMPLING

This section begins with an overview of CTHMMs and introduces the notation used throughout the paper. It next

considers the use of clustering to account for heterogeneity between different independent sequences when fitting a CTHMM. The use of a pre-trained CTHMM to predict the underlying state for a new response sequence, both at and beyond the latest sampling time, is described. Finally, the proposed adaptive sampling scheme is then introduced.

A. Overview of CTHMMs

CTHMMs are a class of state-switching models for irregularly-observed temporal data, where the distribution of the (observed) response is conditional on the underlying (unobserved) state. HMMs are more flexible than standard Markov models as they allow for latent state membership. This allows the models to be applied where no simple state label pre-exists in the data. Transitions between states follow a continuous-time Markov process, i.e., the probability of moving to another state is independent of the past state sequence given the current state. States are inferred from the observed responses. For a pre-determined number of states, each assumed to correspond to a given distribution family, the fitted model identifies the set of state-specific probability distributions and state transition intensities that fit to the observed sequence best (usually with maximum likelihood).

An accessible overview of HMMs is provided in [14], [15] and an in-depth text provided in [6].

The notation is as follows: A CTHMM describes the observed response sequence $Y(t)$ and a hidden state sequence $Z(t)$ evolve according to a transition matrix ϕ and a response distribution F_θ . Specifically, we have:

- Response sequence $Y(t)$, which could be multivariate. If $Y(t)$ is a D -variate response then $Y(t) = [Y^1(t), \dots, Y^D(t)]$. Here we assume the $Y^d(t)$ s are independent conditional on the hidden state $Z(t)$, i.e., $\Pr(Y^d(t) \leq y | Z(t), Y^1(t), \dots, Y^{d-1}(t), Y^{d+1}(t), \dots, Y^D(t)) = \Pr(Y^d(t) \leq y | Z(t))$. Joint multivariate distributions are possible but not considered here.
- Hidden state sequence $Z(t) \in \mathcal{K}$, for some finite state space $\mathcal{K}$ with $|\mathcal{K}| = K$. $Z(t)$ is a continuous-time Markov chain.
- State transition intensity matrix $\phi = \{\phi_{kl}\}$, for $k, l \in \mathcal{K}$, where $\phi_{kl} = \lim_{\epsilon \rightarrow 0^+} \Pr(Z(t + \epsilon) = l | Z(t) = k) / \epsilon$ is the (k, l) entry of ϕ and represents the instantaneous risk of $Z(t)$ moving from state k to state l for $k \neq l$, and $\phi_{kk} = -\sum_{l \neq k} \phi_{kl}$.
- Response (or emission) distributions F_{θ_k} for $k \in \mathcal{K}$, where $F_{\theta_k}(y) = \Pr(Y(t) \leq y | Z(t) = k)$ is the distribution function of $Y(t)$ for each state k , and θ_k is the set of distribution parameters for state k .
- Initial state occupancy probability vector $\alpha = [\alpha_1, \dots, \alpha_K]$, where $\alpha_k = \Pr(Z(t_1) = k)$.

$\mathcal{M} = (F_\theta, \phi, \alpha)$ defines the CTHMM. For $Y(t)$ observed at sampling times $t_1, \dots, J$, we have a vector (or matrix in the multivariate case) of observations $y_{1:J} = [y(t_1), \dots, y(t_J)]$. $\mathcal{M}$ is estimated by maximising over $Z, \phi, \theta_k, \alpha$ the joint likelihood of observations $y_{1:J}$ and

hidden states $Z(t)$ conditional on the response distributions F_θ and initial state probabilities α .

Model estimation is extended here to include multiple independent sequences with independent sampling times over index i , giving t_{ij} , J_i , $y_i(t_j)$, $Z_i(t)$, and possibly $\pi_i(t_1)$, for $i = 1, \dots, N$. Further extensions that are not considered here include allowing transition intensities and response densities to depend on (possibly time-varying) covariates $X_i(t)$, so that $\phi = \phi(X_i(t))$ and $\theta = \theta(X_i(t))$.

B. Missing values

Since CTHMMs deal with data arriving at irregular times, missing values are dealt with naturally; if $y(t_j)$ is not available at sampling time t_j then $y(t_j)$ is removed from the set of responses $y_{1:J}$ and the index-set $\{j\}$ contracts by one. If $y(t_j)$ is partially observed (some but not all of $y(t_j)$ is available at t_j), then $y(t_j)$ still contributes to the likelihood and can still be used for state prediction, again with missing values assumed to be missing at random.

However, this does not exploit information in a missing-not-at-random scenario, which is likely to hold in many scenarios involving remote monitoring. For instance in ClinTouch, our motivating example, a person might be less likely to respond if they are experiencing heightened delusions. For simplicity, we assume non-response occurs non-informatively within this study, but revisit how models might be extended to handle other missingness mechanisms in the Discussion.

C. Clustering

A one-size-fits-all approach for predicting unobserved states and future sampling times may be unsuitable amongst a set of sequences that exhibit significant heterogeneity. For instance, psychotic symptoms may have unpredictable timings, duration, and severity, and can differ substantially from person to person. It is important therefore to account for these clinical subtypes, particularly for prediction problems where population effects are not of interest, as is the case for individualised adaptive sampling.

The approach taken here is to identify homogeneous clusters of people, each representing the same phenotype, for whom a shared model adequately describes $y_i(t_j)$. This can be achieved by restricting the transition matrix ϕ to be block-diagonal, so that transitions from one state to another are only permitted between states belonging to the same block[16], thus forcing the CTHMM to identify clusters of people with shared response distributions. For example, a two-cluster, five-state model with two and three states in each cluster respectively, will have a transition matrix $\phi = \begin{bmatrix} \phi^1 & 0 \\ 0 & \phi^2 \end{bmatrix}$ where ϕ^1 is the 2×2 transition matrix for cluster 1 and ϕ^2 is the 3×3 transition matrix for cluster 2. Since transitions between clusters are not possible and the response distributions are independent, each cluster c can be considered as a CTHMM on its own, $\mathcal{M}^c$, with its own transition matrix $\phi^c = \{\phi_{lk}^c\}$ and set of response distributions $F_{\theta^c} = \{F_{\theta_k^c}\}$ for $k, l \in 1, \dots, K_c$. As

with states, cluster membership is estimated from the data rather than prescribed; configurations of cluster numbers and states per cluster are pre-defined. Note that setting the number of clusters to the number of independent sequences is equivalent to having one model per series.

For our application, this approach strikes a pragmatic balance between plausibility and practicality owing to the small number of sequences available for clustering and the short sequence length in our dataset. It also provides multiple phenotype-specific models on which to assess the adaptive sampling scheme that is proposed.

It is important to be clear about the distinction between the role of states and clusters. States partition the person-specific response sequence $Y(t)$ through time and correspond to homogeneous clinical statuses, such that responses observed whilst in the same state arise from the same distribution. Clusters partition the set of independent sequences into distinct sub-models, corresponding to homogeneous clinical phenotypes, with people in the same phenotype moving through the same set of states.

D. State prediction

The following section outlines how to use a known CTHMM $\mathcal{M}$ to estimate the underlying state $Z(t)$ for any time t , including beyond the latest sampling time t_J . The model $\mathcal{M}$ is assumed fixed and we are not learning more from future data. Here, $y_{1:J}$ can be any response sequence whose distribution is assumed to be adequately captured by $\mathcal{M}$ (i.e., $\mathcal{M}$ is well-validated).

1) *State prediction at a sample time t_j* : Here we consider the estimation of the states $Z(t_j)$ at each sampling time t_j for $j = 1, \dots, J$ using model $\mathcal{M}$ and data $y_{1:J}$. Let

$$p(t_j, y_{1:J}) := \left\{ \Pr(Z(t_j) = k \mid y_{1:J}, \mathcal{M}) \right\}_{k=1:K},$$

that is, the probability that $Z(t_j)$ is in state k conditional on $y_{1:J}$ and the model $\mathcal{M}$. This problem is known as *smoothing* in the HMM literature. This becomes analytically intractable when J is large since we need to marginalise over all possible state sequences, which increases exponentially with J (the number of permutations for $Z(t)$ at times $t_1, \dots, t_J$ is K^J). In this case, $p(t_j, y_{1:J})$ is obtained using the *Baum-Welch*, or *forward-backward*, algorithm (see [17], [18] for an implementation in discrete-time HMMs and [19] for an extension in continuous-time). For estimation of $p(t_J, y_{1:J})$ (the last sampling time), known as *filtering* in the HMM literature, this algorithm reduces to the *forward* algorithm.

If required, p can be calculated using less than all available samples. For example, $p(t_J, y_{j:J})$ estimates $Z(t_J)$ using only samples from t_j onwards. This may be useful if there are storage or retrieval costs for $y(t)$ or if computation time for the forward-backward algorithm is high. If only a single sample y_j is used, then the problem reduces to a finite mixture model with $p(t_j, y_j) = \left\{ \frac{f_k(y_j \mid \theta_k)}{\sum_{l \in 1:K} f_l(y_j \mid \theta_l)} \right\}_{k=1:K}$ where f_{θ_k} is the probability density function corresponding

to F_{θ_k} . ϕ is not needed here since we are only using information at one time-point, t_j , and so transition probabilities become irrelevant.

2) *State prediction at time $t_J + u$* : Let $p_J = p(t_J, y_{1:J})$ be the vector of state probabilities at time t_J conditional on responses $y_{1:J}$, with elements of p_J summing to 1, i.e., $\mathbf{1}^\top p_J = 1$. Consider the estimation of future states beyond the latest response time. That is, $Z(t_J + u)$ for some $u > 0$ using model $\mathcal{M}$ and data $y_{1:J}$. This is a simple multiplication of p_J with $\text{Exp}(u\phi)$, where Exp is the matrix exponential and $[\text{Exp}(u\phi)]_{kl} = \Pr(Z(t_J + u) = l \mid Z(t_J) = k)$ is the probability of being in state l in u days time given that the process is in state k at time t_J , giving the *state probabilities at time $t_J + u$ given p_J* ,

$$\begin{aligned} SP(p_J, u) &= \left\{ \Pr(Z(t_J + u) = k \mid p_J, \mathcal{M}) \right\}_{k=1:K} \\ &= p_J^\top \text{Exp}(u\phi). \end{aligned}$$

Since $\text{Exp}(0\phi)$ is the identity matrix, $SP(p_J, 0) = p_J$.

Now consider the probability that a state k has been passed through at least once in the period $[t_J, t_J + u]$. This is implemented by setting row k of ϕ to zero, denoted ϕ^{0_k} , so that it becomes an absorbing state (i.e., a state that cannot be exited), giving the *passage probabilities at time $t_J + u$ given p_J* ,

$$\begin{aligned} PP(p_J, u) &= \left\{ \bigcup_{v \in [0, u]} \Pr(Z(t_J + v) = k \mid p_J, \mathcal{M}) \right\}_{k=1:K} \\ &= p_J^\top \text{Exp}(u\phi^{0_k}). \end{aligned}$$

Again, when $u = 0$, $PP(p_J, 0) = p_J$. This can be generalised to a set of states by setting all rows in ϕ corresponding to those states equal to zero. If ϕ includes no absorbing states then PP converges to $\mathbf{1}$.

E. The proposed adaptive sampling method

Here we introduce the method to adaptively select the next sampling time such that the sampling frequency is high during, or prior to, alert state $\bar{k}$ being reached, and low during non-alert states. Recall that the alert state corresponds to a high-risk or otherwise interesting state capturing the behaviour of the person under observation. Because the states are identified in an unsupervised fashion, expert examination of the model fit is required to determine which states should become alert states. The next sampling time is chosen based on the first passage probability $PP(p_J, u)$ exceeding a pre-defined threshold τ . Thus, the system will sample with higher frequency during periods when the person under observation is experiencing, or anticipated to experience, an alert state (for example as indicated by unhealthy responses), but with lower frequency when not in an alert state.

Formally, our next sampling time $t_{J+1} = t_J + u_J(\tau)$ is selected such that it is the earliest time where $PP_{\bar{k}}(p_J, u)$ is estimated to be equal to or greater than τ , for some pre-specified threshold $0 < \tau < 1$. That is,

$$\begin{aligned} u_J(\tau) &= \min_u \left\{ u > 0 \mid PP_{\bar{k}}(p_J, u) \geq \tau \right\} \\ &= \min_u \left\{ \left| p_J^\top \text{Exp}(u\phi^{0_{\bar{k}}}) \mathbf{1}_{\bar{k}} - \tau \right| \right\}. \end{aligned}$$

This second equality is because $PP_{\bar{k}}(p_J, u)$ is non-decreasing over $u_J(\tau) > 0$, so has a unique solution. Analytically, this is the solution to $p_J^\top \text{Exp}(u\phi^{0_{\bar{k}}}) \mathbf{1}_{\bar{k}} = \tau$ with respect to u but this has no closed-form solution, so must be solved numerically.

A pre-specified maximum sampling interval can be set to ensure that samples are not too infrequent, even if the risk of an alert state $\bar{k}$ is low. This ensures that unexpected transitions to alert states are not missed and that future state estimation is based on sufficiently many recent samples. Similarly, a minimum sampling interval can be set to ensure the next sample does not occur too soon after the current sample if the risk of alert state $\bar{k}$ is high (in our ClinTouch example below, requesting responses every minute is a way to guarantee app-fatigue).

III. AN APPLICATION TO ECOLOGICAL MOMMENTARY ASSESSMENT DATA

In this section the dataset from our ClinTouch example is described, followed by the model configurations that are considered as candidates for our simulations. Finally, we describe the methods used to assess the properties of the proposed adaptive sampling scheme in comparison to a fixed sampling scheme. We consider both simulated response data where the underlying state sequence is known and the actual ClinTouch response data where the underlying state sequence is not known.

A. ClinTouch data

ClinTouch is an EMA-type mobile app designed to monitor the symptoms of people who experience psychosis [13], [20]. It collects responses to a short mental health questionnaire, asking the respondent about the presence and severity of various psychological and behavioural symptoms experienced since the previous response. The user responds on a touch-screen slider with responses from 1 (healthy, no symptoms) to 7 (unhealthy, intense symptoms), with questions derived from the Positive and Negative Syndrome Scale (PANSS) [21]. Prompts for responses are sent at one of four pseudo-random sampling times per day, covering a subset of the questions each time.

Response histories are available within the app and are also uploaded to a central server for remote monitoring by healthcare professionals. An alert system is used which informs the care-team if a user has met pre-defined criteria indicating unhealthy mental status (for instance 3 consecutive days with a high score anxiety).

This paper uses data from the ClinTouch arm of two trials, CareLoop [ISRCTN88145142] and Actissist [ISRCTN34966555], [22]. The CareLoop trial assessed the safety of ClinTouch-enhanced symptom monitoring compared with routine monitoring in people with schizophrenia across two sites; one an early-intervention cohort in typically younger people experiencing their first interaction with psychotherapeutic services and the other a cohort of older people with long-term mental health issues. The single-site Actissist trial compared ClinTouch with the EMI

Fig. 1. Example ClinTouch user trajectories

app Actissist. In total, 49 ClinTouch users are available from the ClinTouch arms of the CareLoop ($n=38$) and Actissist ($n=11$) trials. This excludes three ClinTouch users who were removed due to drop-out within a week.

In both trials, users were instructed to use the app for 12 weeks, though non-response and early drop-out was common, with 33 of 49 participants (67%) still using the app in the final week of the 12-week observation period. The total number of non-missing responses across all users is 7719 out of a total of 14412 possible responses up to the last response day for each user, representing an average response rate of 53.6%.

For simplicity, the collected responses from these trials have been simplified into symptom severities for four *items*: Anxiety, Depression, Hallucinations, and Delusions, with symptoms kept on the original ordinal scale from 1 to 7. The simplified responses are equal to the mean value of all core (non-branching) questions, rounded upwards to the nearest integer. Responses thus remain on an ordinal scale from 1 to 7 which ensures the methods remain adaptable to individual ClinTouch responses in future work. More formally, for each participant we have observed response data of the form $Y = ([Y^1(t_1), \dots, Y^D(t_1)], \dots, [Y^1(t_J), \dots, Y^D(t_J)])$, where D is the number of items, J is the number of response times, $Y^d(t_j) \in \otimes^d = \{\emptyset, 1, 2, 3, 4, 5, 6, 7\}$ is the user's response at time t_j for item d and $\emptyset$ indicates either no question asked or no response (which are indistinguishable in the dataset provided).

Example response trajectories from two users are presented in Figure 1.

B. Modelling the data

In summary we fit the model using all subjects. We fit K -state models with C clusters for a variety of choices of C and K , using BIC to select an optimal model.

1) *Response distributions*: We consider two approaches to modelling the distribution of the response variable $Y(t)$.

The first is to treat the responses as categorical, with $\theta_r = \Pr(Y(t) = r)$, for $r = 1, \dots, 7$ and $\sum_r \theta_r = 1$ for

each item. This ignores the ordinality inherent in $Y(t)$ but is a conceptually simple model and is maximally flexible, since it is capable of representing all possible response distributions for the given response-space $\{1, \dots, 7\}$.

The second is to treat the response variable as ordinal using Cumulative Link Models (CLMs) [23] with a latent Normal response. The advantage of CLMs in this context is that, rather than estimating $7 - 1 = 6$ independent parameters for each state in each cluster, we can estimate a set of 6 shared threshold parameters across each cluster with one additional location parameter per state, which reduces the parameter space considerably. More details of this approach are available in the Appendix.

2) *Clusters and states*: For a K -state model with C clusters and K_c states in cluster $c = 1, \dots, C$, with $\sum_c K_c = K$, there are $\sum_c K_c(K_c - 1)$ parameters to be estimated for the transition matrix ϕ . For a 7-category categorical model, there are 6 parameters for each of the 4 items to be estimated for each state, giving $24K$ parameters for θ . For a 7-category ordinal model, there are 6 parameters for each of the 4 items to be estimated for each cluster, plus $K_c - 1$ location parameters per cluster, giving $\sum_c 4(6 + (K_c - 1)) = 20C + 4K$. Finally, there are an additional $K - 1$ parameters per model to estimate the initial state probabilities.

We consider between 2 and 3 states per cluster and between 1 and 3 clusters per model. These constraints give 18 configurations in total, 9 each for the categorical and ordinal models. Single-state clusters are not considered to ensure the probability of moving out of or into a state is non-zero (otherwise there would be no 'alert' state, sampling times would not be adaptive, and the premise of early detection of unhealthy psychotic events would be lost). However, such configurations are possible and this is not a limitation of the method. Models with more states or clusters are not considered here due to limitations imposed by the number of independent sequences and responses per sequence. The total number of parameters to be estimated for each model is given in Appendix Table I.

C. Model selection and assessment

For convenience, we use the model with the smallest Bayesian Information Criterion (BIC), a likelihood-based statistic used to guide model selection amongst a finite set of non-nested models, to demonstrate the properties of the adaptive sampling scheme outlined below. Whilst it is known that the BIC may be unsuitable for HMMs [24], the interest here is less in selecting an optimal model for future use than in demonstrating the properties of the adaptive sampling scheme.

Model diagnostics such as PP-plots and QQ-plots are not considered here as the model is intended primarily to illustrate the proposed adaptive sampling method, rather than for direct clinical description or prediction purposes. In the results that follow we make the assumption that the underlying model has been robustly tested in terms of its performance and goodness-of-fit

properties. We stress that this is not true in practice and models should not be used for clinical decision-making.

D. Comparison of fixed versus adaptive sampling

After identifying of the best model via minimum BIC and identifying the alert states in each cluster (corresponding to states with a high chance of unhealthy responses), the adaptive sampling scheme was assessed as follows.

First, the fitted model was used to simulate new responses under the proposed adaptive sampling scheme. The scheme chose, for the latest sampling time t_j , a new sampling time $t_{j+1} = t_j + u_j(\tau)$ such that $PP_k(p(t_j), u_j) = \tau$, for $\tau \in \{0, \frac{1}{32}, \frac{1}{16}, \frac{1}{8}, \frac{1}{4}, \frac{1}{2}\}$ and $p(t_j) = p(t_j, y_{1:j})$. Further, $u_j(\tau)$ was modified to ensure responses were not sampled before 3 hours nor after 3 days after the previous sample, i.e., u_j becomes $\max(\min(u_j, 72 \text{ hrs}), 3 \text{ hrs})$. Finally, if after this adjustment $t_j + u_j(\tau)$ was out-of-hours (between 21:00 and 08:59) then $u_j(\tau)$ was increased so that t_{j+1} occurred at 09:00 the next day, thereby mimicking the ClinTouch sampling scheme. Note that $\lim_{\tau \rightarrow 0^+} u_j(\tau) = 0$, i.e., there is non-zero probability of being in or moving to alert state $\bar{k}$ at any time $t_j \geq 0$ and so setting $\tau = 0$ will lead to an immediate request for a new sample. After adjustment for the minimum sampling interval, $\tau = 0$ is equivalent to a fixed sampling scheme with four observations per day at 09:00, 12:00, 15:00, and 18:00.

For each τ , 500 replications were generated, each for 100 days. The primary assessment measure is the sampling frequency for fixed ($\tau = 0$) versus adaptive ($\tau > 0$) sampling schemes. This is calculated for the whole sampling period, and separately for periods when in and not in the alert state. The duration from transition to the alert state until the next sample, the detection delay, is also calculated.

Second, real responses from the two ClinTouch trials are used, where all available sampling times are fixed and known. After calculation of u^j for $\tau \in \{0, \frac{1}{32}, \frac{1}{16}, \frac{1}{8}, \frac{1}{4}, \frac{1}{2}\}$, the soonest available sample after $t_j + u_j$ was taken as the actual sample time. This mimics a scenario where sampling times are pre-specified, but are only used if the $PP_k(p(t_j), u_j) \geq \tau$. Again, $\tau = 0$ is equivalent to sampling at the maximum frequency possible, so will use all available samples. This scheme is applied for each ClinTouch user and the sampling frequency used for fixed ($\tau = 0$) versus adaptive ($\tau > 0$) sampling scheme is compared.

E. Software

All data processing, modelling, and graphs were produced in the statistical computing software R [25]. CTHMMs were fitted using the `msm` package [15]. Analysis code and data are available at <https://github.com/wjchulme/CTHMMs-for-adaptive-sampling>.

IV. RESULTS

In this section we describe the best-fitting model from the configurations considered and describe the performance of the adaptive sampling versus fixed sampling scheme,

Fig. 2. Response distribution for the 2-2-3 model

with respect to average sampling frequency, the alert-state detection delay, and how this changes with different values of the alert threshold τ .

A. Chosen model

The BIC for all 18 models (9 each for categorical and ordinal variants) is provided in Table I. The best model, according to the minimum BIC, was the 2-2-3 ordinal model (3 clusters with 2, 2, and 3 states per cluster respectively). Recall that the states are identified in an unsupervised fashion and require expert examination to determine their meaning. To this end, the response distributions for the 2-2-3 model are provided in Figure 2. The states C1S2, C2S2, and C3S2 were chosen as the alert states as these corresponded either to the states with the highest chance of unhealthy symptoms, or for cluster 3, to the state with intense hallucination symptoms.

B. Adaptive sampling in simulated responses

The 2-2-3 ordinal model was used to simulate new 100-day response sequences and assess the effect of adaptive sampling on the sampling frequency and detection delay. Simulation results are provided in Figures 3 and 4 and tabulated in Appendix Table II.

The distribution of the delay between entering the alert state and taking a new sample is given in Figure 4.

As expected, there is a clear inverse relationship between the alert threshold τ and the sampling frequency per day, and a positive relationship between τ and the alert state delay. For clusters 1 and 3, even for a very small alert state threshold ($\tau = \frac{1}{32}$) the sampling frequency reduced from 4 to around 2.9 and 2.7 samples per day respectively, but remained high (3.9 samples per day) when in the alert state. The cost of this is an increase in the alert-state detection delay. For fixed sampling this delay is 0.12 days (just under

Fig. 3. Sampling frequency for each tau.

Fig. 4. Distribution of the mean detection delay for different values of τ , i.e., the duration from entering the alert state until the next sampling time.

3 hours) but for $\tau = \frac{1}{32}$ this increases to 0.47 and 0.24 days (11 and 6 hours) for C1 and C3, respectively. For cluster 2, though there is little decrease in the delay before sampling during an alert state (an increase of 0.02 days or half an hour), there is only a very small reduction in the overall sampling frequency, even for $\tau = \frac{1}{2}$.

The reduction in the sampling frequency is strongly dependent on the underlying CTHMM and the incoming data. If the modelled transition rates are high then PP will be higher for a given response trajectory, since movement between states is assumed to be more frequent.

Fig. 5. Proportion of available samples used by adaptive sampling. Each line represents data from one ClinTouch user. A small amount of random noise is added to the vertical dimension to show overlapping lines.

C. Adaptive sampling in real responses

While it is not possible to know the underlying states in real responses, we can approximate how a given CTHMM would select sample times under different values of the alert threshold τ . Figure 5 shows how the proportion of used versus available samples reduces as τ increases, for each ClinTouch user.

When $\tau = 0$, all available samples are used which is equivalent to having fixed sampling times. But even at $\tau = \frac{1}{32}$, there are some users who are observed much less frequently than when samples are fixed. This means that for those people the risk of an alert state at the next available sampling time is often predicted to be less than $\frac{1}{32}$ and so the sample is not used. In effect, these are the people who are almost always in non-alert states and are considered low risk. Conversely, there are some people who continue to be observed at maximum frequency even when $\tau = \frac{1}{2}$, because their risk of being in or moving to an alert state is persistently high. The average increase in the detection delay is largest in cluster 1, from 0.3 days when $\tau = \frac{1}{32}$ to from 1.4 days when $\tau = \frac{1}{2}$. In the other clusters, the average increase in the detection delay does not exceed 0.4 days.

V. DISCUSSION

We have presented a model of intensive longitudinal monitoring that uses continuous-time hidden Markov models to capture the underlying state-switching process, which is then used to adaptively select the next sampling time according to the predicted risk of being in a pre-defined alert state.

Unlike passive monitoring systems such as those using wrist-worn accelerometry or photoplethysmography devices,

active monitoring systems frequently demand the user's time and attention, making them more vulnerable to app-fatigue, increasing the risk of temporary or permanent disengagement (non-response or drop-out). In the CareLoop and Actissist trials from which our dataset was derived, 16 of 49 participants (33%) had stopped using the app before the final week of the 12-week observation period. Sub-optimal app tolerability may have contributed to these early drop-outs. On the other hand, too infrequent samples will not provide sufficient information to monitor or model the symptom processes under consideration. By adaptively sampling, it is possible to reduce unnecessary sampling during low risk periods whilst maintaining or increasing the sampling frequency during high risk periods.

We show that, in some cases, our proposed adaptive scheme substantially reduces the sampling frequency compared with a fixed scheme even for very low risk thresholds, leading to substantial sample savings compared with fixed sampling times and removing the need for over 75% of samples in some cases (see Figure 5). We chose 4 samples per day (every 3 hours in a 12-hour window) as our maximum sampling frequency as this is the fixed frequency available in the ClinTouch dataset, but it is possible that increasing the maximum frequency further will still result in sample savings overall under an adaptive scheme. The cost for fewer samples is an increase in the time taken to observe the user during an alert state, where there may be an immediate need to intervene. However, the relatively recent introduction of intensive longitudinal monitoring systems, if implemented successfully, could represent a vast sampling improvement on traditional monitoring for health conditions, where a person may only be seen every couple of months. Set in this context, we believe that the potential benefits gained through increased tolerability, and therefore prolonged engagement, due to reductions in the average sampling frequency are likely to outweigh the harms resulting from a delay in identifying an ongoing alert state.

Our choice of CTHMMs for modelling sequential EMA-type data was made for a number of reasons. First, they handle irregular sampling times which are typical in EMA sequences, for instance due to pseudo-random sampling, user-initiated responses, connectivity issues. This is also necessary if the sampling times are adaptively chosen. Second, they can accommodate multivariate response data, where there is more than one response variable of interest that can be used to predict the state sequence. Third, a suitable CTHMM can predict physiological, psychological, or behavioural states in real-time, which can be used to dynamically decide when the next sample should be taken, and when to intervene if necessary. Finally, they can identify symptom phenotypes and person-specific traits from historical EMA data, so that heterogeneity between sequences can be incorporated into state predictions to improve state prediction for individuals. CTHMMs therefore provide a unified framework for trait and state identification, and for risk-aware sampling and intervention times.

Our adaptive sampling approach can be particularly

useful in settings where people are being monitored for the purpose of delivering the right intervention at the right time. Our use-case is in relation to psychotic relapse risk in schizophrenia, but the principles of adaptive sampling using CTHMMs can be applied more broadly. Further applications in mental health might include prediction of suicidal ideation or self-harm. There are potential applications for home-based, low-frequency monitoring of chronic diseases; samples may ordinarily be taken weekly or monthly but are increased during signs of symptom exacerbation, accelerating detection of disease progression and increasing sample size for subsequent analyses. Another use-case is to sample more frequently during periods where there is a high rate of change or fluctuation of response values, which could improve the accuracy of offline interpolation of symptoms between sampling times. In this case, the response variable would be a transformation of the raw symptom data that captures symptom volatility or variation, rather than the symptoms themselves. More generally, it might be applied in settings such as activity or location monitoring to ensure more frequent samples during active or moving states.

EMA systems are deployed for monitoring symptoms or behaviours in everyday settings and offer improved ecological validity for research in various domains. The rich longitudinal datasets generated by EMA monitoring present many new opportunities for clinical phenotyping. Phenotyping studies using EMA data exist, though often involve observation periods that are too short to reasonably detect or exploit temporal patterns, and so the time dimension is simplified (e.g., clock-time or morning/evening) or else ignored completely [26]–[29]. Where longer sequences of data are available, incorporating temporality through the use of state transition models may lead to more useful phenotypes [30], [31]. However, collecting longer sequences is only possible if the monitoring system is tolerable to the user for a prolonged period. While the ClinTouch dataset used in this paper boasts relatively long observation periods compared with many other EMA phenotyping studies, we recognise that there are limitations to its usefulness in identifying clinical phenotypes. We therefore do not attempt, and indeed advise against, a thorough clinical interpretation of the phenotypes in the 2-2-3 CTHMM identified as the best model.

A. Limitations and future work

This work represents the first step in the application of CTHMMs to model EMA-type for adaptive sampling in real-time, whose potential has not been fully explored here. There are a number of limitations, which themselves may motivate future research.

The models presented here used the method proposed by Smyth[16] to identify phenotypes of people with a diagnosis of schizophrenia assumed to be governed by the same transition and response distributions. However, for people with sufficiently long histories of EMA-acquired symptom data it may be possible to develop personalised models

that are unique to each person. Indeed, such an approach could be deployed in an EMA framework that dynamically updates the model with the incoming data, either with each new sample or in batches, so that sampling times and alerts are informed by a model that is continually learning, reducing dependence of the system on a possibly poorly-suited prior model [32]. Bayesian HMMs are relevant here, particularly early in the observation process where little or no data has been collected as there is no dependence on asymptotic distributions. In practice, however, this is a difficult approach to implement [33]. Each update requires some guarantee of model convergence, which is notoriously difficult for complex Bayesian models in research, let alone in real-time. There may be challenges arising from the label-switching problem (e.g., states diverging from their prior healthy/unhealthy labels over time), causing problems in the interpretation of predicted hidden states [34], [35].

For any models developed with the intention to be deployed in real EMA settings, it is important to assess the goodness-of-fit properties to understand if and where the model may fail to capture the underlying state transition process adequately. In general, however, model selection and validation is a challenge for HMMs. The BIC was used to choose the configuration of the final model, and whilst likelihood-based information criteria are commonly used to select from a set of candidate models for HMMs, the appropriateness of these measures is questionable since the autocorrelation in the responses is not accounted for. See [24] and references therein for some attempts to remedy this in specific modelling scenarios. The chosen model is therefore selected more for expediency than for satisfying any theoretically-justified optimality criteria.

Hidden Markov models typically deal with multivariate data by assuming independence conditionally on the hidden state, as has been done in this paper. Whilst this may be reasonable in certain applications, extensions that permit dependencies between multivariate responses will provide more flexible, and therefore potentially better-fitting, models. Since our cumulative link models assume a latent Normal variable, permitting conditional dependency between ordinal variables is conceptually straight-forward: simply discretise the multivariable Normal response [36]. Unfortunately, the likelihood function is difficult to evaluate [37], particularly when combined with HMMs, and there is no readily-available software to handle this for the applied researcher.

We assumed that missingness was non-informative in our model and illustrative example. However, in remote monitoring contexts, missing data (data not being recorded at an expected time) and indeed the observation process (the rate at which unscheduled data are collected) are likely to provide information in themselves that can be predictive of the underlying state [38], [39].

It would also be of interest to incorporate data from other sources when determining states of interest. This could be achieved in a Bayesian setting by placing prior distributions on the hidden states or cluster parameters. Similarly, other data on the individual patients, such as comorbidity or

demographic information, could also straightforwardly be incorporated into the model.

In this work, alert states were identified by an experienced consultant in psychiatry (author SL). However, validation of the choice of alert states should also be considered. One way to do this would be to use hard outcome data: e.g. relapse events in the example presented in this paper. Alert states should then be states that are highly predictive of relapse events.

The illustration of the methods developed here used a small dataset of 49 ClinTouch users observed for a maximum of 12 weeks, with users asked to answer the full questionnaire two times per day (over four sampling times). However, due to non-response the average number of samples per user was 157.2, which is typically too small for fitting all but the smallest CTHMMs reliably. This exceeds the parameter space for many of the CTHMMs considered in this paper (see Appendix Table I) and so personalised CTHMMs were not explored.

Whilst this work demonstrates the principle of CTHMMs for monitoring EMA data in real-time, the adaptive sampling process was not implemented for real, relying instead on simulations to assess its performance and characteristics. This approach should therefore be trialed and tested in other real datasets and settings before being relied upon for safe monitoring of psychotic symptoms or any other monitoring setting. There are clear implementation challenges for the proposed system, and careful failure checks are required to ensure the system can be implemented safely and effectively. Perhaps the biggest challenge is that of non-response due to a lack of app engagement. This poses a similar practical challenge to medication adherence from the patient side, and a missing data challenge from the analysis side.

VI. CONCLUSION

Maintaining user engagement with potentially inconvenient or stressful technology is a persistent issue in research that demands time and effort on the part of the user. Safely reducing the sampling frequency has the potential to improve tolerability of EMA systems and improve the effectiveness of data collection for real-time monitoring and research. We have proposed an extension to EMA systems such that the sampling frequency is increased if the risk of an alert state is high, as predicted using a continuous-time hidden Markov model. Extensions that consider personalised models, updated in real-time, should be considered in future to improve model reliability and relevance.

REFERENCES

- [1] S. Shiffman, A. A. Stone, and M. R. Hufford, "Ecological momentary assessment," *Annual Review of Clinical Psychology*, vol. 4, pp. 1–32, 2008.
- [2] K. E. Heron and J. M. Smyth, "Ecological momentary interventions: Incorporating mobile technology into psychosocial and health behaviour treatments," *British Journal of Health Psychology*, vol. 15, no. 1, pp. 1–39, 2010.

- [3] D. L. Littlewood, S. D. Kyle, L.-A. Carter, S. Peters, D. Pratt, and P. Gooding, "Short sleep duration and poor sleep quality predict next-day suicidal ideation: An ecological momentary assessment study," *Psychological Medicine*, vol. 49, no. 3, pp. 403–411, Feb. 2019.
- [4] M. Vigo, L. Hassan, W. Vance, C. Jay, A. Brass, and S. Cruickshank, "Britain Breathing: Using the experience sampling method to collect the seasonal allergy symptoms of a country," *Journal of the American Medical Informatics Association*, vol. 25, no. 1, pp. 88–92, Jan. 2018.
- [5] K. L. Druce, J. McBeth, S. N. van der Veer, D. A. Selby, B. Vidgen, K. Georgatzis, B. Hellman, R. Lakshminarayana, A. Chowdhury, D. M. Schultz, C. Sanders, J. C. Sergeant, and W. G. Dixon, "Recruitment and Ongoing Engagement in a UK Smartphone Study Examining the Association Between Weather and Pain: Cohort Study," *JMIR mHealth and uHealth*, vol. 5, no. 11, p. e168, Nov. 2017.
- [6] I. L. MacDonald and W. Zucchini, *Hidden Markov and Other Models for Discrete-valued Time Series*. CRC Press, 1997.
- [7] Z. Yan, V. Subbaraju, D. Chakraborty, A. Misra, and K. Aberer, "Energy-Efficient Continuous Activity Recognition on Mobile Phones: An Activity-Adaptive Approach," in *2012 16th International Symposium on Wearable Computers*, 2012, pp. 17–24.
- [8] A. Krause, M. Ihmig, E. Rankin, D. Leong, Smriti Gupta, D. Siewiorek, A. Smailagic, M. Deisher, and U. Sengupta, "Trading off prediction accuracy and power consumption for context-aware wearable computing," in *Ninth IEEE International Symposium on Wearable Computers (ISWC'05)*, 2005, pp. 20–26.
- [9] J. Paek, J. Kim, and R. Govindan, "Energy-efficient Rate-adaptive GPS-based Positioning for Smartphones," in *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services*, 2010, pp. 299–314.
- [10] K. K. Rachuri, "Smartphones based Social Sensing: Adaptive Sampling, Sensing and Computation Offloading," p. 204.
- [11] K. Vansteelandt, F. Rijmen, G. Pieters, M. Probst, and J. Vanderlinden, "Drive for thinness, affect regulation and physical activity in eating disorders: A daily life study," *Behaviour Research and Therapy*, vol. 45, no. 8, pp. 1717–1734, Aug. 2007.
- [12] B. Hosenfeld, E. H. Bos, K. J. Wardenaar, H. J. Conradi, H. L. J. van der Maas, I. Visser, and P. de Jonge, "Major depressive disorder as a nonlinear dynamic system: Bimodality in the frequency distribution of depressive symptoms over time," *BMC Psychiatry*, vol. 15, no. 1, p. 222, Sep. 2015.
- [13] P. Whelan, M. Machin, S. Lewis, I. Buchan, C. Sanders, E. Applegate, C. Stockton, S. Preston, R. A. Bowen, Z. Ze, C. Roberts, L. Davies, T. Wykes, N. Tarrier, S. Kapur, and J. Ainsworth, "Mobile early detection and connected intervention to coproduce better care in severe mental illness," *Stud Health Technol Inform*, vol. 216, pp. 123–126, 2015.
- [14] I. Visser, "Seven things to remember about hidden Markov models: A tutorial on Markovian models for time series," *Journal of Mathematical Psychology*, vol. 55, no. 6, pp. 403–415, Dec. 2011.
- [15] C. Jackson, "Multi-State Models for Panel Data: The msm Package for R," *Journal of Statistical Software*, vol. 38, no. 1, pp. 1–28, Jan. 2011.
- [16] P. Smyth, "Clustering Sequences with Hidden Markov Models," in *Advances in Neural Information Processing Systems*, 1997, pp. 648–654.
- [17] L. E. Baum, T. Petrie, G. Soules, and N. Weiss, "A Maximization Technique Occurring in the Statistical Analysis of Probabilistic Functions of Markov Chains," *The Annals of Mathematical Statistics*, vol. 41, no. 1, pp. 164–171, 1970.
- [18] L. E. Baum and T. Petrie, "Statistical Inference for Probabilistic Functions of Finite State Markov Chains," *The Annals of Mathematical Statistics*, vol. 37, no. 6, pp. 1554–1563, Dec. 1966.
- [19] A. Bureau, J. P. Hughes, and S. C. Shiboski, "An S-plus implementation of hidden markov models in continuous time," *Journal of Computational and Graphical Statistics*, vol. 9, no. 4, pp. 621–632, Jan. 2000.
- [20] J. E. Palmier-Claus, J. Ainsworth, M. Machin, C. Barrowclough, G. Dunn, E. Barkus, A. Rogers, T. Wykes, S. Kapur, I. Buchan, E. Salter, and S. W. Lewis, "The feasibility and validity of ambulatory self-report of psychotic symptoms using a smartphone software application," *BMC Psychiatry*, vol. 12, no. 1, p. 172, Oct. 2012.
- [21] S. R. Kay, A. Fiszbein, and L. A. Opler, "The Positive and Negative Syndrome Scale (PANSS) for Schizophrenia," *Schizophr Bull*, vol. 13, no. 2, pp. 261–276, Jan. 1987.
- [22] S. Bucci, C. Barrowclough, J. Ainsworth, M. Machin, R. Morris, K. Berry, R. Emsley, S. Lewis, D. Edge, I. Buchan, and G. Haddock, "Actissist: Proof-of-Concept Trial of a Theory-Driven Digital Intervention for Psychosis," *Schizophr Bull*, vol. 44, no. 5, pp. 1070–1080, Aug. 2018.
- [23] A. Agresti, *Categorical data analysis*, 2nd ed. New York: Wiley-Interscience, 2002.
- [24] R. J. Mackay, "Estimating the order of a hidden markov model," *Canadian Journal of Statistics*, vol. 30, no. 4, pp. 573–589, 2002.
- [25] R Core Team, *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria, 2018.
- [26] C. C. Cushing, A. M. Marker, C. M. Bejarano, C. J. Crick, and L. P. Huffhines, "Latent Variable Mixture Modeling of Ecological Momentary Assessment Data: Implications for Screening and Adolescent Mood Profiles," *Journal of Child and Family Studies*, vol. 26, no. 6, pp. 1565–1572, Jun. 2017.
- [27] C. Crayen, M. Eid, T. Lischetzke, D. Courvoisier, and J. Vermunt, "Exploring Dynamics in Mood Regulation Mixture Latent Markov Modeling of Ambulatory Assessment Data," *Psychosomatic Medicine*, vol. 74, no. 4, pp. 366–376, May 2012.
- [28] S. A. Wonderlich, R. D. Crosby, S. G. Engel, J. E. Mitchell, J. Smyth, and R. Miltenberger, "Personality-

Based Clusters in Bulimia Nervosa: Differences in Clinical Variables and Ecological Momentary Assessment,” *Journal of Personality Disorders*, vol. 21, no. 3, pp. 340–357, May 2007.

[29] S. Shiffman, “Conceptualizing Analyses of Ecological Momentary Assessment Data,” *Nicotine & Tobacco Research*, vol. 16, no. Suppl 2, pp. S76–S87, May 2014.

[30] C. Crayen, M. Eid, T. Lischetzke, and J. K. Vermunt, “A Continuous-Time Mixture Latent-State-Trait Markov Model for Experience Sampling Data,” *European Journal of Psychological Assessment*, vol. 33, no. 4, pp. 296–311, Jul. 2017.

[31] T. Asparouhov, E. L. Hamaker, and B. Muthén, “Dynamic Latent Class Analysis,” *Structural Equation Modeling: A Multidisciplinary Journal*, vol. 24, no. 2, pp. 257–269, Mar. 2017.

[32] D. A. Jenkins, M. Sperrin, G. P. Martin, and N. Peek, “Dynamic models to predict health outcomes: Current status and methodological challenges,” *Diagnostic and Prognostic Research*, vol. 2, no. 1, p. 23, Dec. 2018.

[33] S. de Haan-Rietdijk, P. Kuppens, C. S. Bergeman, L. B. Sheeber, N. B. Allen, and E. L. Hamaker, “On the Use of Mixed Markov Models for Intensive Longitudinal Data,” *Multivariate behavioral research*, vol. 52, no. 6, pp. 747–767, 2017.

[34] G. Celeux, “Bayesian Inference for Mixture: The Label Switching Problem,” in *COMPSTAT*, R. Payne and P. Green, Eds. Heidelberg: Physica-Verlag HD, 1998, pp. 227–232.

[35] M. Stephens, “Dealing with label switching in mixture models,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 62, no. 4, pp. 795–809, Nov. 2000.

[36] I. Visser and M. Speekenbrink, “depmixS4: An R Package for Hidden Markov Models,” *Journal of Statistical Software*, vol. 36, no. 1, pp. 1–21, Aug. 2010.

[37] S. Chib and E. Greenberg, “Analysis of Multivariate Probit Models,” *Biometrika*, vol. 85, no. 2, pp. 347–361, 1998.

[38] M. Sperrin, G. P. Martin, R. Sisk, and N. Peek, “Missing data should be handled differently for prediction than for description or causal explanation,” *Journal of Clinical Epidemiology*, vol. 0, no. 0, Jun. 2020.

[39] E. M. Pullenayegum and L. S. Lim, “Longitudinal data subject to irregular observation: A review of methods with a focus on visit processes, assumptions, and study design,” *Statistical Methods in Medical Research*, May 2014.

given the underlying state and that the state sequence is Markovian, it follows that

$$L_i = \sum_{Z_i^1} f_{Z_i^1}(Y_i^1) \Pr(Z_i^1) \times \sum_{Z_i^2} f_{Z_i^2}(Y_i^2) \Pr(Z_i^1, Z_i^2) \times \dots \times \sum_{Z_i^{J_i}} f_{Z_i^{J_i}}(Y_i^{J_i}) \Pr(Z_i^{J_i-1}, Z_i^{J_i}).$$

See [15] for more details on how this likelihood is calculated in the *msm* package, including details on constrained parameter spaces necessary to implement the clustering.

B. Ordinal response distribution

We use Cumulative Link Models (CLMs) [23] to model the distribution of the response variable Y . A CLM assumes an unobserved response variable y^* on $\mathbb{R}$ (note this is different to the latent state-space z) with a known parametric distribution. The observed response y takes value r if y^* lies in $[\theta_{r-1}, \theta_r]$. θ_r is simply a partition of $\mathbb{R}$, with $\theta_r \geq \theta_{r-1}$ and $\theta_0 \equiv -\infty$ and $\theta_7 \equiv \infty$. The advantage of CLMs is that covariate effects can be placed on (a subset of) the parameters of the continuous latent variable rather than on the θ_r , reducing the number of parameters to estimate and ensuring that changes in covariates effect the probabilities of r smoothly. For example, if a standard Normal latent response distribution $N(\mu, 1)$ is assumed, then covariate are modelled to alter the location parameter μ rather than θ_r .

We exploit CLMs as follows. Within each cluster, a latent Normal response distribution is assumed, the thresholds θ_r are fixed, and the location parameter μ is dependent on the current state k . μ is fixed at zero for the first state in each cluster for identifiability. The dispersion parameter σ is fixed at one and is not re-estimated across states. To enforce model identifiability if a contiguous edge set (e.g., $\{1\}$, $\{7\}$, or $\{1, 2\}$) is not observed (leading to edge thresholds being unbounded in the likelihood maximisation procedure), a truncated standard Normal model was used with truncation at ± 4 , with the probability density beyond these limits set to zero. The probability of a random Normal variable being less than -4 or more than 4 is approximately 1 in 30000, which far exceeds the effective sample-size and so estimated parameters are sufficient to approximate the standard Normal model.

C. Tables

APPENDIX

A. Likelihood

The likelihood function for each sequence i is

$$L_i = \Pr(Y_i^1, \dots, Y_i^{J_i}) \\ = \sum \Pr(Y_i^1, \dots, Y_i^{J_i} \mid Z_i^1, \dots, Z_i^{J_i}) \Pr(Z_i^1, \dots, Z_i^{J_i}),$$

summed over all possible state sequences $Z_i^1, \dots, Z_i^{J_i}$. Assuming that the responses are conditionally independent

TABLE I
BIC FOR EACH MODEL

Model	Configuration	transition params	inital state params	response distribution params	Total params	BIC
categorical	2	2	1	48	51	48022
categorical	3	6	2	72	80	47693
categorical	2-2	4	3	96	103	46118
categorical	2-3	8	4	120	132	47667
categorical	3-3	12	5	144	161	50322
categorical	2-2-2	6	5	144	155	49311
categorical	2-2-3	10	6	168	184	51724
categorical	2-3-3	14	7	192	213	54600
categorical	3-3-3	18	8	216	242	52278
ordinal	2	2	1	28	31	47778
ordinal	3	6	2	32	40	46127
ordinal	2-2	4	3	56	63	44739
ordinal	2-3	8	4	60	72	43828
ordinal	3-3	12	5	64	81	43838
ordinal	2-2-2	6	5	84	95	42947
ordinal	2-2-3	10	6	88	104	42555
ordinal	2-3-3	14	7	92	113	43248
ordinal	3-3-3	18	8	96	122	44295

TABLE II
ADAPTIVE SAMPLING PERFORMANCE

τ	Mean frequency	Mean frequency in alert state	Mean frequency not in alert state	Mean alert state detection delay
Cluster 1				
0	4.0 (4.0, 4.0)	4.0 (4.0, 4.0)	4.0 (4.0, 4.0)	0.1 (0.1, 0.4)
1/32	2.9 (2.4, 3.3)	3.9 (3.9, 4.0)	1.3 (1.2, 1.4)	0.5 (0.2, 0.7)
1/16	2.6 (2.1, 3.1)	3.8 (3.7, 3.9)	0.8 (0.8, 0.9)	0.7 (0.3, 1.1)
1/8	2.3 (1.7, 2.8)	3.5 (3.3, 3.7)	0.5 (0.4, 0.5)	1.4 (0.7, 2.1)
1/4	2.0 (1.5, 2.5)	3.1 (2.8, 3.4)	0.4 (0.4, 0.4)	1.5 (0.8, 2.3)
1/2	1.7 (1.2, 2.1)	2.6 (2.3, 3.0)	0.4 (0.4, 0.4)	1.5 (0.8, 2.3)
Cluster 2				
0	4.0 (4.0, 4.0)	4.0 (4.0, 4.0)	4.0 (4.0, 4.1)	0.1 (0.1, 0.4)
1/32	4.0 (4.0, 4.0)	4.0 (4.0, 4.0)	4.0 (3.8, 4.0)	0.2 (0.1, 0.4)
1/16	4.0 (3.9, 4.0)	4.0 (4.0, 4.0)	3.9 (3.5, 4.0)	0.2 (0.1, 0.4)
1/8	3.9 (3.8, 4.0)	4.0 (4.0, 4.0)	3.6 (2.9, 4.0)	0.2 (0.1, 0.5)
1/4	3.8 (3.6, 3.9)	4.0 (3.9, 4.0)	3.0 (2.0, 3.6)	0.3 (0.1, 0.6)
1/2	3.7 (3.4, 3.8)	3.9 (3.8, 4.0)	2.7 (1.4, 3.4)	0.3 (0.1, 0.9)
Cluster 3				
0	4.0 (4.0, 4.0)	4.0 (3.9, 4.1)	4.0 (4.0, 4.0)	0.1 (0.1, 0.4)
1/32	2.7 (2.0, 3.5)	3.9 (3.8, 4.0)	2.2 (1.5, 3.4)	0.2 (0.1, 0.6)
1/16	2.3 (1.7, 3.1)	3.8 (3.6, 4.0)	1.7 (1.1, 2.9)	0.4 (0.1, 0.8)
1/8	2.0 (1.5, 2.9)	3.8 (3.4, 4.0)	1.4 (0.9, 2.5)	0.4 (0.1, 1.1)
1/4	1.9 (1.4, 2.8)	3.6 (3.3, 3.9)	1.2 (0.7, 2.6)	0.5 (0.1, 1.5)
1/2	1.8 (1.2, 2.8)	3.6 (3.2, 3.9)	1.1 (0.6, 2.5)	0.5 (0.1, 1.6)

Units are samples per day