Published in final edited form as:

J Am Stat Assoc. 2013 January 1; 108(504): 1147–1162. doi:10.1080/01621459.2013.779838.

Statistical Learning with Time Series Dependence: An Application to Scoring Sleep in Mice

Blakeley B. McShane¹, Shane T. Jensen², Allan I. Pack³, and Abraham J. Wyner²

¹Kellogg School of Management, Northwestern University

²The Wharton School, University of Pennsylvania

³Center for Sleep and Circadian Neurobiology, University of Pennsylvania

Abstract

We develop methodology which combines statistical learning methods with generalized Markov models, thereby enhancing the former to account for time series dependence. Our methodology can accommodate very general and very long-term time dependence structures in an easily estimable and computationally tractable fashion. We apply our methodology to the scoring of sleep behavior in mice. As currently used methods are expensive, invasive, and labor intensive, there is considerable interest in high-throughput automated systems which would allow many mice to be scored cheaply and quickly. Previous efforts have been able to differentiate sleep from wakefulness, but they are unable to differentiate the rare and important state of REM sleep from non-REM sleep. Key difficulties in detecting REM are that (i) REM is much rarer than non-REM and wakefulness, (ii) REM looks similar to non-REM in terms of the observed covariates, (iii) the data are noisy, and (iv) the data contain strong time dependence structures crucial for differentiating REM from non-REM. Our new approach (i) shows improved differentiation of REM from non-REM sleep and (ii) accurately estimates aggregate quantities of sleep in our application to video-based sleep scoring of mice.

Keywords

sleep; REM; classification; Markov; time series

1. INTRODUCTION: MOTIVATION AND CHALLENGES

1.1 The Science and Scoring of Sleep

Roughly 70 million Americans suffer from chronic sleep loss or sleep disorders, costing the nation approximately \$16 billion per year in medical expenses and a further \$50 billion per year in lost productivity (Patlak 2005). In total, over seventy such disorders afflict about 40 million Americans (Patlak 2005). As knowledge of these problems grows amongst the populace, sleep is becoming an increasingly important field of medical inquiry.

A major area of focus is in determining the genetic basis of sleep behaviors: non-rapid eye movement sleep (NREM), rapid eye movement sleep (REM), and wakefulness (WAKE). Of particular interest is REM sleep which occurs much less frequently than either NREM sleep or WAKE and has an important role in learning and memory (Stickgold et al. 2001). Sleep scientists have shown that various aspects of sleep are heritable including sleep duration

(Partinen et al. 1983; Heath et al. 1990), the timing of sleep (i.e., circadian phase (Heath et al. 1990; Vink et al. 2001)), and the response to sleep deprivation (Franken et al. 2001). While some gene variants affecting these traits have been identified in human studies (He et al. 2009), present knowledge is limited because it is too costly and labor intensive to phenotype a large number of human subjects.

The principal alternative to human studies is to study gene variants in mice. Mice are specially bred to allow the mapping of quantitative trait loci to small regions of the genome (Churchill et al. 2004; Chesler et al. 2008) and their behavior is observed. Unfortunately, studies in mice are also quite limited in size (and consequently scope) because the “gold standard” methodology for studying sleep even in mice is expensive, invasive, and time-consuming. First, electrodes are surgically implanted into the head of a mouse. Then, after waiting ten to fourteen days for the mouse to recover from surgery, electroencephalographic (EEG) and electromyographic (EMG) waves are recorded at 256 Hz for twenty-four hours. The EEG/EMG data are then broken into distinct ten-second blocks (termed “epochs”) and each (i.e., 8,640 epochs per mouse per day) is *manually* scored as NREM, REM, or WAKE by specially-trained technologists.

Not only expensive and laborious, this process can also be quite inconsistent. If the same mouse is scored independently by two different technologists, there is disagreement on approximately 6% of epochs and up to 15% within the important REM stage (Guan et al. 2010; McShane et al. 2012). Moreover, when a given technologist revisits the same data at a later time period, the original scores and new scores fail to match at rates that are only mildly better than those for two different technologists.

Consequently, sleep scientists have sought high-throughput automated systems that would avoid both the surgery and the labor associated with EEG/EMG-based manual scoring. Indeed, sleep scientists have already made initial efforts towards this end. A leading approach, termed the “40-second Rule” (Pack et al. 2007), uses video recordings (or, alternatively, electronic beam splits) to determine whether a mouse is moving or not. Any duration of inactivity lasting forty seconds or more is considered sleep. The 40-second Rule has been validated by comparison to manual scores based on EEG/EMG recordings in both young (Pack et al. 2007) and old (Naidoo et al. 2008) mice of the strain C57BL/6J. An alternative approach (Flores et al. 2007) also relies on mouse movements but instead uses piezoelectric sensors implanted into the floor of the mouse cage to detect movement. The data recorded by such sensors contains patterns that are characteristic of sleep versus wakefulness.

While these methods are able to differentiate sleep from wakefulness, they have no ability to detect the substages of sleep (i.e., NREM sleep and REM sleep). However, there are known physical manifestations differentiating NREM sleep and REM sleep that should induce subtle signals in digital video recordings (for a full review, see Steriade (2005)). For example, while sleep in general is associated with a reduction in muscle tone (atonia), this reduction is far more pronounced in REM sleep as compared to NREM sleep with a near complete absence of muscle tone in REM sleep. Consequently, as the mouse transitions from NREM sleep to REM sleep, there is an onset of near complete atonia leading to a change in the shape of the mouse. In particular, the mouse “flattens out” such that there is a decrease in its aspect ratio (i.e., the ratio of the mouse’s length to its width) and an increase in its area (i.e., the size of the mouse in a two-dimensional video recording). Both of these features are detectable by video. For example, the decrease in aspect ratio as the mouse transition from NREM sleep to REM sleep is visible in Figure 1 which plots the aspect ratio (derived from video recordings) and sleep state (determined by EEG/EMG-based manual scoring) of a mouse over several hours. While this change and similar ones are subtle, they

are detectable and it is these features in combination with the striking temporal dependencies evident in Figure 1 which allow us to achieve our goal, namely the development of a model that can identify NREM sleep versus REM sleep in mice based on digital video recordings.

1.2 Challenges for Statistical Learning Methods

To date, sleep scientists have focused on movement-based measures (e.g., activity versus inactivity) obtained from video cameras, electronic beams, or piezoelectric sensors. However, Figure 1 and similar plots suggest that additional covariates such as aspect ratio are relevant for classifying sleep stages. A natural approach, therefore, would be to (i) cull additional covariates from video data and (ii) employ statistical learning procedures such as logistic regression, AdaBoost (Freund and Schapire 1996), or random forests (Breiman 2001).

While statistical learning procedures have proven extremely successful at minimizing classification error on a wide variety of problems, they are known to have reduced performance in certain more difficult classification situations. For example, AdaBoost can produce “overfit” conditional class probability estimates which tend towards zero or one (Mease et al. 2007; Mease and Wyner 2008), a problem that is particularly acute in noisy environments where the Bayes’ error is high. The high rate of disagreement amongst manual scorers and the similarity of the NREM and REM states in video recordings both suggest that the Bayes’ error is high in our setting.

Another difficulty is that many learning methods were developed for the binary classification setting and do not naturally accommodate a multi-class setting like ours. Even approaches which do (e.g., random forests which does so by using a voting rule over the generated trees) can fail in situations where multiple states have similar observed covariates or when one or more states are rare. In our application, we face both of these concerns: not only are REM and NREM relatively undifferentiated in terms of the observed covariates but also REM occurs in a mere 5% of epochs (compared to about 45% for NREM and 50% for WAKE).

A final difficulty posed by our application—and the one that will be the focus of our efforts—concerns the fact that most statistical learning methods assume independently and identically distributed data. Specifically, the conditional class probability function $\mathbb{P}(Y_t|X_t)$ for outcome Y_t given covariates X_t is assumed to be independent of the rest of the data (Y_{-t} , X_{-t}). As illustrated in Figure 1, our data form a non-*i.i.d.* time series with strong dependencies in both the response (i.e., the sleep states) and the covariates. It should be possible to gain additional discriminatory power by modeling these dependencies.

In summary, our application presents three challenges to standard statistical learning procedures: we require estimates of the conditional state probabilities in a (i) noisy (ii) multi-class setting with rare and undifferentiated states and (iii) time dependencies. We address these challenges by combining conventional classification methods with a very general form of the Markov model. In particular, we introduce methodology that enhances the conditional class probability estimates produced by standard statistical learning procedures so that they are able to take account of a very general class of time dependence structures. Nevertheless, by embedding these general dependence structures in a first order Markov structure, our approach remains computationally feasible, parsimonious, and easily estimable.

The more general approach considered here is critical for our application as the time dependence structures present in sleep data are not suitably addressed by prior techniques. For example, modeling strategies that combine classification methods with a first order

Markov model (Rabiner 1989; Smyth 1994) are ill-suited to sleep data as the first order Markov model implies that bout durations (i.e., the individual state holding times) are (i) geometrically distributed and (ii) do not depend on the previous state. Prior literature has found both of these assumptions untenable (McShane et al. 2010), and, indeed, the geometric distribution provides a poor fit to our data as evidenced by the fact that the black Q-Q lines in Figure 2 exceed the gray null bands for each of the three states. Further, the bout duration distributions of each state depend strongly on the previous state (i.e., they are “transition-dependent”) as the black Q-Q lines in Figure 3 exceed the gray null bands in each plot. In contrast to the first order Markov model, our more general methodology can easily accommodate both of these features.

The empirical duration distributions plotted in Figures 2 and 3 point to a further difficulty posed by sleep data: in addition to their non-geometric shape and their transition-dependence, these distributions have extremely long (if not unbounded) support thus suggesting that a high order Markov model may be necessary to adequately reflect the time dependence structures present in the data. A challenge, then, is to allow for the extension of the basic first order Markov model to higher orders without the concomitant explosion in the number of parameters associated with standard higher order Markov models. Building on related techniques such as non-stationary Markov models (Sin and Kim 1995; Vaseghi 1995; Djuric and Chun 2002), semi-Markov models (Janssen and Limnios 1999), variable duration Markov models (Ferguson 1980; Levinson 1986), and variable length Markov models (Buhlmann and Wyner 1999), our methodology overcomes this challenge by using a principled and pragmatic parsimonious parametric approach that naturally allows for the long state durations present in sleep data.

The remainder of our paper is organized as follows. We introduce our methodology in Section 2. In Section 3, we evaluate its performance relative to simpler alternatives in an extensive simulation study. Section 4 demonstrates that exploiting time dependencies substantially adds to our ability to detect the signal extant in our video recordings of mice, particularly given the high level of noise inherent in the problem; our procedure is able to identify sleep and its substages with reasonable accuracy when validated by comparison to EEG/EMG assessments in C57BL/6J male mice. Section 5 discusses the principal benefits of our new sleep scoring methodology, including its ability (i) to detect the rare and subtle REM state which is of special interest to sleep researchers and (ii) to analyze the large numbers of mice required for genetic analysis.

2. METHODS

2.1 Data Structure

In the standard classification setting, we have a response variable Y which takes on a finite set of values $Y \in S = \{1, \dots, k\}$ depending on a set of covariates $\mathbf{X}$. Typically, the goal is to estimate the conditional class probability function $\mathbb{P}(Y = y | \mathbf{X} = \mathbf{x})$ based on *i.i.d.* training data $(y_i, \mathbf{x}_i)_{i=1}^N$ drawn from the joint distribution $\mathbb{P}(Y, \mathbf{X})$. Our sleep data contains strong sequential dependencies, forcing us to move beyond this *i.i.d.* assumption.

In particular, our dataset consists of (i) a vector of length T of response variables $\mathbf{Y}_{1:T} = (Y_1, \dots, Y_T)$ with $Y_t \in S$ and (ii) a $T \times p$ matrix of covariates $\mathbf{X}_{1:T}$ whose rows are the video-based covariates $\mathbf{X}_t$ for time t . The goal is to either (i) estimate the joint distribution $\mathbb{P}(\mathbf{Y}_{1:T}, \mathbf{X}_{1:T})$ or (ii) predict a new response sequence $\mathbf{Y}^*$ given a new covariate matrix $\mathbf{X}^*$ based on the conditional distribution $\mathbb{P}(\mathbf{Y}_{1:T} | \mathbf{X}_{1:T})$.

2.2 Markov Models for Local Time Dependence

A standard approach for sequential data is the first order Markov model (Rabiner 1989) shown in Figure 4. This model decomposes the joint distribution $\mathbb{P}(\mathbf{Y}_{1:T}, \mathbf{X}_{1:T})$ into three components:

1. An initialization distribution, $\boldsymbol{\pi} = (\pi_1, \dots, \pi_k)$ where $\pi_i = \mathbb{P}(Y_1 = i)$.
2. A time-homogeneous transition matrix $\mathbf{A} = (a_{ij})$ where $a_{ij} = \mathbb{P}(Y_{t+1} = j | Y_t = i)$.
3. A set of multivariate time-homogeneous covariate emission distributions $\boldsymbol{\mu} = (\mu_1(\mathbf{x}), \dots, \mu_k(\mathbf{x}))$ where $\mu_i(\mathbf{x}) = \mathbb{P}(\mathbf{X}_t = \mathbf{x} | Y_t = i)$.

A sequence of data is generated from this first order Markov model by (i) drawing an initial state $Y_1 = i$ from $\boldsymbol{\pi}$, (ii) drawing $\mathbf{X}_1 = \mathbf{x}$ from μ_i , (iii) transitioning to state $Y_2 = j$ according to the transition probabilities given by the i^{th} row of $\mathbf{A}$, and (iv) repeating steps (ii)–(iii) *mutatis mutandis* until time T is reached (for full details, see our supplementary materials).

Collecting our model parameters into $\Delta = (\boldsymbol{\pi}, \mathbf{A}, \boldsymbol{\mu})$, the likelihood is given by

$$L(\mathbf{Y}_{1:T}, \mathbf{X}_{1:T} | \Delta) = \mathbb{P}(Y_1) \mathbb{P}(\mathbf{X}_1 | Y_1) \cdot \mathbb{P}(Y_2 | Y_1) \mathbb{P}(\mathbf{X}_2 | Y_2) \cdots \mathbb{P}(Y_T | Y_{T-1}) \mathbb{P}(\mathbf{X}_T | Y_T) = \pi_{Y_1} \left[\prod_{t=2}^T a_{Y_{t-1} Y_t} \right] \left[\prod_{t=1}^T \mu_{Y_t}(\mathbf{X}_t) \right]. \quad (1)$$

Given observed $\mathbf{Y}_{1:T}$, one typically estimates the values of $\mathbf{A} = (a_{ij})$ by the empirical rate of transitions from state i to state j . To estimate $\boldsymbol{\pi}$, one usually uses either (i) the empirical frequency of each of the k classes or the (ii) stationary distribution of the estimated transition matrix¹. Finally, one can estimate each $\mu_i(\mathbf{x})$ independently using the $\mathbf{X}_t$ for which $Y_t = i$.

The overall purpose of this model is the scoring of “out-of-sample” sleep states. That is, given a covariate sequence $\mathbf{X}^* = (\mathbf{X}_{T+1}, \dots, \mathbf{X}_{T+T^*})$ and estimated parameters $\hat{\Delta}$, we want to calculate the conditional class probabilities for $\mathbf{Y}^* = (Y_{T+1}, \dots, Y_{T+T^*})$. The probabilities, denoted $\gamma_t(i) \equiv \mathbb{P}(Y_t = i | \mathbf{X}^*, \hat{\Delta})$, can be calculated using the forward-backward algorithm (Rabiner 1989). Our actual classification $\hat{Y}_t$ for of the unobserved time periods is the state with maximum probability, $\hat{Y}_t = \arg\max_i \gamma_t(i)$; classification by the modal state $\hat{Y}_t$ is optimal in terms of minimizing the out-of-sample classification error. While we also calculate the most likely path $\hat{\mathbf{Y}}^* = \arg\max_{\mathbf{Y}^*} \mathbb{P}(\mathbf{Y}^* | \mathbf{X}^*, \hat{\Delta})$ using the Viterbi algorithm, we do not use it for classification purposes as $\hat{\mathbf{Y}}^*$ is suboptimal under classification error loss (Rabiner 1989).

In the context of our sleep application, the first order Markov model suffers from two serious limitations. First, the estimation of the k multivariate probabilities $\mu_i(\mathbf{x})$ is a difficult task, particularly since we have a large number of covariates and there are rare states such as REM. Second, long-term dependencies in the Y_t cannot be captured due to the first order Markov property: any relationship between Y_t and Y_{t+k} must be “communicated” via $Y_{t+1}, \dots, Y_{t+k-1}$. The first order Markov assumption also implies the holding times in each state are geometrically distributed, an assumption at odds with both data and common sense since sleep bouts cannot be memoryless (see Figure 2).

Our proposed methodology addresses both of these limitations. In Section 2.3, we address the estimation of $\mu_i(\mathbf{x})$ by estimating the model in a *discriminative* fashion. In Sections 2.4–

¹Estimating $\boldsymbol{\pi}$ may or may not be necessary for out-of-sample prediction. When the out-of-sample sequence “continues on” from the in-sample sequence as in Figure 4, the initialization distribution for Y_{T+1} is simply the row of the transition probability matrix $\mathbf{A}$ corresponding to the observed Y_T and thus no estimate of $\boldsymbol{\pi}$ is required. When the out-of-sample sequence is an entirely new sequence, an estimate of $\boldsymbol{\pi}$ is necessary.

2.5, we introduce longer-term dependence structures and more general holding time distributions directly into the Y_t themselves.

2.3 Discriminative Markov Models

The estimation of the k multivariate covariate emission distributions $\mu_i(\mathbf{x})$ is difficult, particularly when the number of covariates p is large, the distributional form of μ_i is unknown, or some of the states i are rare. However, it is potentially much easier to address the inverse problem of classifying categorical Y_t based on high-dimensional $\mathbf{X}_t$. This is the setting of various classification methods such as logistic regression and random forests. In particular, Bayes' theorem links classification methods to the covariate emission distributions μ_i as

$$\mu_i(\mathbf{x}) = \mathbb{P}(\mathbf{X}_t = \mathbf{x} | Y_t = i) = \frac{\overbrace{\mathbb{P}(Y_t = i | \mathbf{X}_t = \mathbf{x})}^{\text{Classification Methods}} \cdot \overbrace{\mathbb{P}(\mathbf{X}_t = \mathbf{x})}^{\text{Constant w.r.t. } Y_t}}{\underbrace{\mathbb{P}(Y_t = i)}_{\text{Marginal Probabilities}}} \propto \frac{f_i(\mathbf{X}_t)}{p_i}$$

where $f_i(\mathbf{X}_t) = \mathbb{P}(Y_t = i | \mathbf{X}_t)$ and $p_i = \mathbb{P}(Y_t = i)$. We can thus rewrite the likelihood given in Equation 1 as

$$L(\mathbf{Y}_{1:T}, \mathbf{X}_{1:T} | \Delta) \propto \pi_{Y_1} \left[\prod_{t=2}^T a_{Y_{t-1}, Y_t} \right] \left[\prod_{t=1}^T \frac{f_{Y_t}(\mathbf{X}_t)}{p_{Y_t}} \right] = L(\mathbf{Y}_{1:T}, \mathbf{X}_{1:T} | \Theta)$$

where $\Theta = (\pi, \mathbf{A}, \mathbf{f}, \mathbf{p})$ with $\mathbf{f} = (f_1(x), \dots, f_k(x))$ and $\mathbf{p} = (p_1, \dots, p_k)$.

We have thus transformed the difficult problem of estimating k multivariate probability distributions μ into the more straightforward problem of estimating of (i) a conditional class probability vector $\mathbf{f}$ and (ii) a marginal probability vector $\mathbf{p}$. For estimating $\mathbf{f}$, any classification method which gives conditional class probability estimates will suffice, though, depending on the data, some may be more appropriate than others; in this paper, we focus on using either logistic regression or random forests. Probabilities $\mathbf{p}$ are estimated by either (i) the empirical frequency of each of the k states or (ii) the stationary distribution corresponding to the estimate of $\mathbf{A}$. The forward-backward algorithms require only slight modifications to calculate $\gamma_t(i) = \mathbb{P}(Y_t = i | \mathbf{X}_t, \Theta)$ under our new parameterization (see Appendix A for details).

In summary, rather than estimating the Markov model *generatively*, we are estimating the model *discriminatively* (Smyth 1994). The discriminative Markov model has many advantages including allowing us (i) to avoid the estimation of k multivariate probability emission distributions, (ii) to enhance *any* standard classification methodology to accommodate sequential data, and (iii) to save computational costs as compared to alternative sequential data methods (e.g., conditional random fields (Lafferty et al. 2001)).

Nonetheless, our discriminative Markov model faces the limitation that only local time dependence is modeled. Any relationship between Y_t and Y_{t+k} must be communicated via $Y_{t+1}, Y_{t+2}, \dots, Y_{t+k-1}$ and holding times in each state must be geometrically distributed. We extend our discriminative Markov approach to accommodate more general and longer-range dependencies in the next two sections.

2.4 Generalized Markov Models

The first order Markov model (denoted IMM) presented in the previous sections has geometrically distributed duration times: the probability of staying in state i for τ epochs conditional on having arrived in state i is $\mathbb{P}_i(\tau) = a_{i,i}^{\tau-1}(1 - a_{i,i})$. This distribution, and its memoryless property, provides a poor fit to our sleep data (see Figure 2). While it is conceptually straightforward to generalize the model to accommodate more general duration distributions, efficient computational algorithms (i.e., forward-backward and Viterbi) are only available for first order Markov models. We overcome this limitation by embedding a *generalized* Markov model (denoted GMM) with non-geometric durations into a IMM structure, thereby retaining the use of the efficient forward-backward and Viterbi algorithms.

A GMM is parameterized by $\Delta = (\pi, \mathbf{A}, \mu, \delta)$ in the generative case and $\Theta = (\pi, \mathbf{A}, \mathbf{f}, \mathbf{p}, \delta)$ in the discriminative case. All parameters remain the same as in the IMM case but (i) each $a_{i,i} = 0$ and (ii) holding times within state i are governed not by $a_{i,i}$ but by an arbitrary duration distribution $\delta_i(\tau) = \mathbb{P}_i(\tau)$ which gives the probability of spending τ periods in state i conditional on having arrived in state i . By δ , we denote the collection $(\delta_1(\tau), \dots, \delta_k(\tau))$ of these distributions.

A sequence of data is generated from a GMM by (i) drawing an initial state $Y_1 = i$ from π , (ii) drawing a duration τ_1 from δ_i and setting $Y_1 = \dots = Y_{\tau_1} = i$, (iii) drawing

$\mathbf{X}_1, \dots, \mathbf{X}_{\tau_1} \stackrel{\text{iid}}{\sim} \mu_i$, (iv) transitioning to state $Y_{\tau_1+1} = j \neq i$ according to the transition probabilities given by the i^{th} row of $\mathbf{A}$, and (v) repeating steps (ii)–(iv) *mutatis mutandis* until time T is reached (for full details, see our supplementary materials). Hence, the likelihood is given by

$$\begin{aligned} L(\mathbf{Y}_{1:T}, \mathbf{X}_{1:T} | \Delta) &= [\mathbb{P}(Y_1) \mathbb{P}_{Y_1}(\tau_1)] \left[\prod_{t|Y_t \neq Y_{t-1}} \mathbb{P}(Y_t | Y_{t-1}) \mathbb{P}_{Y_t}(\tau_t) \right] \left[\prod_{t=1}^T \mathbb{P}(\mathbf{X}_t | Y_t) \right] = [\pi_{Y_1} \delta_{Y_1}(\tau_1)] \left[\prod_{t|Y_t \neq Y_{t-1}} a_{Y_t|Y_{t-1}} \delta_{Y_t}(\tau_t) \right] \left[\prod_{t=1}^T \mu_{Y_t}(\mathbf{X}_t) \right] \propto [\pi \prod_{t=1}^T \mu_{Y_t}(\mathbf{X}_t)] \quad (2) \\ &= L(\mathbf{Y}_{1:T}, \mathbf{X}_{1:T} | \Theta). \end{aligned}$$

In order to embed the GMM in a IMM, we must model the $\delta_i(\tau)$ as

$$\delta_i(\tau) = d_i(\tau) \cdot \mathbf{I}(\tau \leq M_i) + (1 - q_i) \cdot s_i^{\tau - M_i - 1} \cdot (1 - s_i) \cdot \mathbf{I}(\tau > M_i) \quad (3)$$

which is a mixture of (i) a distribution $d_i(\tau)$ which gives the probability mass for $\tau = 1, 2, \dots, M_i$ and (ii) a shifted geometric distribution which gives the probability mass for values of τ greater than M_i where M_i separates the “head” and “tail” of the distribution. More specifically, the head of the distribution has total mass given by the mixing proportion

$q_i = \sum_{\tau=1}^{M_i} d_i(\tau)$ and this mass is spread over arbitrary shape $d_i(\tau)$ and arbitrary length M_i ; on the other hand, the tail has total mass $1 - q_i$ and this mass is of geometric shape and unbounded length.

We write the full set of parameters for our GMM as $\Theta = (\pi, \mathbf{A}, \mathbf{f}, \mathbf{p}, \mathbf{d}, \mathbf{s}, \mathbf{q})$ replacing δ by the three equivalent parameters (i) $\mathbf{d} = (d_1(\tau), \dots, d_k(\tau))$ where each d_i is a vector of length M_i giving the probability of remaining in state i conditional on having arrived there for $\tau = 1, \dots, M_i$ periods, (ii) $\mathbf{s} = (s_1, \dots, s_k)$ where each s_i is the geometric distribution parameter for state i , and (iii) $\mathbf{q} = (q_1, \dots, q_k)$ where each $q_i = \sum_{\tau=1}^{M_i} d_i(\tau)$ is the mixing proportion between the head and tail of the duration distribution for state i .

We can embed our GMM into a IMM through a set of augmented variables. Given an observed sequence $\mathbf{Y}_{1:T} = (Y_1, \dots, Y_T)$, we construct $R_t = \arg\max_{\tau} \{Y_t = Y_{t+1} = \dots = Y_{t+\tau-1} Y_{t+\tau}\}$ and $Z_t = \min(R_t, M_{Y_t} + 1)$. R_t gives how much longer the sequence remains in the current state and Z_t caps R_t at the maximal size. We denote our augmented variable as $Y'_t = (Y_t, Z_t)$ and our augmented state space as $S' = \{(i, n) | i \in S, n = 1, \dots, M_i + 1\}$ (i.e., all possible values of the couplet Y'_t).

As a simple example, consider the observed sequence $\mathbf{Y}_{1:T} = \{a, a, a, a, b, b, b, b, c, c, a\}$ on $S = \{a, b, c\}$ with $M_a = 2, M_b = M_c = 3$. Our process converts $\mathbf{Y}_{1:T}$ to the augmented sequence

$$\mathbf{Y}'_{1:T} = \{(a, 3), (a, 3), (a, 2), (a, 1), (b, 4), (b, 4), (b, 3), (b, 2), (b, 1), (c, 2), (c, 1), (a, 1)\}.$$

Our augmented state space S' requires an augmented transition probability matrix $\mathbf{A}'$ which we construct as follows (see Figure 5 for a simple example based on the above):

1. *Transitions away from states:* For $Z_t = 1$, there are three cases:
 - a. Transitions from $Y'_t = (i, 1)$ to $Y'_{t+1} = (i, n)$ have probability zero for $n = 1, \dots, M_i$.
 - b. Transitions from $Y'_t = (i, 1)$ to $Y'_{t+1} = (j, n)$ have probability $a_{ij} \cdot d_i(n)$ for j and $n = 1, \dots, M_j$.
 - c. Transitions from $Y'_t = (i, 1)$ to $Y'_{t+1} = (j, M_j + 1)$ have probability $a_{ij} \cdot (1 - q_j)$ for j .
2. *Transitions within “head” duration of states:* For $Z_t = 2, \dots, M_i$, the transition from $Y'_t = (i, Z_t)$ to $Y'_{t+1} = (i, Z_t - 1)$ has probability one; all other transitions have probability zero.
3. *Transitions within “tail” duration of states:* For $Z_t = M_i + 1$, the self-transition from $Y'_t = (i, M_i + 1)$ to $Y'_{t+1} = (i, M_i + 1)$ has probability s_i and the transition from $Y'_t = (i, M_i + 1)$ to $Y'_{t+1} = (i, M_i)$ has probability $(1 - s_i)$; all other transitions have probability zero.

We can also easily augment the remaining parameters for our new state space S' . Marginal probabilities $\mathbf{p}'$ are obtained from the stationary distribution of $\mathbf{A}'$. To obtain the conditional class probabilities $\mathbf{f}'$ and the initialization distribution $\boldsymbol{\pi}'$, let us first define $Y(i)$ and $Z(i)$ as the first and second elements of the state couplet $i = (Y_i, Z_i) \in S'$. Then, the $\mathbf{f}'$ are obtained by

setting $f'_i(\mathbf{x}) = \mathbb{P}(Y'_t = i | \mathbf{X}_t = \mathbf{x}) = f_{Y(i)}(\mathbf{x}) \cdot p'_i / \sum_{j | Y(j) = Y(i)} p'_j$. Similarly, each $\pi'_i = \mathbb{P}(Y'_1 = i)$ is given by $\pi'_i = \pi_{Y(i)} \cdot d_{Y(i)}(Z(i))$ if $Z(i) = 1, \dots, M_i$ or $\pi'_i = \pi_{Y(i)} \cdot (1 - q_{Y(i)})$ if $Z(i) = M_i + 1$. Thus, the f'_i and π'_i are simply reweighted versions of the original $f_{Y(i)}$ and $\pi_{Y(i)}$.

In summary, we have embedded a GMM on state space S parameterized by $\boldsymbol{\Theta} = (\boldsymbol{\pi}, \mathbf{A}, \mathbf{f}, \mathbf{p}, \mathbf{d}, \mathbf{s}, \mathbf{q})$ into a IMM on state space S' parameterized by $\boldsymbol{\Theta}' = (\boldsymbol{\pi}', \mathbf{A}', \mathbf{f}', \mathbf{p}')$. Because this is a true embedding, the likelihood $L(\mathbf{Y}_{1:T} | \mathbf{X}_{1:T}; \boldsymbol{\Theta})$ given in Equation 2 is equivalent to the likelihood under the new parameterization, $L(\mathbf{Y}'_{1:T}, \mathbf{X}_{1:T} | \boldsymbol{\Theta}')$.

Estimation of the parameters $\boldsymbol{\Theta}'$ is straightforward and we begin with the transition probability matrix $\mathbf{A}'$. While the MLE of $\mathbf{A}'$ is given by the empirical frequencies of the Y'_t transitions, this can be an inefficient estimate since it ignores the known structure in $\mathbf{A}'$ (i.e.,

entries are known to be zero, one, or products of the original $\mathbf{A}$ and the parameters $\mathbf{d}$, $\mathbf{s}$, and $\mathbf{q}$). Instead, we prefer (i) to estimate the transition probability a_{ij} from state i to state j using empirical frequencies and (ii) to parameterize the $d_i(\tau)$ and estimate the parameters d_i , s_i , and q_i using the observed duration times for state i . We can then combine estimates (i) and (ii) as detailed above to form our estimate of $\mathbf{A}'$.

The $\mathbf{p}'$ are estimated by using the stationary distribution of the estimate of $\mathbf{A}'$. The augmented conditional class probabilities $\mathbf{f}'$ are estimated by (i) estimating $\mathbf{f}$ as in the 1MM case (i.e., by using standard classification procedures such as random forests or logistic regression) and then (ii) reweighting these estimates as above (i.e.,

$\hat{f}'_i(\mathbf{x}) = \hat{f}_{Y(i)}(\mathbf{x}) \cdot \hat{p}'_i / \sum_{j|Y(j)=Y(i)} \hat{p}'_j$). Similarly, the initialization parameters π are estimated by (i) estimating π as in the 1MM case and then (ii) reweighting these estimates as above

(i.e., $\hat{\pi}'_i = \hat{\pi}_{Y(i)} \cdot \hat{d}_{Y(i)}(Z(i))$ if $Z(i) = 1, \dots, M_i$ or $\hat{\pi}'_i = \hat{\pi}_{Y(i)} \cdot (1 - \hat{d}_{Y(i)})$ if $Z(i) = M_i + 1$).

Given these estimates of a GMM embedded into a 1MM, we can apply the forward-backward and Viterbi algorithms for the discriminative 1MM just as in Section 2.3

2.5 Transition-Dependent Generalized Markov Model

Although the GMM presented in section 2.4 is much more general than a 1MM, certain properties of the GMM are not realistic in our application. Specifically, the duration distributions in a GMM are “unconditional”: the duration distribution for state i does not depend on the prior state j . However, we have observed in our sleep data that the duration of, for example, WAKE bouts which follow from NREM are longer than WAKE bouts which follow from REM (see Figure 3).

We thus introduce a transition-dependent generalized Markov model (TDGMM) where the duration distributions of the current state depend on the previous state. In this new model, δ contains $k \cdot (k - 1)$ duration distributions $\delta_{ij}(\tau)$ instead of the k duration distributions $\delta_i(\tau)$ of the GMM; each $\delta_{ij}(\tau) = \mathbb{P}_{ij}(\tau)$ is the probability of spending τ periods in state j conditional on having arrived in state j from state i . The data generation process and likelihood for the TDGMM are identical to those of the GMM with the appropriate δ_{ij} used in place of δ_j .

We model the δ_{ij} using the specification of Equation 3 and, hence, our duration distribution parameters are again given by $(\mathbf{d}, \mathbf{s}, \mathbf{q})$. In contrast to the GMM where each of these parameters was of length k (i.e., one for each δ_j), in the TDGMM each of these parameters is of length $k \cdot (k - 1)$ (i.e., one for each δ_{ij}).

As with the GMM, we can embed the new TDGMM within the first order Markov structure by augmenting the state space. Assuming we have observed $\mathbf{Y} = (Y_1, \dots, Y_T)$, we define $P_t = Y_{\sigma_t}$ where $\sigma_t = \operatorname{argmin}_s (Y_s = Y_{s+1} = Y_{s+2} = \dots = Y_t)$ (i.e., the state the sequence was in prior to the current state). We then construct R_t as we did in for the GMM (i.e., $R_t = \operatorname{argmax}_{\tau} \{Y_t = Y_{t+1} = \dots = Y_{t+\tau-1} = Y_{t+\tau}\}$) and again cap it with $Z_t = \min(R_t, M_{P_t} + 1)$. Finally, we set the triplet $Y'_t = (P_t, Y_t, Z_t)$ thus collecting the previous state, the current state, and the capped length of time remaining in the current state. The augmented state space is $S' = \{(i, j, n) | i \in S, j \in S, i \neq j, n = 1, \dots, M_{ij} + 1\}$ (i.e., all possible values of the triplet Y'_t).

Consider again the example $\mathbf{Y}_{1:T} = \{a, a, a, a, b, b, b, b, c, c, a\}$ on $S = \{a, b, c\}$ with $M_{b,a} = M_{a,b} = M_{c,b} = 3$, $M_{c,a} = M_{a,c} = M_{b,c} = 2$. Supposing $Y_0 = c$, our process converts $\mathbf{Y}_{1:T}$ into the augmented sequence $\mathbf{Y}' = \{(c, a, 3), (c, a, 3), (c, a, 2), (c, a, 1), (a, b, 4), (a, b, 4), (a, b, 3), (a, b, 2), (a, b, 1), (b, c, 2), (b, c, 1), (c, a, 1)\}$.

As before, augmented state space S' requires an augmented transition probability matrix $\mathbf{A}'$ which can be constructed as follows (see Figure 6 for a simple example based on the above):

1. *Transitions away from states:* For $Z_t = 1$, there are three cases:
 - a. Transitions from $Y'_t = (i, j, 1)$ to $Y'_{t+1} = (i', j', n)$ have probability zero for $n = 1, \dots, M_{i',j'}$ and $i' = j$.
 - b. Transitions from $Y'_t = (i, j, 1)$ to $Y'_{t+1} = (i', j', n)$ have probability $a_{i',j'} \cdot d_{i',j'}(n)$ for $n = 1, \dots, M_{i',j'}$, $i' = j$, and $j' = j$.
 - c. Transitions from $Y'_t = (i, j, 1)$ to $Y'_{t+1} = (i', j', M_{i',j'} + 1)$ have probability $a_{i',j'} \cdot (1 - q_{i',j'})$ for $i' = j$ and $j' = j$.
2. *Transitions within “head” duration of states:* For $Z_t = 2, \dots, M_{i,j}$, the transition from $Y'_t = (i, j, Z_t)$ to $Y'_{t+1} = (i, j, Z_t - 1)$ has probability one; all other transitions have probability zero.
3. *Transitions within “tail” duration of states:* For $Z_t = M_{i,j} + 1$, the self-transition from $Y'_t = (i, j, M_{i,j} + 1)$ to $Y'_{t+1} = (i, j, M_{i,j} + 1)$ has probability $s_{i,j}$ and the transition from $Y'_t = (i, j, M_{i,j} + 1)$ to $Y'_{t+1} = (i, j, M_{i,j})$ has probability $(1 - s_{i,j})$; all other transitions have probability zero.

We have thus embedded a TDGMM on state space S parameterized by $\Theta = (\pi, \mathbf{A}, \mathbf{f}, \mathbf{p}, \mathbf{d}, \mathbf{s}, \mathbf{q})$ into a IMM on state space S' parameterized by $\Theta' = (\pi', \mathbf{A}', \mathbf{f}', \mathbf{p}')$. Due to the similarity between the GMM and the TDGMM, we can again use the estimation strategy outlined Section 2.4 (the only difference is that the duration distribution parameters $(\mathbf{d}, \mathbf{s}, \mathbf{q})$ are estimated based on the observed duration times for state j conditional on entering state j from state i rather than unconditionally and we use the above procedure to construct estimates of $\mathbf{A}'$; these differences propagate in the obvious manner into the estimates of $\mathbf{p}', \mathbf{f}'$, and π').

3. SIMULATION EVALUATION OF MODELS

We evaluate the performance of several models presented in Section 2 via simulation where the true model is a TDGMM. Our simulation state space is $S = \{a, b, c\}$ and the initialization distribution is the uniform distribution $\pi = (1/3, 1/3, 1/3)$. We set the covariate emission distributions μ equal to $X_t \sim N(\mu Y_t, \sigma^2)$ where $\mu_a = 0$, $\mu_b = 1$, and $\mu_c = 2$ and σ is set to thirteen different values ranging from 0.01 to 3. The transition probability distributions are given by

$$\mathbf{A} = \begin{pmatrix} 0 & 1/2 & 1/2 \\ 3/4 & 0 & 1/4 \\ 1/3 & 2/3 & 0 \end{pmatrix}.$$

The duration distributions $\delta_{i,j}$ are as in Equation 3 with $M_{i,j}$ set to $M = 10$ and $d_{i,j}$, the “head” components of the duration distributions, set to the Beta Negative Binomial distribution with finite support,

$$d(\tau | \alpha, \beta, r) = \frac{1}{c(\alpha, \beta, r, M)} \frac{\Gamma(\alpha + \beta) \Gamma(\alpha + r) \Gamma(\tau + r - 1) \Gamma(\tau + \beta - 1)}{\Gamma(r) \Gamma(\alpha) \Gamma(\beta) \Gamma(\tau) \Gamma(\tau + r + \alpha + \beta - 1)}, \tau = 1, \dots, M$$

where $c(\alpha, \beta, r, M)$ is a normalizing constant ensuring that each $\delta_{i,j}$ sums to one (i.e., $d_{i,j}$ sums to $q_{i,j}$). The parameters (α, β, r) of the “head” components and (q, s) of the “tail” components of the duration distributions $\delta_{i,j}$ were set to

State	α	β	r	q	s
$a \rightarrow b$	0.00	1.00	442413	1.00	0.00
$a \rightarrow c$	0.00	1.65	0.45	1.00	0.00
$b \rightarrow a$	1808	148.41	33.12	1.00	0.00
$b \rightarrow c$	0.00	7.39	7.39	0.50	0.69
$c \rightarrow a$	0.00	22026	0.61	0.62	0.90
$c \rightarrow b$	0.00	1.00	442413	0.50	0.90

The parameters settings were chosen to provide a diversity of shapes as illustrated in Figure 7.

Our study uses three different training set sizes ($T = 100$, $T = 1,000$, and $T = 10,000$). The test set size is always fixed at $T^* = 200$ and “continues” from the training data as in Figure 4. For each value of σ and T , results are averaged over 100 simulated datasets. The competing models we consider are 1. multinomial logistic regression (MLR), 2. MLR enhanced by a 1MM (MLR+1MM), 3. MLR enhanced by a GMM (MLR+GMM), and 4. MLR enhanced by a TDGMM (MLR+TDGMM).

Performance is evaluated in three ways. First, we examine the classification error of each of the four methods as well as that of the Bayes’ Rule which classifies based on the modal true marginal probability $\gamma_T(i) = \mathbb{P}(Y_i | \mathbf{X}^*, \Delta)$. Second, since the Bayes’ Rule gives the optimal classifications which minimize the error rate and since it is free of the noise inherent in the Y_i , we also examine the error rate of each of the four methods relative to the classes predicted by the Bayes’ Rule. Finally, we examine the RMSE of the marginal probability estimates $\gamma_T(i) \equiv \mathbb{P}(Y = i | \mathbf{X}^*, \Theta)$ from the four models².

We present our simulation results in Figure 8. In terms of classification error, the correctly-specified TDGMM model does best with a moderate to large amount of in-sample data. The incorrectly-specified models do not seem to improve as T is increased from 1,000 to 10,000 whereas the correctly-specified TDGMM does. However, even for $T = 10,000$, the TDGMM does not match the Bayes’ Rule at high noise levels; this suggests that massive amounts of training data may be required in very high noise settings. The GMM model performs poorly because it estimates a single duration distribution for each state when the reality is a mixture of transition-specific duration distributions. It is interesting that the 1MM outperforms the more general GMM here, even with large amounts of training data.

In terms of the RMSE of the estimated probabilities, the TDGMM gives the best results in large training samples and appears to be converging on the true probabilities as the training set size grows. However, the simpler 1MM model seems to provide better estimates in small sample sizes or with very high noise levels. Again, the incorrectly-specified models do not seem to improve as T is increased from 1,000 to 10,000 whereas the correctly-specified TDGMM does—an encouraging result for our approach.

²We also examined other proper scoring rules such as log loss and exponential loss (Savage 1971; Buja et al. 2005). All yielded results which were qualitatively similar to those presented for squared error loss.

4. APPLICATION TO SLEEP SCORING

4.1 Data Description and Summary Statistics

Our sleep data come from eight mice of the strain C57BL/6J. Each mouse was manually scored for sleep by the invasive and laborious method described in Section 1.1. That is, the mice were surgically implanted with EEG/EMG electrodes and allowed a ten to fourteen day post surgery recovery and habituation period. Next, EEG/EMG signals were recorded for twenty-four hours, digitized at 256 Hz, and broken up into ten-second epochs (8,640 in total per mouse) for manual scoring. Each of the $8 \times 8,640$ epochs was manually scored as NREM, REM, or WAKE by two independent scorers.

The independent scorers disagreed on about $\approx 6\%$ of the ten-second epochs and disagreement rates were substantially higher among those epochs where the sleep stage was REM or transitional. For each epoch where the two scorers disagreed, we used an independent third scorer to break the tie and to thereby determine the “true” state.

The goal of our methodology is to predict sleep scores using video data which closely match the “gold standard” EEG/EMG-based manual scores. Our digital video consists of 100 frames per epoch and two example frames, one typical of wakefulness and one typical of sleep, are shown in Figure 9. Tracking software was used to fit an elliptical approximation to the mouse in each frame.

Using this ellipse, the tracking software calculates six continuous covariates for each epoch: the intra-epoch mean and standard deviation of the velocity, aspect ratio, and size (area) of the mouse. We also have a single binary covariate which indicates whether or not a light in the cage was turned on in each epoch (lights were on from 7AM – 7PM). Thus, we have seven covariates X_t with which to predict the sleep stage Y_t for each of 8,640 epochs.

To give an overall sense of the data, we provide summary statistics for the three states in Table 1. The mice spend about half of the day awake and the other half in the two stages of sleep. As mentioned previously, REM sleep is relatively rare, accounting for fewer than 5% of epochs. Furthermore, REM has many fewer bouts and a lower average bout duration.

As bout durations in a given state depend on the prior state (see Figure 3), we provide in Table 2 the same summary statistics for the three sleep states conditional on the previous state. We see that almost all bouts of REM come from NREM. This is not surprising since transitions from WAKE to REM are an extremely rare physiological occurrence³. For NREM, the bouts seem to be longer when entered from REM compared to when entered from WAKE. For WAKE, the bouts tend to be longer when entered from NREM rather than entered from REM. These characteristics of the data motivated our consideration of the transition-dependent generalized Markov model.

Another important feature of the duration distributions is heavy skewness, with some bouts lasting many times the length of the mean durations presented in Tables 1 and 2. As mentioned previously, sleep bout durations are distributed according to a “head and tail” distribution (McShane et al. 2010) which features (i) most of the probability mass in the first few epochs and (ii) the remaining mass in a long right tail. We found that the beta negative binomial distribution with geometric tail used in Section 3 could accommodate these

³Such transitions are indicative of sleep disorders, incorrectly scored epochs, or DREM. DREM is a direct transition from WAKE to REM that occasionally occurs in some mice. Such episodes occur almost exclusively during the lights on period and are the result of brief awakenings interrupting a sustained period of REM sleep (Fujiki et al. 2009).

features and provide a good fit to the conditional bout duration distributions of our TDGMM. Details of this estimation are provided in our supplementary materials.

4.2 Internal Model Evaluation

Our principal goal is to build a model based on video data which best matches EEG/EMG-based manual scores on a epoch-by-epoch basis. We thus use a cross-validation approach to evaluating our various models: we take the full set of epochs for one mouse as our training data and test each model on the full set of epochs from the other seven mice. We then repeat this procedure using each mouse as the “in-sample” mouse and then averaging our results over all 483,840 out-of-sample epochs.

Before trying to discriminate REM from NREM, we first consider the simpler two-state problem of scoring SLEEP versus WAKE. We assign the “true” two-state score for a given epoch to be SLEEP if the manual scorers scored the epoch as NREM or REM and assign it to be WAKE otherwise (breaking ties, as above, by using an independent, third scorer when necessary). Our performance metric is the out-of-sample classification error rate of our predictions as compared to the manual scores.

For this two-state task, we compare the performance of six models: 1. multinomial logistic regression (MLR), 2. MLR enhanced by a 1MM (MLR+1MM), 3. MLR enhanced by a GMM (MLR+GMM), 4. random forests (RF), 5. random forests enhanced by a 1MM (RF+1MM), and 6. random forests enhanced by a GMM (RF+GMM). For data consisting only of two states, the TDGMM reduces to a GMM and so is excluded from this evaluation.

We also compare our model choices to the “40-second rule” common in sleep science (Pack et al. 2007). The 40-second rule judges a mouse “inactive” in a given epoch if the mean intra-epoch velocity is less than 3 pixels/s and scores an epoch as SLEEP when there are four or more consecutive inactive epochs (i.e., forty or more seconds of inactivity); otherwise, it scores the epoch as WAKE.

The second column of Table 3 gives the classification error rate for each of these methods as well as the “gold-standard” rate of disagreement between manual scorers (which can be viewed as the minimal achievable error rate). We see that our 1MM and GMM enhancements of MLR and RF achieve error rates lower than the “40-second rule” that is currently used by sleep scientists. The best overall error rates are achieved by the GMMs which have the most flexible model for the duration distributions. The MLR+GMM model provides a 15% raw improvement over MLR (which does not model time-series dependence); this equates to a 31% improvement relative to the minimum achievable error rate. It also provides a 19% raw improvement over the 40-second rule or a 36% improvement relative to the minimum achievable. We also note that ROC curves (not shown) demonstrate a uniform improvement of the GMMs relative to the 1MMs and base MLR and RF models.

We now consider the more difficult (and important) task of classifying each epoch into NREM versus REM versus WAKE. As discriminating between SLEEP and WAKE is comparably easy, the primary consideration is in classifying NREM versus REM. As mentioned, REM is a relatively rare state that is biologically important. Furthermore, it is the focus of many inquiries in sleep science. Consequently, in typical applications, the cost of misclassifying a NREM epoch as REM is much lower than the cost of misclassifying a REM epoch as NREM. Indeed, the rarity and importance of REM suggests the desirability of a relatively high false positive rate and a comparably low false negative rate.

For this three-state task, we compare the performance of 1. MLR, 2. MLR+1MM, 3. MLR enhanced by a TDGMM (MLR+TDGMM), 4. RF, 5. RF+1MM, and RF enhanced by a TDGMM (RF+TDGMM). We consider a TDGMM but not a GMM for this three-state task because the unconditional duration distributions of the GMM provided a poor fit to the data (see our supplementary materials). We also exclude the “40-second rule” from this comparison since this rule is incapable of differentiating NREM from REM sleep. In fact, there is no existing procedure in the sleep literature which classifies NREM versus REM sleep.

Classification error rates for the three-state task are presented in the third column of Table 3. The three-state task is complicated by the rarity of the REM state and the similarity between REM and NREM on the observed covariates, thus making overall error rates much higher than in two-state task (i.e., column two versus column three). In addition to the overall error rates, we also include rate of REM prediction (column four) as well as the false positive and false negative rate for REM (columns five and six respectively) since this state is especially important to sleep researchers.

Table 3 gives some unsurprising results: the relatively rare REM state is difficult to classify correctly with REM false negative rates being especially large. It is a challenge to discover any method with power to detect REM sleep. That is, there is an inherent trade-off between (i) obtaining a low REM false negative rate accompanied by a higher overall error rate and a higher REM false positive rate or (ii) obtaining a lower overall error rate and a lower REM false positive rate while having a high REM false negative rate.

The simple MLR and RF have the lowest overall error rates and REM false positive rates simply because they choose alternative (ii) and more or less ignore the REM state. Epochs are very rarely classified as REM by these simple methods, resulting in extremely high REM false negative rates. Due to the asymmetry in the cost of errors discussed above, sleep researchers would have a strong preference against such classifications.

In contrast, the TDGMM enhancements of MLR and RF are able to achieve a better balance for REM prediction: they have a much lower REM false negative rate while remaining competitive on the overall error rate and REM false positive rates. By accounting for the time-dependence in the data, our TDGMM procedure is able to capture a greater proportion of the subtle REM signal.

We also examine ROC curves for the three states in Figure 10. MLR, MLR+1MM, and MLR+TDGMM all perform strongly at discriminating WAKE from the other states though MLR+TDGMM uniformly dominates. On the other hand, there are substantial differences for REM and NREM. MLR+TDGMM uniformly dominates the other two models on the rare and important REM state. For NREM, MLR and MLR+TDGMM perform similarly and both uniformly dominate MLR+1MM.

4.3 External Model Evaluation

In this section, we compare our top methodology from the previous section (MLR+TDGMM) to several other classification techniques. Specifically, we consider AdaBoost (Freund and Schapire 1996), LogitBoost (Friedman et al. 2000), and Bagged Trees (Breiman 1996). We also examined two implementations of conditional random fields—the linear MALLT system (McCallum 2003) and the TreeCRF approach (Dietterich et al. 2004)—both of which are popular in the computer science literature⁴.

Classification error rates for these methods applied to the three-state NREM vs. REM vs. WAKE task are presented in Table 4. In general, the alternative classification methods

perform similarly to the simple MLR and RF models in Table 3 in that the important REM state is generally ignored. The result is lower overall error rates and REM false positive rates but extremely high REM false negative rates. Our MLR+TDGMM methodology achieves a much lower REM false negative rate with only modest increases in overall error rate and REM false positive rates.

4.4 Aggregate Measures of Sleep

Sections 4.2–4.3 focused on our ability to match the EEG/EMG-based manual scores on an epoch-by-epoch basis. However, sleep scientists are also interested in *aggregate* measures of sleep behavior such as the total number of minutes spent in each sleep state at various times of the day. While matching manual scores on an epoch-by-epoch basis is a sufficient condition for estimating these aggregate measures, it is not a necessary one and accurate estimates of aggregate quantities can be obtained from models which are less precise on an epoch-by-epoch basis.

A principal quantity of interest to sleep scientists is N_i^j , the average number of minutes mice spend in state i during each two-hour block j of the day. This value can be calculated from EEG/EMG-based manual scores as

$$N_i^j = \frac{c}{8} \sum_{m=1}^8 \sum_{t \in \text{Block}_j} \mathbf{I}(Y_{m,t}=i)$$

where $Y_{m,t}$ is the manual score for mouse m at epoch t and $c = 10/60$ converts time measured in ten-second epochs to time measured in minutes.

Our methodology based on video data can be used to estimate N_i^j by summing our estimated probabilities $\gamma_t(i) = \mathbb{P}(Y_t = i | \mathbf{X}, \hat{\Theta})$ over all epochs t in block j . We average these probabilities over all combinations of one training (in-sample) mouse $\times$ seven testing (out-of-sample) mice

$$\hat{N}_i^j = \frac{c}{8} \frac{1}{7} \sum_{m_{in}=1}^8 \sum_{m_{out} \neq m_{in}} \sum_{t \in \text{Block}_j} \hat{\gamma}_{m_{in}, m_{out}, t}(i)$$

where $\hat{\gamma}_{m_{in}, m_{out}, t}(i) = \mathbb{P}(Y_{m_{out}, t} = i | \mathbf{X}_{m_{out}}, \hat{\Theta}_{m_{in}})$ is the probability that mouse m_{out} is in state i at epoch t given covariates from mouse m_{out} and parameters estimated on mouse m_{in} .

This simple estimation is used for our two-state task of classifying SLEEP vs. WAKE. For the more difficult three-state task of classifying NREM vs. REM vs. WAKE, we introduce additional parameters w_m which partition the time spent REM and NREM for each mouse m . Specifically, we have

⁴We used default parameter settings for MALLiT. TreeCRF requires discrete data and so our covariates were binned using different numbers of quantiles (5, 10, 25 and 100). TreeCRF also requires that the number of leaves be specified (we tried 8, 16, 32, 64, 128 and 256). The TreeCRF results given are the best-case over all these parameter settings (b and l are the corresponding number of bins and leaves).

$$\begin{aligned}
\hat{N}_{REM}^j &= \frac{c}{8} \frac{1}{7} \sum_{m_{in}=1}^8 \sum_{m_{out} \neq m_{in}} \sum_{t \in \text{Block}_j} \hat{w}_{m_{in}} \\
&\quad \cdot \hat{\gamma}_{m_{in}, m_{out}, t}(REM) \hat{N}_{NREM}^j \\
&= \frac{c}{8} \frac{1}{7} \sum_{m_{in}=1}^8 \sum_{m_{out} \neq m_{in}} \sum_{t \in \text{Block}_j} \left[\hat{\gamma}_{m_{in}, m_{out}, t}(NREM) + (1 - \hat{w}_{m_{in}}) \cdot \hat{\gamma}_{m_{in}, m_{out}, t}(REM) \right].
\end{aligned}$$

We let $\hat{N}_{WAKE}^j$ remain as before. Each of the w_m parameters lies in $[0, 1]$ and allows the model to “give” a portion of the time spent in REM to NREM for mouse m . We estimate each w_m on only the in-sample mouse so that our evaluation remains entirely out of sample.

In Figure 11, we compare the N_i^j from EEG/EMG-based manual scoring to the estimated $\hat{N}_i^j$ from our video-based methodology. We examine the match for both the two-state task of classifying SLEEP vs. WAKE (left panel) and the three-state task of classifying NREM vs. REM vs. WAKE (right panel). For both tasks, our estimated number of minutes spent in each state provides an excellent match to the number based on manual scores.

5. DISCUSSION

Our proposed methodology enhances standard statistical learning methods such as logistic regression and random forests to account for time dependence. This approach is very general in that it can be used to enhance any statistical learning method which produces conditional class probability estimates. Our embedding methodology for Markov models makes a very rich and general class of dependence structures computationally feasible and easily estimable; though we have focused on first order Markov models, generalized Markov models, and transition-dependent generalized Markov models, our strategy can also accommodate higher order Markov models, variable length Markov models, and other related models.

When applied to the classification of sleep states in mice based on video data, our procedure shows increased ability to discriminate NREM from REM sleep relative to alternative classification methods. This is a considerable advance over current approaches used by sleep scientists which can only discriminate SLEEP from WAKE (Pack et al. 2007; Flores et al. 2007). Furthermore, our procedure also allows for very accurate assessment of the total amount of NREM sleep, REM sleep, and wakefulness.

This is an improvement over the current approach to the assessment of this phenotype in mice which requires (i) surgery following anesthesia for the insertion of small screws in the scalp for assessment of the electroencephalogram; (ii) insertion of a wire in the neck muscles for assessment of the electromyogram; (iii) recovery from surgery for several days; (iv) assessment of stages of sleep or wakefulness by direct visual inspection and scoring of 8,640 ten-second epochs per mouse per day. This is expensive, invasive, and time-consuming thereby limiting the number of mice that can be studied. Thus, our approach has particular benefit when a large number of mice need to be studied for reasons of, for example, statistical power.

High-throughput screening has several immediate applications. First, in studies in which mRNA changes or protein changes with sleep and wakefulness are being assessed, our approach is much more cost-effective for estimating sleep states. Second, there are several new mouse resources for studying the genetic basis of behaviors such as sleep including (i)

the panel of Collaborative Cross mice which consists of 300 lines of mice derived from eight founder strains (Churchill et al. 2004; Chesler et al. 2008; Collaborative Cross Consortium 2012); (ii) the panel of Diversity Outbred mice which is based on random mating of Collaborative Cross mice to provide genetic heterogeneity (Svenson et al. 2012); and (iii) other large panels of mice with individual genes knockout (e.g., Skarnes et al. (2011)). Our methods could be applied to these large numbers of genetically heterogeneous mice in order to assess the genetic basis of various sleep behaviors. Third, high-throughput screening has advantages in screening libraries of compounds to detect novel drugs that alter sleep and wake. Indeed, high-throughput screening based on video data has already been employed in zebra-fish to identify such compounds (Rihel et al. 2009) and the methodology developed here for mice facilitates this approach to mammalian species.

In many of aforementioned applications, the focus is on obtaining accurate estimates of aggregate measures of sleep as opposed to epoch-by-epoch scores. In addition to metrics such as the amount of time spent in each state at various points throughout the day, scientists are also interested in estimating quantities such as the number and typical length of bouts of each state in order to understand how they vary across the day (Saper et al. 2005), particularly for mice that differ in terms of (i) genetic background (Mochizuki et al. 2011; Sehgal and Mignot 2011; Naidoo et al. 2012), (ii) age (Naidoo et al. 2008; Hasan et al. 2012), (iii) level of sleep deprivation (Franken et al. 1999), and (iv) exposure to various compounds (Vienne et al. 2010). Our proposed methodology shows the ability to provide accurate assessments of these aggregate quantities in a high-throughput setting (McShane et al. 2012).

In addition to these immediate applications, sleep science suggests several potential future applications and enhancements of our proposed methodology. First, a three-dimensional recording of the mouse using high-definition video cameras as well as side-angle cameras in addition to overhead cameras would likely improve the assessment of mouse behavior; in particular, such a system could likely track breathing as well as the small twitches that occur in REM sleep, thus yielding improved differentiation of NREM sleep versus REM sleep on an epoch-by-epoch basis. Second, augmenting or substituting video data with piezoelectric data (i.e., sensors located in the floor of the mouse cage that measure pressure changes produced by movement of the mouse) could yield improved scores as there are (i) highly variable signals during wakefulness as the mouse moves around the cage and (ii) subtle signals during sleep and its substages which reflect breathing patterns (Friedman et al. 2004). Augmenting our model with covariates from both of these sources could potentially allow for the wholesale replacement of manual scoring as such covariates may be sufficiently precise to allow model-based error rates to approach the inter-human level of $\approx 5\%$.

In conclusion, video-based analyses show the ability to distinguish REM from NREM sleep in mice. While future elaborations of this technological approach could lead to further improvements in these estimates, high-throughput phenotyping of sleep in mice is feasible and will facilitate genetic studies and the investigation of chemical libraries to determine compounds that affect NREM, REM, and WAKE.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Financial support: This research was supported in part by NIH grants T32 HL07713, P01 AG17628, and MH081491.

APPENDIX A. DISCRIMINATIVE MARKOV ALGORITHMS

In Section 2.3, we introduced a discriminative approach to estimating a Markov model. Here we modify the forward-backward and Viterbi algorithms to accommodate our discriminative parameter set $\hat{\Theta} = (\hat{\pi}, \hat{\mathbf{A}}, \hat{\mathbf{f}}, \hat{\mathbf{p}})$ rather than generative parameter set $\hat{\Delta} = (\hat{\pi}, \hat{\mathbf{A}}, \hat{\mu})$. We can use the results of the modified forward-backward algorithms (Algorithms 1–2) to estimate the conditional class probabilities $\gamma_t(i) = \mathbb{P}(Y_t = i | \mathbf{X}, \hat{\Theta})$ for each $i \in S$. In particular,

$$\hat{\gamma}_t(i) = \frac{\alpha_t(i)\beta_t(i)}{\mathbb{P}(\mathbf{X}_{1:T} = \mathbf{X} | \hat{\Theta})} = \frac{\alpha_t(i)\beta_t(i)}{\sum_{j=1}^k \alpha_t(j)\beta_t(j)}.$$

As with conventional generative Markov models, our maximum likelihood estimate of Y_t is $\hat{Y}_t = \operatorname{argmax}_{i \in S} \gamma_t(i)$ and our maximum likelihood estimate of the sequence $\mathbf{Y} = (Y_1, \dots, Y_T)$ is given by the Viterbi Algorithm (Algorithm 3).

Algorithm 1

The Discriminative Forward Algorithm.

Begin with estimates $\hat{\Theta} = (\hat{\pi}, \hat{\mathbf{A}}, \hat{\mathbf{f}}, \hat{\mathbf{p}})$ and define $\alpha_t(i) = \mathbb{P}(\mathbf{X}_{1:t}, Y_t = i | \hat{\Theta})$.

- 1 Initialization: $\alpha_1(i) = \pi_i \frac{f_i(x_1)}{p_i}$ for $i = 1, \dots, k$.
 - 2 Induction: $\alpha_{t+1}(j) = \left[\sum_{i=1}^k \alpha_t(i) \cdot \hat{a}_{i,j} \right] \frac{f_j(x_{t+1})}{p_j}$ for $t = 1, \dots, T-1$.
 - 3 Termination: $\mathbb{P}(\mathbf{X}_{1:T} = \mathbf{X} | \hat{\Theta}) = \sum_{i=1}^k \alpha_T(i)$.
-

Algorithm 2

The Discriminative Backward Algorithm.

Begin with estimates $\hat{\Theta} = (\hat{\pi}, \hat{\mathbf{A}}, \hat{\mathbf{f}}, \hat{\mathbf{p}})$ and define $\beta_t(i) = \mathbb{P}(\mathbf{X}_{t+1:T} | Y_t = i, \hat{\Theta})$.

- 1 Initialization: $\beta_T(i) = 1$ for $i = 1, \dots, k$.
 - 2 Induction: $\beta_t(i) = \sum_{j=1}^k \hat{a}_{i,j} \beta_{t+1}(j) \frac{f_j(x_{t+1})}{p_j}$ for $t = T-1, \dots, 1$.
-

Algorithm 3

The Discriminative Viterbi Algorithm.

Begin with estimates $\hat{\Theta} = (\hat{\pi}, \hat{\mathbf{A}}, \hat{\mathbf{f}}, \hat{\mathbf{p}})$ and define $\delta_t(i) = \max_{Y_{1:t-1}} \mathbb{P}(Y_{1:t-1}, Y_t = i, \mathbf{X}_{1:t} | \hat{\Theta})$.

- 1 Initialization: $\delta_1(i) = \pi_i \frac{f_i(x_1)}{p_i}$ and $\psi_1(i) = 0$ for $i = 1, \dots, k$
- 2 Recursion:

$$\delta_t(j) = \max_{i=1,\dots,k} [\delta_{t-1}(i) \hat{a}_{i,j}] \frac{f_j(x_t)}{p_j} \quad t = 2, \dots, T; \quad j = 1, \dots, k \quad \psi_t(j) = \operatorname{argmax}_{i=1,\dots,k} [\delta_{t-1}(i) \hat{a}_{i,j}] \quad t = 2, \dots, T; \quad j = 1, \dots, k$$

- 3 Termination: $P^* = \max_{i=1,\dots,k} \delta_T(i)$ and $Y_T^* = \operatorname{argmax}_{i=1,\dots,k} \delta_T(i)$
- 4 Path (state sequence) backtracking: $Y_t^* = \psi_{t+1}(Y_{t+1}^*)$ for $t = T-1, \dots, 1$.

REFERENCES

- Breiman L. Bagging Predictors. *Machine Learning*. 1996; 24:123–140.
- Breiman L. Random Forests. *Machine Learning*. 2001; 45:5–32.
- Buhlmann P, Wyner AJ. Variable Length Markov Chains. *Annals of Statistics*. 1999; 27(2):480–513.
- Buja A, Stuetzle W, Shen Y. Loss Functions for Binary Probability Estimation and Classification: Structure and Applications. 2005 <http://www-stat.wharton.upenn.edu/buja/PAPERS/paper-proper-scoring.pdf> (Under review),.
- Chesler EJ, Miller DR, Branstetter LR, Galloway LD, Jackson BL, et al. The Collaborative Cross at Oak Ridge National Laboratory: developing a powerful resource for systems genetics. *Mammalian Genome*. 2008; 19(6):382–389. [PubMed: 18716833]
- Churchill GA, Airey DC, Allayee H, Angel JM, Attie AD, et al. The Collaborative Cross, a community resource for the genetic analysis of complex traits. *Nature Genetics*. 2004; 36(11):1133–1137. [PubMed: 15514660]
- Collaborative Cross Consortium. The Genome Architecture of the Collaborative Cross Mouse Genetic Reference Population. *Genetics*. 2012; 190(2):389–491. [PubMed: 22345608]
- Dietterich, TG.; Ashenfelder, A.; Bulatov, Y. Training Conditional Random Fields via Gradient Tree Boosting; *Proceedings of the 21st International Conference on Machine Learning (ICML 2004)*; 2004.
- Djuric PM, Chun J-H. An MCMC Sampling Approach to Estimation of Nonstational Hidden Markov Models. *IEEE Transactions on Signal Processing*. 2002; 50(5):1113–1123.
- Ferguson, JD. Variable Duration Models for Speech; *Proceedings of Symposium on the Application of Hidden Markov Models to Text and Speech*; 1980. p. 143-179.
- Flores A, Flores J, Deshpande H, Picazo J, Xie X, Franken P, Heller H, Grahn D, O'Hare B. Pattern recognition of sleep in rodents using piezoelectric signals generated by gross body movements. *IEEE Transactions on Biomedical Engineering*. 2007; 54(2):225–233. [PubMed: 17278579]
- Franken P, Chollet D, Tafti M. The homeostatic regulation of sleep need is under genetic control. *Journal of Neuroscience*. 2001; 21(8):2610–2621. [PubMed: 11306614]
- Franken P, Malafosse A, Tafti M. Genetic determinants of sleep regulation in inbred mice. *Sleep*. 1999; 22(2):155–169. [PubMed: 10201060]
- Freund, Y.; Schapire, R. Experiments with a New Boosting Algorithm; *Machine Learning: Proceedings of the Thirteenth International Conference*; 1996. p. 148-156.
- Friedman J, Hastie T, Tibshirani R. Additive logistic regression: a statistical view of boosting. *Annals of Statistics*. 2000; 28:337–374.
- Friedman L, Haines A, Klann K, Gallagher L, Salibra L, Han F, Strohl KP. Ventilatory behavior during sleep among A/J and C57BL/6J mouse strains. *Journal of Applied Physiology*. 2004; 97(5):1787–1795. [PubMed: 15475556]
- Fujiki N, Cheng T, Yoshino F, Nishino S. Specificity of direct transition from wake to REM sleep in orexin/ataxin-3 transgenic narcoleptic mice. *Experimental Neurology*. 2009; 217(1):46–54. [PubMed: 19416673]
- Guan C, Ye C, Yand X, Gao J. A review of current large-scale mouse knockout efforts. *Genesis*. 2010; 48(2):73–85. [PubMed: 20095055]

- Hasan S, Dauvilliers Y, Mongrain V, Franken P, Tafti M. Age-related changes in sleep in inbred mice are genotype dependent. *Neurobiology of Aging*. 2012; 33(1):195.e13–195.e26. [PubMed: 20619936]
- He Y, Jones CR, Fujiki N, Xu Y, Guo BJL, Holder J, Rossner MJ, Nishino S, Fu YH. The transcriptional repressor DEC2 regulates sleep length in mammals. *Science*. 2009; 325(5942):866–870. [PubMed: 19679812]
- Heath AC, Kendler KS, Eaves LJ, Martin NG. Evidence for genetic influences on sleep disturbance and sleep pattern in twins. *Sleep*. 1990; 13(4):318–335. [PubMed: 2267475]
- Janssen, J.; Limnios, N., editors. *Semi-Markov Models and Applications*. Dordrecht, The Netherlands: Kluwer Academic; 1999.
- Lafferty J, McCallum A, Pereira F. Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. *Proceedings of the 18th International Conference on Machine Learning (ICML 2001)*. 1999
- Levinson SE. Continuously variable duration hidden Markov models for automatic speech recognition. *Computer Speech and Language*. 1986; 1(1):29–45.
- McCallum, A. Efficiently inducing features of conditional random fields; *Proceedings of the 19th Conference on Uncertainty in Artificial Intelligence*; 2003. p. 403–410.
- McShane BB, Galante RJ, Biber MP, Jensen ST, Wyner AJ, Pack AI. Assessing REM Sleep in Mice Using Video Data. *Sleep*. 2012; 35(3):433–442. [PubMed: 22379250]
- McShane BB, Galante RJ, Jensen ST, Naidoo N, Pack AI, Wyner A. Characterization of the Bout Durations of Sleep and Wakefulness. *Journal of Neuroscience Methods*. 2010; 193(2):321–333. [PubMed: 20817037]
- Mease D, Wyner A. Evidence Contrary to the Statistical View of Boosting. *Journal of Machine Learning Research*. 2008; 9:131–156.
- Mease D, Wyner A, Buja A. Boosted Classification Trees and Class Probability/Quantile Estimation. *Journal of Machine Learning Research*. 2007; 8:409–439.
- Mochizuki T, Arrigoni E, Marcus JN, Clark EL, Yamamoto M, Honer M, Borroni E, Lowell BB, Elmquist JK, Scammell TE. Orexin receptor 2 expression in the posterior hypothalamus rescues sleepiness in narcoleptic mice. *Proceedings of the National Academy of Sciences*. 2011; 108(11):4471–4476.
- Naidoo N, Ferber M, Galante RJ, McShane BB, Hu JH, Zimmerman J, Maislin G, Cater J, Wyner AJ, Worley P, Pack AI. Role of Homer Proteins in the Maintenance of Sleep-Wake States. *PLoS One*. 2012; 7(4):e35174. [PubMed: 22532843]
- Naidoo N, Ferber M, Master M, Zhu Y, Pack AI. Aging impairs the unfolded protein response to sleep deprivation and leads to proapoptotic signaling. *Journal of Neuroscience*. 2008; 28(26):6539–6548. [PubMed: 18579727]
- Pack AI, Galante RJ, Maislin G, Cater J, Metaxas D, Lu S, Zhang L, Smith RV, Kay T, Lian J, Svenson K, Peters LL. Novel Method for High-Throughput Phenotyping of Sleep in Mice. *Physiological Genomics*. 2007; 28:232–238. [PubMed: 16985007]
- Partinen M, Kaprio J, Koskenvuo M, Putkonen P, Langinvainio H. Genetic and environmental determination of human sleep. *Sleep*. 1983; 6(3):179–185. [PubMed: 6684786]
- Patlak, M. *Your Guide to Healthy Sleep*. Bethesda, MD: U.S. Department of Health and Human Services; 2005. (*NIH Publication No. 06-5271*)
- Rabiner LR. A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. *Proceedings of the IEEE*. 1989; 77:257–286.
- Rihel J, Prober DA, Arvanites A, Lam K, Zimmerman S, et al. Zebrafish behavioral profiling links drugs to biological targets and rest/wake regulation. *Science*. 2009; 327(5963):348–351. [PubMed: 20075256]
- Saper CB, Scammell TE, Lu J. Hypothalamic regulation of sleep and circadian rhythms. *Nature*. 2005; 437(7063):1257–1263. [PubMed: 16251950]
- Savage LJ. Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association*. 1971; 66(336):783–801.
- Sehgal A, Mignot E. Genetics of sleep and sleep disorders. *Cell*. 2011; 146(2):194–207. [PubMed: 21784243]

- Sin B, Kim JH. Nonstationary hidden Markov model. *Signal Processing*. 1995; 46:31–46.
- Skarnes WC, Rosen B, West AP, Koutsourakis M, Bushell W, et al. A conditional knockout resource for the genome-wide study of mouse gene function. *Nature*. 2011; 474(7351):337–342. [PubMed: 21677750]
- Smyth P. Markov monitoring with unknown states. *IEEE Journal of Selected Areas in Communications, Special Issue on Intelligent Signal Processing for Communications*. 1994; 12:1600–1612.
- Steriade, BM. Brain electrical activity and sensory processing during waking and sleep states. In: Kryger, MH.; Roth, T.; Dement, WC., editors. *Principles and Practice of Sleep Medicine*. 4 edn. Saunders: Elsevier; 2005. p. 101-119.
- Stickgold R, Hobson J, Fosse R, Fosse M. Sleep, Learning, and Dreams: Off-line Memory Reprocessing. *Science*. 2001; 294:1052–1057. [PubMed: 11691983]
- Svenson KL, Gatti DM, Valdar W, Welsh CE, Cheng R, et al. High-Resolution Genetic Mapping Using the Mouse Diversity Outbred Population. *Genetics*. 2012; 190(2):437–448. [PubMed: 22345611]
- Vaseghi SV. State duration modeling in hidden Markov models. *Signal Processing*. 1995; 41:31–41.
- Vienne J, Bettler B, Franken P, Tafti M. Differential effects of GABAB receptor subtypes, gamma-hydroxybutyric Acid, and Baclofen on EEG activity and sleep regulation. *Journal of Neuroscience*. 2010; 30(42):14194–14204. [PubMed: 20962240]
- Vink JM, Groot AS, Kerkhof GA, Boomsma DI. Genetic analysis of morningness and eveningness. *Chronobiology International*. 2001; 18(5):809–822. [PubMed: 11763988]

Figure 1.
Mean Intra-epoch Aspect Ratio. A time-series plot of the mean intra-epoch aspect ratio for one mouse.

Figure 2.

Q-Q plots for the Geometric Distribution Fit to NREM, REM, and WAKE. The Q-Q line is given by the solid line and the $y=x$ line by the dashed line. The gray region is a null interval created using parametric bootstrap samples based on the maximum likelihood estimate.

Figure 3.

Bout Duration Q-Q plots for Each State Conditional on the Previous State. The Q-Q line is given by the solid line and the $y=x$ line by the dashed line. The gray region is a null interval created using non-parametric bootstrap samples which permute the true labels.

Figure 4.

A First Order Markov Model. In the first part, we observe both the Y_t and the X_t . In the second part, we observe the X_t and must predict the Y_t .

Figure 5. GMM Transition Diagram. The original state space is $S = \{a, b, c\}$ and the head sizes are $M_a = 2$, $M_b = M_c = 3$. All transitions within the “head” and “tail” durations are given. Of transitions away from states, only transitions away from $(b, 1)$ are given; similar transitions from $(a, 1)$ to each (b, n) and (c, n) and from $(c, 1)$ to each (a, n) and (b, n) are omitted for aesthetic reasons.

Figure 6.

TDGMM Transition Diagram. The original state space is $S = \{a, b, c\}$ and the head sizes are $M_{b,a} = M_{a,b} = M_{c,b} = 3$, $M_{c,a} = M_{a,c} = M_{b,c} = 2$. All transitions within the “head” and “tail” durations are given. Of transitions away from states, only transitions away from $(b, a, 1)$ are given; similar transitions from $(a, c, 1)$ to each (c, a, n) and (c, b, n) ; from $(a, b, 1)$ to each (b, a, n) and (b, c, n) ; from $(c, a, 1)$ to each (a, b, n) and (a, c, n) ; from $(c, b, 1)$ to each (b, a, n) and (b, c, n) ; and from $(b, c, 1)$ to each (c, a, n) and (c, b, n) are omitted for aesthetic reasons.

Figure 7.
Transition-Dependent Duration Distributions for the Simulation.

Figure 8. Simulation Results. The x -axis is the noise level σ of the covariate emission distributions. The y -axis is the classification error (row 1), the classification error relative to Bayes' Rule (row 2), and the RMSE of the probability estimates (row 3). The columns give the three training set sizes.

Figure 9.

Two Frames of Video Data. The left frame illustrates a typical awake posture whereas the right frame illustrates a typical sleep posture. Our tracking software imposes an ellipse upon the mouse which is used to calculate its size (area), aspect ratio, and velocity.

Figure 10.

Out-of-Sample ROC Curves for the Three-state REM vs. NREM vs. WAKE Task. The curves are for the variants of multinomial logistic regression. Those for the variants of random forests are nearly identical.

Figure 11. Minutes Spent in Each State. For the two-state SLEEP vs. WAKE task (left panel), we compare EEG/EMG-based manual scores to our video-based MLR+GMM. For the three-state NREM vs. REM vs. WAKE task (right panel), we compare EEG/EMG-based manual scores to our video-based MLR+TDGMM.

Table 1

Summary Statistics By Sleep State. Average duration is given in number of ten-second epochs.

Sleep State	Fraction of Epochs	Number of Bouts	Average Duration
NREM	43.99%	2,002	15.19
REM	4.82%	452	7.38
WAKE	51.19%	1,907	18.55

Table 2

Summary Statistics By Conditional Sleep State. Average duration is given in number of ten-second epochs.

Sleep State	Fraction of Epochs	Number of Bouts	Average Duration
NREM → REM	4.86%	446	7.45
WAKE → REM	0.01%	5	1.80
REM → NREM	3.12%	102	20.94
WAKE → NREM	41.25%	1,898	14.88
NREM → WAKE	48.19%	1,552	21.26
REM → WAKE	2.58%	350	5.04

Table 3

Out-of-sample Classification Error Rates. The first column gives the methodology and the second column gives the overall classification error rate for the two-state SLEEP vs. WAKE task. The third through sixth columns give respectively the overall classification error rate, the rate of REM prediction, the REM false positive rate, and the REM false negative rate for the three-state NREM vs. REM vs. WAKE task. We use a GMM for the two-state task and a TDGMM for the three-state task.

Method	Two-State Error Rate	Three-State			
		Error Rate	REM Rate	REM FP	REM FN
MLR	9.7%	14.9%	1.4%	1.2%	95.3%
MLR+IMM	8.4%	25.1%	18.3%	16.8%	52.4%
MLR+(TD)GMM	8.2%	21.9%	14.5%	13.0%	55.8%
RF	10.4%	16.2%	2.3%	1.9%	90.9%
RF+IMM	8.9%	24.7%	17.4%	15.9%	53.0%
RF+(TD)GMM	8.8%	23.3%	15.6%	14.0%	54.2%
40-Second Rule	10.1%	NA	NA	NA	NA
Manual Scores	4.8%	5.8%	NA	NA	NA

Table 4

Out-of-sample Classification Error Rates for Various External Methods on the Three-state NREM vs. REM vs. WAKE Task. The first column gives the methodology, the second column the overall classification error rate, the third column the rate of REM prediction, the fourth column the REM false positive rate, and the fifth column the REM false negative rate.

Method	Error Rate	REM Rate	REM FP	REM FN
AdaBoost	17.4%	2.7%	2.3%	89.9%
LogitBoost	20.7%	4.0%	3.7%	88.7%
Bagged Trees	15.2%	1.5%	1.2%	93.4%
MALLET CRF	15.1%	3.9%	3.3%	85.2%
TreeCRF 10b; 64l	16.9%	3.7%	3.4%	89.3%
MLR+TDGMM	21.9%	14.5%	13.0%	55.8%