A Comparison of Hidden Markov Model Sleep State Annotation using EEG and Cardiovascular Features.

Ghassemi M.M

Abstract - **We performed feature extraction from EEG and cardiovascular waveforms for use in a Hidden Markov Model (HMM) sleep state classifier. We tested the efficacy of the sleep classifier for a simple binary case, sleep versus non-sleep, and a more complex case with four classes: Sleep-REM sleep non-REM, Apnea event, awake. As our aim was to compare the efficacy of cardiovascular features as compared to EEG features we trained two distinctive classifiers, one which used only EEG features, and another which used only cardiovascular features. In addition to the standard HMM approach; where each state generates one observation vector, we also implemented a clustered observation HMM, where each state was modelled as generating a continuous set of observation vectors. We noticed that this approach significantly improved the performance of the EEG based classifier. Lastly, we demonstrated that a classifier which included both EEG and cardiovascular components outperformed classifier with only one set of features. To gauge the efficacy of our classifier, We compared it's performance against a recently used method proposed by Zhang et al. We observed that our HMM classifier outperformed the Zhang classifier in a binary classification task (sleep versus awake) with accuracy levels of 99% as compared to Zhang's 89%.**

1. Introduction

Since the introduction of the electronic medical record over a decade ago, retrospective analysis of medical data has provided a unique, and very low-cost way to perform hypothesis generation and testing in the medical community. There currently exist several clinical databases, with varying levels of information density. Of these, one of the richest and highest resolution publically available databases is MIMIC-II, which contains physiological recordings from several thousand patients at the Beth Israel Deaconess medical hospital from 2000 onwards, with over 700 users worldwide (Saeed & Al., 2011). One of the central issues with the MIMIC database, and other databases of its kind are a lack of comprehensive clinical annotation on the dynamic state of the patient during their stay. While the reason for hospital admission is clearly denoted, and daily summary scores of patient well-being such as SOFA or SAPS are available, it is difficult to extract patient cohorts on the basis of features which are not clearly measurable, but rather, require inference. Several of these hidden features are phenomenon of the mind and include a patient's state of mind, stress level or sleep stage.

We strongly believe that the inability to access, evaluate, or observes the dynamics of these hidden states is detrimental to efficacious retrospective clinical research (McEwen, 2008). Addressing this issue will require the creation of a clinical feature annotator, with an ability to learn the characteristics of hidden features, based on empirical measurements and training data from experienced clinical staff. This project makes the first and most basic step in this direction by tacking the issue of sleep state annotation.

To illustrate in a more concrete way the potential impact of hidden features on patient diagnosis we will briefly relate a project which is currently which is currently underway in the Laboratory of Computational Physiology, here at MIT. The team at the LCP is performing an investigation into the effects of circadian variation on patient state of health in the ICU. There is some evidence that circadian variation may be a means of understanding a patient's overall state of health (Lanthier et al., 2011; Vincent, 2011). A circadian variation can be defined as any biological process that periodically varies over a given time period. Human heart rate and blood pressure, for instance, are known to be subject to this form of variation, being higher while patients are awake, and lower when patients are asleep (Dean et al., 2012). There is evidence that the magnitude of this variation may be correlated with a patient's overall state of health, with higher variation tending to be better. Hence, if monitored, circadian variation could potentially allow clinicians to measure if a patient's health is deteriorating, or improving, over the course of their hospital stay.

The central challenge in retrospective analysis of circadian variation is figuring out when the patients were asleep. One recently used strategy by Zheng et al is naive in that the sense that it identifies a single, continuous timeframe for when the patient was likely asleep (Zheng et al, 2013). There are many ways to infer this timeframe, but one proposed method utilizes differential blood pressure and heart rate values, averaged across days. This strategy is sub-optimal for two reasons. Firstly, it assumes that critically ill patients will be sleeping and awake at consistent time intervals, which is clearly unlikely. Secondly, the identification of the sleep state currently relies exclusively on the blood pressure and heart rate of the patient, which wrongfully assumes that the cardiovascular dynamics of a critically ill patient will be similar to those of a healthy one. (Baumgart et al., 1991)

Hence, our specific aim in this paper is to outline an improved method for sleep state identification with hopes of facilitating more accurate analysis of the effects of circadian variation on clinical outcomes in an intensive care setting. To perform sleep state identification in a principled way, we propose a Hidden Markov Model approach, which would use features from an EEG monitoring system in addition to features from the blood pressure and EKG waveforms. We advocate for the use of EEG in this setting as sleep has several well-known EEG related signatures, and furthermore we expect less severe fluctuations in the state of a patient's brain during critical illness as compared to fluctuations in cardiovascular features.

2. Data

2.1 Source: All data used in this study was collected from the MIT-BIH Polysomnographic Database, which is publically accessible via the Physionet project, an open source archive of recorded physiological signals (Goldberger et al., 2000). Our dataset included 13 subjects which were monitored in the Beth Israel Hospital Sleep Laboratory. Subjects were monitored for the evaluation of chronic obstructive sleep apnea syndrome. The subjects contained anywhere from 4-9 physiological signal recordings, which included the Electrocardiogram (ECG), Blood pressure (BP), C4/A1 Electroencephalogram lead (shown in the figure below), nasal and abdominal respiration values among others signals. The signals of interest which all subjects had in common were EEG, ECG and BP waveforms. Physiological signals were sampled by a 12 bit analog-to-digital convertor at a sampling rate of 250Hz for a total of 6 hours (5.4 million samples)

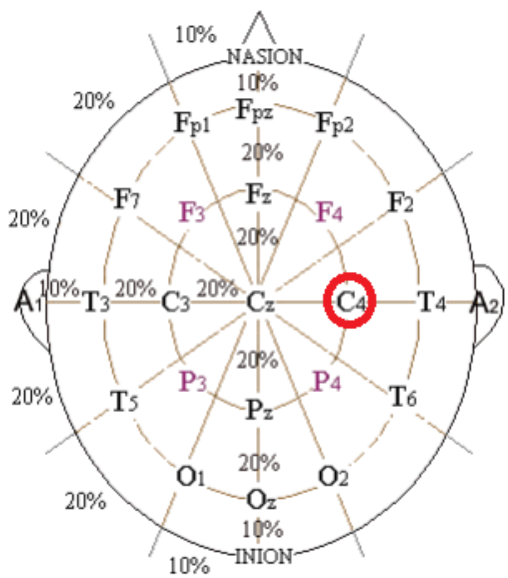

Figure 1: An illustration showing the location of the C4/A1 Electroencephalogram lead.

2.2 Feature Extraction When performing time series analysis, the classical approach in signal processing and machine learning literature is to extract relevant features from the waveforms of interest. These features may then be used to train models for the purpose of prediction. If done properly, feature extraction maintains most of the information content in the original signals while compressing data size, and allowing for a more specific understanding of the features which are predictive of certain phenomenon.

The primary burden of feature selection, however, is that it often requires special knowledge to be done most effectively. Being told that the alpha band of an EEG is a useful feature for coma detection and discovering that the 8-13Hz band of an unlabelled signal is important for predicting coma are clearly different challenges. Fortunately in our case there is vast realm of literature which characterizes the physiological changes that accompany sleep state, and this can guide our feature extraction. The most commonly cited examples of features that accompany sleep state are reductions in blood pressure and heart rate as well as a progressive decrease in the activation rate of neurons in the brain (as sleep progresses from wakefulness to non-rem sleep) (Dean et al., 2012; Saper, Chou, & Scammell, 2001). Neuronal activity is also known to change from chaotic to more coordinated during the non-rem portion of sleep (Maquet, 2001). The classification of REM sleep, on the other hand, is quite a different challenge and not easily inferred as it exhibits characteristics of the wakeful state at nearly all levels. We based our own features on those described above. While there are a variety of features that could potentially be extracted from the EKG waveform. We suspect that the most information dense feature for the type of classification we wished to perform was the RR-interval. The RR-interval describes the beat-to-beat distance and therefore provides information on the pace of cardiac contraction, which is known to vary as a function of sleep state. like EKG, there is more than one feature that can be extracted from a blood pressure waveform. We note here that the cardiovascular system adjusts the RR interval, primarily to control changes in blood pressure. Hence, we decided to extract the median blood pressure value over the RR-interval as a ways of approximating the waveform's behaviour. The dataset contains a single EEG recording from the C4/A1 lead, shown pictorially in Figure 1. We utilized a 6 level wavelet based decomposition of the EEG signal into its corresponding alpha, beta, gamma, theta, and delta bands. Finally we used the power content of the bands over the periods defined by the RR interval as our EEG features.

2.3 Annotations and Preprocessing: As the data was collected from a sleep laboratory, it came with multiple annotation files for each subject. Among other things, these annotation files described the onset of an apnea related events, movements, and the sleep stage of the subject according to the Manual of Sleep Classification by Rechtschaffen and Kales in which recordings are divided into 7 discrete stages (wake, stage 1, stage 2, stage 3, stage 4, stage REM, and movement time). According to the metric outlined by Rechtschaffen and Kales, sleep annotations were valid for 30 sec intervals. Hence, we reformatted the annotation files such that they described the annotation at a per-sample level, and thereby allowed for their use in the analysis.

2.5 Analysis: All analysis for this project was done in the Matlab (2012b) technical computing language. In addition to the use of several Mathwork's toolboxes, our approach utilized the Hidden Markov Model (HMM) and Bayes Net Toolbox authored by Kevin Murphy.

3. Methods

3.1 Classification Model We constructed our activity classifier using a 2-component Mixture of Gaussians model for each of the four classes $S = \{nonrem, rem, awake, apnea\}$

3.2 General HMM approach: The parameters of *K* mixture of Gaussians models may be optimized via the Expectation Maximization algorithm given that a set of emitted observations $O_{1 \dots T}$ and annotated hidden state values $S_{1 \dots T}$ are known. Once the maximum likelihood values for the parameter sets $\theta_{1 \dots K}$ have been identified, classification of novel data points can then be performed by comparing $P(O| \, S, \theta_i) \; for \; all \; i \in \{1 \dots K\}$. In the case of a binary classification task (sleep versus awake) this would equate to comparing $P\big(O \big| \, sleep, \theta_{sleep}\big)$ to $P(O| \, awake, \theta_{awake})$. While this model is certainly valid, it has the distinctive disadvantage that it assumes complete independence of S_{t+1} from S_t. A more realistic assumption would be that the system stochastically transitions to a new state according to the probability distribution$P(S_{t+1}| \, S_t)$. Fortunately Hidden Markov Models (HMMs) allow us to introduce this state dependency across time. Furthermore, it is conceptually trivial to extend a MoG framework to a HMM. As the name implies, an HMM is a Markov chain, where each state generates an observation vector and may be represented with the following graphical model.

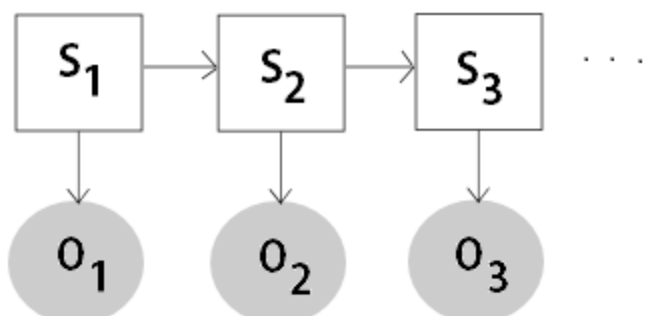

Figure 2: A Hidden Markov Model

The Markov chain encodes the joint distribution. $P(S|O) = P(S_1|O_1)P(S_2|O_2)P(S_3|O_3) \dots$ by assuming that the value of future states is conditionally independent of the past states, given the value of the current state. Given the markovian property, HMMs are quite useful for time-series modelling and are immensely well documented machines learning techniques that have been employed for large realm of classification purposes. Indeed, this is certainly not the first time that HMMs have been used for sleep classification (Doroshenkov, Konyshev, & Selishchev, 2007; A. Flexer, Sykacek, Rezek, & Dorffner, 2000; Arthur Flexer, Gruber, & Dorffner, 2005).

The goal in our case is to train the HMM such that we may infer a hidden state sequence given a novel set of retrospective observations$o_1 \dots o_T$. Given a set of training data and observations as before (as well as initializations of parameter values), The Baum Welsh algorithm (which is an

EM algorithm for HMM parameters) can be used to maximize the likelihood of model parameters. Once the parameters of the model are optimized, the Viterbi algorithm may be used to identify the most likely corresponding state sequence for the observations. An important thing to note here is that the central parameter that distinguishes an HMM from a naive MoG classifier is it's state transition probability matrix, $P(S_{t+1}|\, S_t)$. As there is nothing particularly novel about our methods, we refer the curious or unfamiliar reader to some well-written sources that explain HMMs and EM in greater detail (BILMES, 2006; Caelli, Amin, Duin, Ridder, & Kamel, 2002; Smyth, Heckerman, & Jordan, 1997). As we have done thus far, we will speak to the HMM approach only so much as it allows us to clearly explain our methodological approach.

3.3 Clustered Observation HMM: In addition to the standard HMM approach where each state generates one observation vector; we also implemented an HMM classifier where each state generates *N* successive observation vectors. This approach is illustrated in the following graphical model:

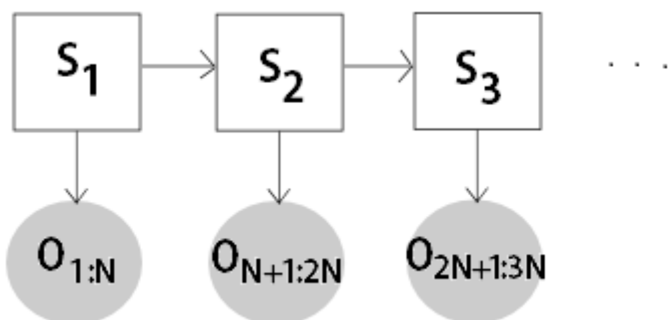

Figure 3: A graphical model for an HMM with a clustered emission vector.

In this approach the hidden state, S_{t+1}, would depend on the previous value of S_{t+1} as before, but would also generate *N* consecutive observation vectors. In this case, the probability can be written as $P(S|O) = P(S_1|O_1)P(S_2|O_2, O_1)P(S_3|O_3\,, O_2, O_1\,) \ldots\ P(S_T|O_T, O_{T-1}, \ldots O_{T-N})$. By modelling each state as outputting a set of consecutive observation vectors, we will allow our classifier to be sensitive to the evolution of observation values, over some time frame as opposed to the observation at a single point in time.

3.4 HMM/EM Parameter Initialization and Choice of Prior:

The initial choice of HMM parameter values is known to have an impact on the convergence of the Baum Welsh algorithm to the actual maximum likelihood parameter values. In our case, we randomly initialized these parameter values. We also set the prior distribution over initial states as uniform. The EM algorithm was allowed to run for 4 iterations or until convergence. The small number of iterations was chosen for the sake of time.

4. Experimental Approach

4.1 Feature Validation: given the potential perils of improper feature selection mentioned above, we began our analysis with validating that the selected features were valid for the purpose of the classification task. To do this, we inspected the mean and standard deviation of the features, across the non-REM sleep, and awake conditions. Table 1 shows this comparison:

Feature	Mean (sleep/awake)	Standard Deviation(sleep/awake)
EEG: alpha	.1 / .12	.12 / .15
EEG:beta	.22 / .24	.21 / .27
EEG:delta	.12 / .11	.21 /.36
EEG:gamma	.33 / .20	.29 / .30
EEG:theta	.06 / .20	.13/ .33
RR Interval	1.09 / .40	.05 /.05
Median BP	.79 / .18	.74 / .36

Table 1: Mean and standard deviation of features across asleep and awake conditions.

We notice in Table 1 that our cardiovascular features show very strong differences across the two states. This implies that the RR interval and Median BP are indeed good features for this binary task. We also notice here that several of the EEG components are similar in terms of their means and standard deviations. This caused concern about the validity of our selected EEG features.

To address this concern, we performed a comparison between a simple binary classifier (awake vs. non-REM sleep) trained on the raw EEG and cardiovascular waveforms, versus one trained using our extracted EEG, and cardiovascular features separately. We trained the feature classifiers using 80% of available data, and tested on 20% of the remaining data. We then compared the effectiveness of the classifiers according to overall accuracy on classifying novel data. See table 2 below for results.

Classification Approach	**Non-Rem Sleep**	**Awake**	**Overall Accuracy**
Zheng et al	93%	85%	89%
Raw waveforms	79%	62%	71%
EEG Features	98%	45%	72%
Extracted Cardiovascular features	99%	99%	99%

Table 2: Comparison of binary classifiers using extracted features versus raw data, versus Zheng et al. method. Numbers shown refer to the portion of correctly classified points.

3.2 Classification using Clustered EEG Observations:

Having validated the efficacy of our cardiovascular features for classification purpose above raw data features and the approach described by by Zheng et al (please see appendix A), the next logical step was to see if we could improve the efficacy of the EEG based classifier by utilizing the clustered observation approach. To do this we varied the cluster size from 0-15, trained the classifier, and computed the classification accuracy for each task on novel data. The results are depicted in figure 4 below:

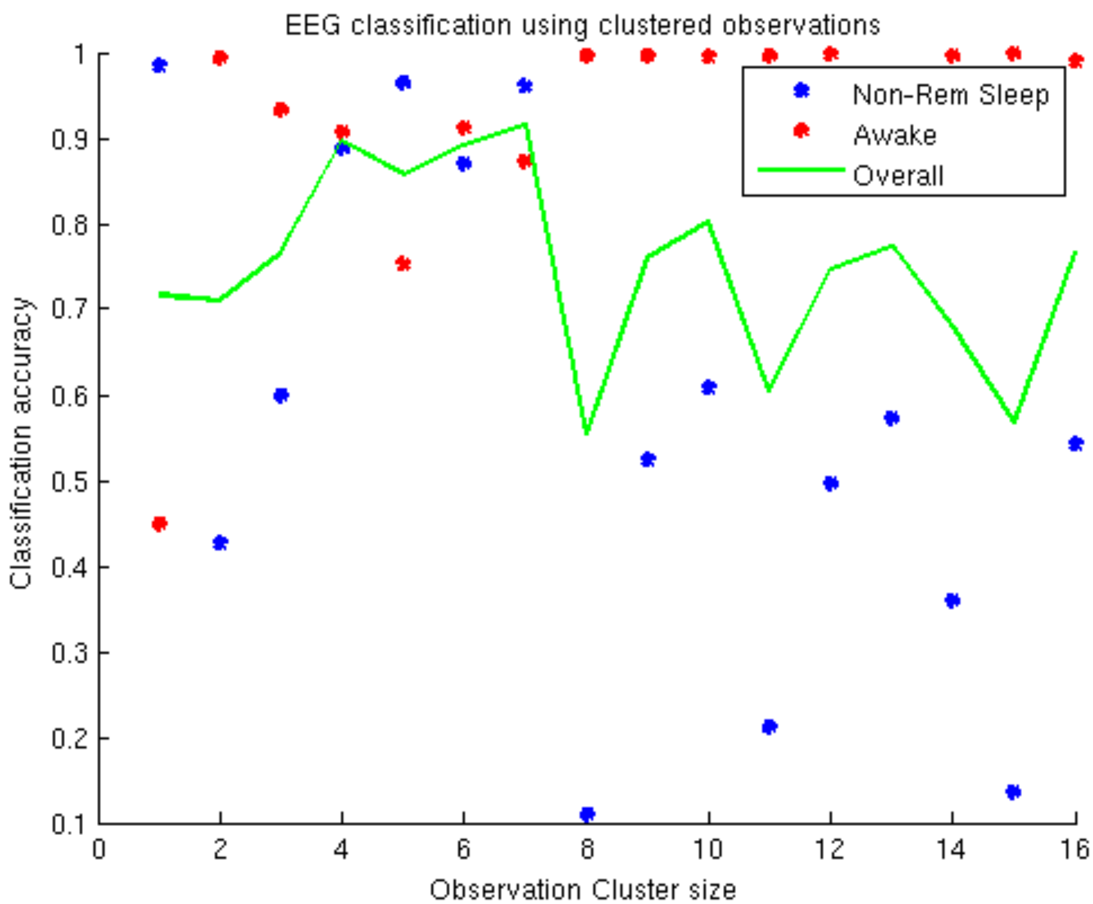

Figure 4: Sleep/awake classification using various clustered observation sizes.

We can see from the figure that the EEG features classifier gains nearly 20% improvement in predictive accuracy when utilizing the clustered observation approach. Notice above that the overall classification accuracy is maximized for a cluster size of 7, yielding a Non-rem sleep state classification of 95%, awake state classification of 87% and an overall classification potential of 92%. Figure 4 also illustrates a rather curious phenomenon, as the observation cluster size is increased; the classifier is better able to detect the sleep state, and less able to detect the awake state.

3.2 Multi-Class Sleep Stage Classifier using Cardiovascular vs. EEG Features:

Given the impressive performance of the sleep state classifier which used cardiovascular features, we thought it might be more interesting to see how well an EEG versus cardiovascular classifier would work for classifying a larger number of sleep-related states. Hence, we decided to add two additional states for classification: REM-sleep and Apnea. It is obvious that such a classifier is more useful in a realistic clinical context. As before, we trained the classifiers using EEG features and cardiovascular features separately utilizing 80% of the available data for training, and the remaining 20% for testing. In the binary case, our cardiovascular based classifier performed so well that we didn't attempt the more computationally intensive clustered observation approach. In this case, we will perform the clustered approach for both to check for trends in performance with observation cluster size. Results are shown in the figure below:

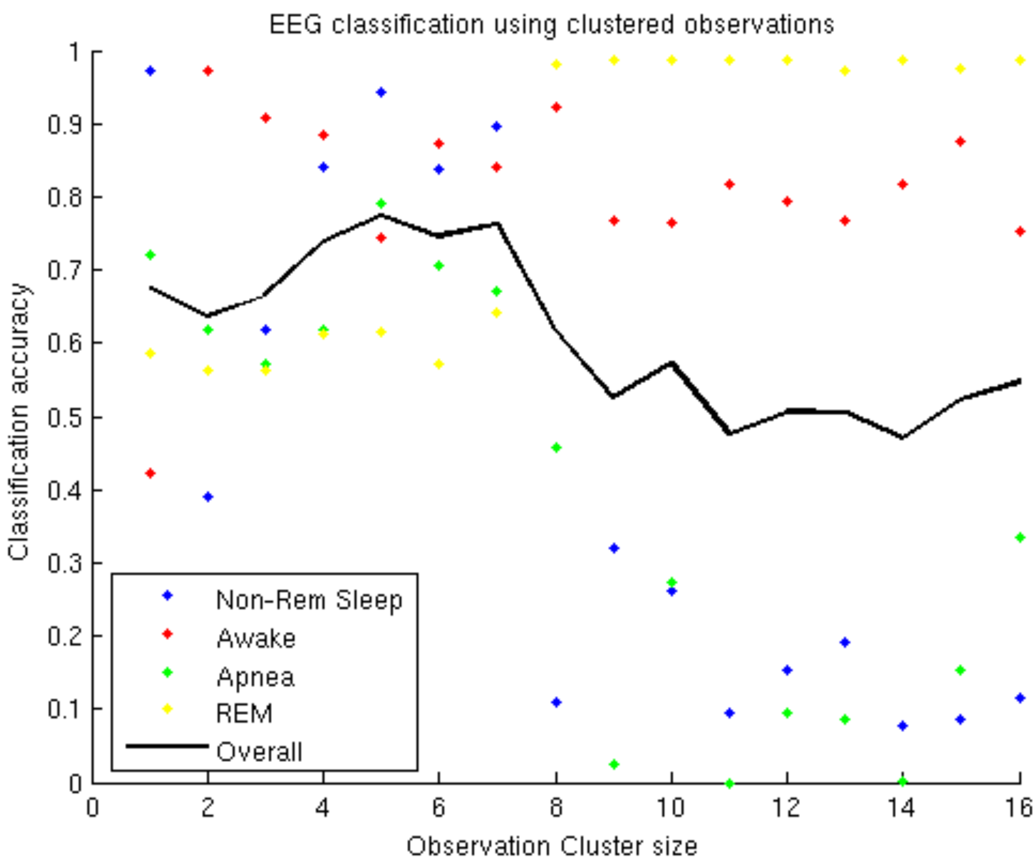

Figure 5: The performance of a multi-class EEG based sleep stage classifier over various observation cluster sizes.

Predicted Actual	Awake	Non-REM sleep	REM	Apnea
Awake	**84%**	14%	2%	0%
Non-REM sleep	0%	**88%**	7%	5%
REM	31%	2%	**63%**	4%
Apnea	14%	13%	6%	**67%**

Table 3: The multi-class EEG based sleep stage classifier's Predicated versus actual class performance with a cluster size of 7.

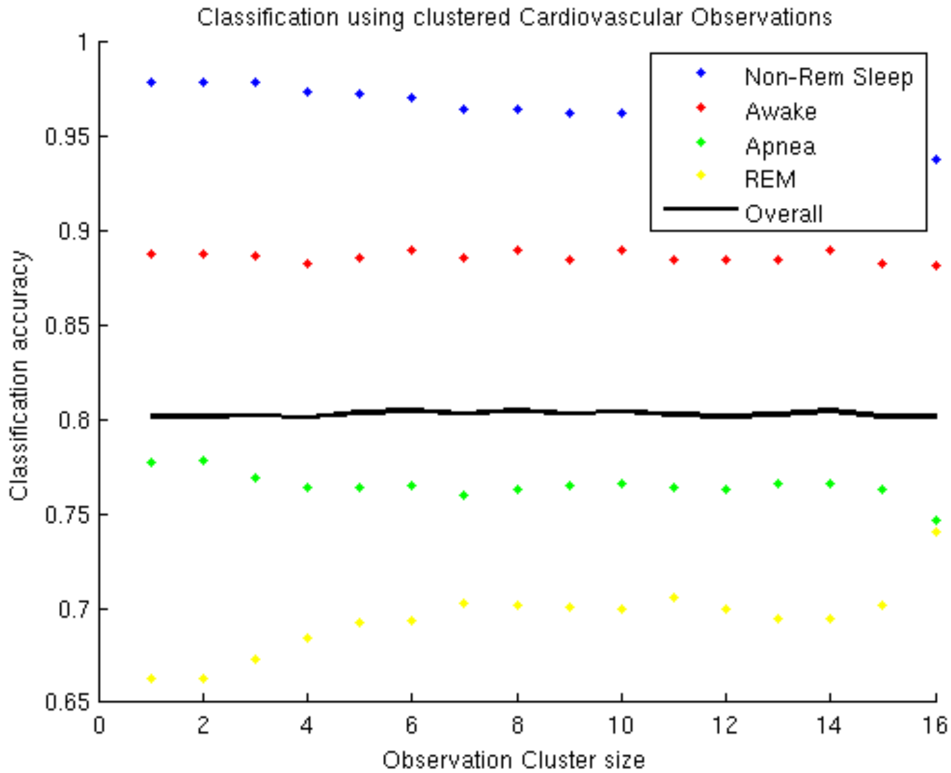

Figure 6: The performance of a multi-class Cardiovascular based sleep stage classifier over various observation cluster sizes

Predicted Actual	Awake	Non-REM sleep	REM	Apnea
Awake	**89%**	7%	4%	0%
Non-REM sleep	>1%	**98%**	>1%	1%
REM	26%	7%	**66%**	>1%
Apnea	18%	0%	4%	**78%**

Table 4: The multi-class cardiovascular based classifier's predicted versus actual class performance with a cluster size of 1.

If we briefly compare the results of figure 5/6, and tables 3/4 we notice immediately that, as before, the EEG based features gain in performance for particular cluster sizes. The cardiovascular features on the other hand, seem unaffected by changes in cluster size. We also notice that, at their best, the cardiovascular based classifier still outperforms the EEG based classifier.

3.2 Multi-Class Sleep Stage Classifier using Cardiovascular and EEG Features:

The results from section 3.1 demonstrate impressive performance levels, but they still fall short of other classifiers in the literature which use the same dataset (Luay et al, 2009). As a final step we included both the EEG and cardiovascular features for the classifier and performed the same approach as section 3.1. The results for the classifier are shown below and include classification, table, performance plots, and the transition matrix for the optimal case.

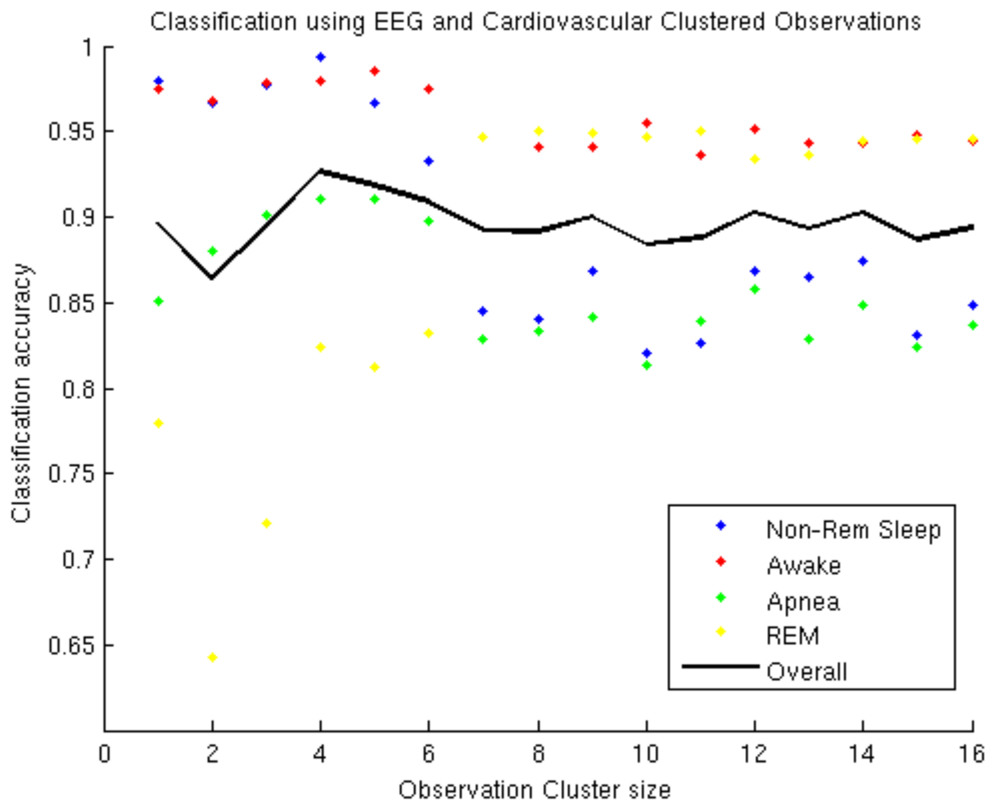

Figure 7: The performance of a multi-class EEG and cardiovascular based sleep stage classifier over various observation cluster sizes. The optimal classification performance occurred at a cluster size of 4, and was 93%.

Predicted Actual	Awake	Non-REM sleep	REM	Apnea
Awake	**97%**	3%	0%	0%
Non-REM sleep	1%	**99%**	0%	0%
REM	1%	6%	**91%**	2%
Apnea	8%	8%	3%	**82%**

Table 5: The multi-class cardiovascular and EEG based classifier's predicted versus actual class performance with a cluster size of 4.

TO FROM	Awake	Non-REM sleep	REM	Apnea
Awake	**82%**	13%	4%	1%
Non-REM sleep	3%	**87%**	5%	5%
REM	1%	1%	**86%**	13%
Apnea	13%	2%	3%	**82%**

Table 6: Transition Probability Matrix of The multi-class cardiovascular and EEG based classifier's with a cluster size of 4.

The classifier described in Table 5/6 demonstrates that the combination of EEG and cardiovascular features significantly improves the performance of the classifier above one which utilizes only cardiovascular or EEG features. The transition probability matrix also matches with our intuitions. It captures that patients who are awake are more likely to enter Non-REM sleep than REM sleep and that patient with Apnea attacks are more likely to wake up than return to Non-REM sleep.

5. Discussion and Future Work

Our results demonstrate that an HMM framework is an effective means to perform classification of hidden features, such as the sleep state of patients. Moreover, our results show that a clustered emissions approach, for certain values of *N*, provides an enhanced ability to perform the classification task, especially in the case of EEG features. When we extended our simple binary classifier to include the Apnea and REM-sleep states, we were presently surprised with the accuracy of the classifier. Indeed, when we compare the results of our classifier to that proposed by others in the literature, we see comparable results.

5.1 Cardiovascular Classification of Sleep State

Table 2 clearly reflects our observations from Table 1, the large difference in the means of the RR Interval and Median BP across the classes makes the cardiovascular features particular potent in the binary classification case. To our surprise, the cardiovascular features also perform classification rather well in the multiclass case. As a follow up to this study, it may be interesting to extract more cardiovascular features and see if we could improve the performance of the cardiovascular classifier to an even greater level.

5.2 EEG Classification of Sleep State

The EEG signal is known to be closely related to the level of consciousness of the person and the C4/A1 lead is often used for sleep state identification. As the activity increases, EEG shifts to higher dominating frequency with lower amplitudes. When the eyes are closed, the alpha waves are known to dominate the EEG signals, which we observed in our case. We also observed that when patients fall asleep, the dominant EEG frequency often decreases. This is one possible reason why the clustered observation HMM works better than the non-clustered HMM for the EEG case, as it can take these changes as an indicator of a change in state.

5.3 EEG vs. Cardiovascular Classification

We noticed that, at their best, the cardiovascular based classifiers always outperformed the EEG based classifier. There are several reasons why this might have occurred. Firstly, it is possible that our selected EEG features were not encoding pertinent information for the classification task. Secondly, we remind the reader that this study employed the use of a single EEG electrode across multiple subjects. It is quite possible that slight differences in positioning of the electrodes across these subjects introduced noise to the result which attenuated the performance of the EEG based classifier. It may be interesting to perform a study with additional electrodes and see to what extent an EEG based classifier alone may be improved.

5.4 Choosing Better Initial HMM Parameters

Our HMM parameters were initialized randomly, it is therefore possible that the HMM parameters were not converging to a global optima. Global optimization grows increasingly important as the dimensionality of the feature space increases. Hence, as an eventual extension to this project we would like to perform global optimization on our HMM parameters via a multi parameter initialization process.

5.5 Application of Classifier to Existing Databases

In large part, this classifier was developed to help annotate subjects in the MIMIC-II clinical database for the purpose of investigating circadian variation. Given the results we have found in this study, we feel comfortable applying this approach for annotation purposes moving forward. As we mentioned earlier, one of the key challenges with extending this onto an ICU database is that patients are critically ill and may feature activity different than that of their healthy counterparts.

6. Acknowledgements:

I would like to sincerely thank the course staff for facilitating a fascinating seminar and providing me with an opportunity to present my work. I would also like to thank Dr. Riccardo Barbieri, and Dr. Mengleng Feng for their advice as I worked through this project.

Appendix A: Zheng et al approach to sleep state annotation:

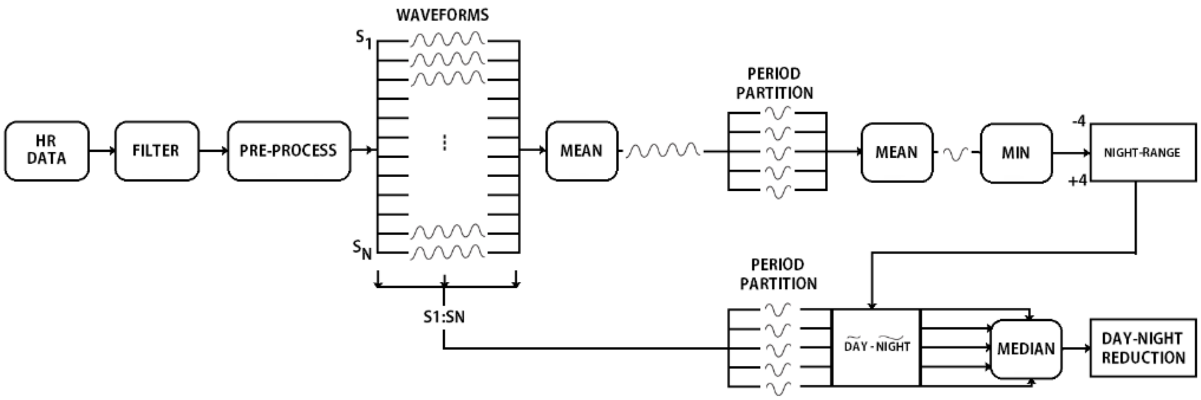

Figure AXXX: A figure depicting the method for sleep state extraction proposed by Zheng et al.

Zheng et al utilized a data-driven approach to infer the sleep and wake cycle on a population level using the fluctuations of the circadian rhythm. After time series alignment, they computed the average hourly signal value. We will refer to these measures as W_k. Next they averaged the values of W_k, across days, to create a single 24 hour waveform which represented each subgroup, called W_d. They computed W_d using the following approach: Let $\widehat{h_i}$ be a vector describing hourly values of the signal on day i Then we can define the average 24 hour waveform for each subgroup as:

$$W_d = \frac{\sum_i \widehat{h_i}}{|i|}$$

The *sleep state* interval was defined as the 8hr window around the minimum of the $W_d$waveform while all remaining hours were classified as *awake state interval*. Given that our dataset was not over a 24 hour period, we rescaled the 8 hour timeframe to the average the average amount of time spend sleeping.

Bibliography

Baumgart, P., Walger, P., Gemen, S., Von Eiff, M., Raidt, H., & Rahn, K. H. (1991). Blood Pressure Elevation during the Night in Chronic Renal Failure, Hemodialysis and after Renal Transplantation. *Nephron*, *57*(3), 293–298. doi:10.1159/000186278

BILMES, J. A. (2006, March 1). What HMMs Can Do. *IEICE TRANSACTIONS on Information and Systems*. The Institute of Electronics, Information and Communication Engineers. Retrieved from http://search.ieice.org/bin/summary.php?id=e89-d_3_869&category=D&year=2006&lang=E&abst=

Caelli, T., Amin, A., Duin, R. P. W., Ridder, D., & Kamel, M. (Eds.). (2002). *Structural, Syntactic, and Statistical Pattern Recognition* (Vol. 2396). Berlin, Heidelberg: Springer Berlin Heidelberg. doi:10.1007/3-540-70659-3

Dean, E., Bloom, A., Cirillo, M., Hong, Q., Jawl, B., Jukes, J., Nijjar, M., et al. (2012). Association between habitual sleep duration and blood pressure and clinical implications: a systematic review. *Blood pressure*, *21*(1), 45–57. doi:10.3109/08037051.2011.596320

Doroshenkov, L. G., Konyshev, V. A., & Selishchev, S. V. (2007). Classification of human sleep stages based on EEG processing using hidden Markov models. *Biomedical Engineering*, *41*(1), 25–28. doi:10.1007/s10527-007-0006-5

Flexer, A., Sykacek, P., Rezek, I., & Dorffner, G. (2000). Using hidden Markov models to build an automatic, continuous and probabilistic sleep stager. *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium* (Vol. 3, pp. 627–631 vol.3). IEEE. doi:10.1109/IJCNN.2000.861392

Flexer, Arthur, Gruber, G., & Dorffner, G. (2005). A reliable probabilistic sleep stager based on a single EEG signal. *Artificial intelligence in medicine*, *33*(3), 199–207. doi:10.1016/j.artmed.2004.04.004

Goldberger, A. L., Amaral, L. A. N., Glass, L., Hausdorff, J. M., Ivanov, P. C., Mark, R. G., Mietus, J. E., et al. (2000). PhysioBank, PhysioToolkit, and PhysioNet : Components of a New Research Resource for Complex Physiologic Signals. *Circulation*, *101*(23), e215–e220. doi:10.1161/01.CIR.101.23.e215

Lanthier, L., Touchette, M., Bourget, P., St-Georges, C., Walker, C., & Tessier, D. M. (2011). [Evaluation of circadian variation of blood pressure by ambulatory blood pressure monitoring in an elderly diabetic population with or without orthostatic hypotension]. *Gériatrie et psychologie neuropsychiatrie du vieillissement*, *9*(1), 59–66. doi:10.1684/pnv.2011.0257

Maquet, P. (2001). The role of sleep in learning and memory. *Science (New York, N.Y.)*, *294*(5544), 1048–52. doi:10.1126/science.1062856

McEwen, B. S. (2008). Central effects of stress hormones in health and disease: Understanding the protective and damaging effects of stress and stress mediators. *European journal of pharmacology*, *583*(2-3), 174–85. doi:10.1016/j.ejphar.2007.11.071

Saeed, M., & Al., E. (2011). Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC-II): A public-access intensive care unit database. *Crit Care Medicine*, *39*(February 2010), 952–960. doi:10.1097/CCM.0b013e31820a92c6.Multiparameter

Saper, C. B., Chou, T. C., & Scammell, T. E. (2001). The sleep switch: hypothalamic control of sleep and wakefulness. *Trends in Neurosciences*, *24*(12), 726–731. doi:10.1016/S0166-2236(00)02002-6

Smyth, P., Heckerman, D., & Jordan, M. I. (1997). Probabilistic Independence Networks for Hidden Markov Probability Models. *Neural Computation*, *9*(2), 227–269. doi:10.1162/neco.1997.9.2.227

Vincent, M. A. Biomarker circadian rhythm profiles in critically ill mechanically ventilated patients. , Texas Medical Center Dissertations (via ProQuest) 1 – 194 (2011). Retrieved from http://digitalcommons.library.tmc.edu/dissertations/AAI3459700