Analysis of Wake/Sleep EEG with Competing Experts

J. Kohlmorgen[1], K.-R. Müller[1], J. Rittweger[2], K. Pawelzik[3]

[1] GMD FIRST, Rudower Chaussee 5, D-12489 Berlin, Germany
e-mail: {jek, klaus}@first.gmd.de
[2] Institut f. Physiologie der Freien Universität Berlin, D-14195 Berlin, Germany
e-mail: ritmus@zedat.fu-berlin.de
[3] MPI für Strömungsforschung and SFB 185: Nonlinear Dynamics
Bunsenstr. 10, D-37073 Göttingen, Germany
e-mail: klaus@chaos.uni-frankfurt.de

Abstract. An analysis of physiological wake/sleep data is presented. We apply a recent method for the analysis of nonstationary time series with multiple operating modes. In particular, it is possible to detect and to model a switching of the dynamics and also a less abrupt, time consuming drift from one mode to another. This is achieved by an unsupervised algorithm that segments the data according to inherent modes, and a subsequent search through the space of possible drifts. The application to wake/sleep data demonstrates that analysis and modeling of real-world time series can be improved when the drift paradigm is taken into account. In the case of wake/sleep data, we hope to gain more insight into the physiological processes that are involved in the transition from wake to sleep.

1 Introduction

Modeling dynamical systems through a measured time series is commonly done by reconstructing the state space with time-delay coordinates [7, 9]. The prediction of the time series can then be accomplished by training neural networks [10]. If, however, a system operates in multiple modes and the dynamics is *drifting* or *switching*, standard approaches like multi-layer perceptrons are likely to fail to represent the underlying input-output relations. Moreover, they do not reveal the dynamical structure of the system. Time series from alternating dynamics can originate from many kinds of systems in physics, biology and engineering. In this contribution, we present an analysis of physiological data (EEG, respiration) from a sleep laboratory.

In [1, 5, 8], we have described a framework for time series from *switching* dynamics, in which an ensemble of neural network predictors specializes on the respective operating modes. Recently, we extended the ability to describe a mode change not only as a switching but – if appropriate – also as a drift from one predictor to another [2]. The application to wake/sleep data indicates that physiological signals contain drifting dynamics, which underlines the potential relevance of our method in time series analysis.

2 Detection of Drifts

The detection and analysis of drifts is performed in two steps. *First*, an unsupervised (hard-)segmentation method is applied [1, 5, 8]. In this approach, an ensemble of prediction experts f_i, $i = 1, ..., N$, is trained by maximizing the likelihood that the ensemble would have generated the time series. For the derivative of the log-likelihood with respect to the output of an expert, we get (cf. [1])

$$\frac{\partial \log L}{\partial f_i} \propto \left[\frac{e^{-\beta(y-f_i)^2}}{\sum_j e^{-\beta(y-f_j)^2}} \right] (y - f_i), \tag{1}$$

where y is a data point to be predicted. This learning rule can be interpreted as a weighting of the learning rate of each expert by the expert's relative prediction performance. Furthermore, we imposed a low-pass filter on the prediction errors and used deterministic annealing in the training process (see [1, 5, 8] for details).[4]

As a prerequisite of this method, mode changes should occur infrequent, i.e. between two mode changes the dynamics should operate stationary in one mode for a certain number of time steps. Applying this method to a time series yields a (hard) segmentation of the series into different operating modes together with prediction experts for each mode. In case of a drift between two modes, the respective segment tends to be subdivided into several parts, because a single predictor is not able to handle the nonstationarity.

The *second* step takes the drift into account. A segmentation algorithm is applied that allows to model drifts between two stationary modes by combining the two respective predictors, f_i and f_j. The drift is modeled by a weighted superposition

$$f(\mathbf{x}_t) \;=\; a(t)\, f_i(\mathbf{x}_t) \;+\; (1 - a(t))\, f_j(\mathbf{x}_t), \quad 0 \le a(t) \le 1, \tag{2}$$

where $a(t)$ is a mixing coefficient and $\mathbf{x}_t = (x_t, x_{t-\tau}, \ldots, x_{t-(m-1)\tau})^T$ is the vector of time-delay coordinates of a (scalar) time series $\{x_t\}$. Furthermore, m is the embedding dimension and τ is the delay parameter of the embedding.

The drift segmentation algorithm is unsupervised and performs a complete search for the optimal segmentation with the lowest average prediction error, i.e. it determines the drift parameters $a(t)$, $i(t)$, $j(t)$. The search is performed efficiently by dynamic programming. A detailed description of the algorithm is beyond the scope of this contribution, however, it can be found in [2].

In the following application we use radial basis function (RBF) networks of the Moody-Darken type [4] as prediction experts, because they offer a fast and robust learning method. A single run of the algorithm typically takes less than 30 seconds for the training step and 1-3 minutes for the segmentation step, measured on a Sun Ultrasparc.

[4] Further information and papers can be found at: `http://www.first.gmd.de/persons/Kohlmorgen.Jens.html`

3 Application to EEG and Respiration Data

In [6], we analyzed physiological data recorded from the wake/sleep transition of a human. The objective was to provide an unsupervised method to detect the sleep onset and to give a detailed approximation of the signal dynamics with a high time resolution, ultimately to be used in diagnosis and treatment of sleep disorders. From the neurophysiological point of view, the sleep onset corresponds to a reorganization of the neural network in the reticular formation of the brain stem, which regulates and integrates cardiovascular, respiratory and somatomotor systems and vigilance [3]. This reorganization can be thought of as a transition to a different mode of its dynamics. The use of the drift segmentation algorithm now yields a more detailed modeling of the dynamical system.

Fig. 1 shows a comparison of drift- and hard-segmentation by the computer, using 8 RBF predictors, versus a manual segmentation by a medical expert. The experimental data was measured during an afternoon nap of a healthy human. The computer-based analysis is performed on a single-channel EEG recording (occipital-1), whereas the manual segmentation was worked out using six physiological signals (EEG, EOG, ECG, heart rate, blood pressure, respiration).

The drift algorithm yields several drift parts. The sleep onset, according to the manual segmentation at $t \approx 4000$, is represented by an exponential drift from a wake-state predictor, net 7, to a sleep-state predictor, net 4. On the other hand, the arousal is introduced at $t \approx 9000$ by a slight drift back to net 7, which holds until the wake-up point is reached ($t \approx 9500$ in the manual segmentation). There, a sudden change of the mixing coefficient gives more weight to wake-state net 7. After $t \approx 9800$ (eyes open), a mixture of two wake-state nets, 2 and 7, is employed for prediction. Compared to both hard segmentations, the drift segmentation reveals the evolution of the dynamical changes at the transitions between different wake/sleep stages. Moreover, when allowing drift, just three nets are sufficient to completely describe the dynamics.

To test the generalization ability of the trained networks, the drift segmentation of a second sleep experiment was computed. A new EEG recording of the same person was analyzed by the ensemble already obtained in the first experiment, without any further training on the new data. The result is shown in Fig. 2: The wake states are modeled again by nets 2 and 7, the sleep states by net 4. Most of the intermediate arousals indicated in the manual segmentation are found as drifts towards wake-state net 7. Furthermore, the final arousal is indicated properly. In summary, the pre-trained ensemble generalizes well and yields good results even on test data.

Next, we analyzed the respective respiration data (Fig. 3, 4). They were obtained in the same two experiments by measuring the thoracic excursion. In both cases, the ensemble of experts was *trained* on the respective data. The resulting hard- and drift segmentations are in good accordance with the manual segmentation. Interestingly, the wake/sleep transitions in the drift segmentation turned out to be more coarse than in the EEG. From the physiological point of view, however, it is known that the respiration does not reveal as many details of the sleep stages as the EEG.

4 Summary and Discussion

An analysis of physiological wake/sleep data was presented. It employs a method for the unsupervised segmentation and identification of nonstationary drifting dynamics. The application to wake/sleep data demonstrated that drift can be found in natural systems. Therefore, it is important to consider this aspect of data description.

In the case of wake/sleep data, where the physiological state transitions are far from being understood, we are able to extract the shape of the dynamical drift from wake to sleep in an unsupervised manner. By applying this new analysis tool, we hope to gain more insights into the underlying physiological processes. Our future work is therefore dedicated to a comprehensive analysis of large physiological datasets.

Acknowledgment

We acknowledge support of the DFG (grant JA379/51).

References

1. Kohlmorgen, J., Müller, K.-R., Pawelzik, K. (1995). Improving short-term prediction with competing experts. Proc. of ICANN'95, EC2 & Cie, Paris, 2:215–220.
2. Kohlmorgen, J., Müller, K.-R., Pawelzik, K. (1997). Segmentation and Identification of Drifting Dynamical Systems. Proceedings of NNSP'97: Workshop on Neural Networks for Signal Processing, IEEE, to appear.
3. Langhorst, P., Schulz, B., Schulz, G., Lambertz, M. (1983). Reticular formation of the lower brainstem. A common system for cardiorespiratory and somatomotor functions: discharge patterns of neighbouring neurons influenced by cardiovascular and respiratory afferents. J.A.N.S., 9: 411–432.
4. Moody, J., C. Darken (1988). Learning with Localized Receptive Fields. In *Proceedings of the 1988 Connectionist Models Summer School*, eds. D. Touretzky, G. Hinton, and T. Sejnowski, 133–143. San Mateo: Morgan Kaufmann.
5. Müller, K.-R., Kohlmorgen, J., Pawelzik, K. (1995). Analysis of Switching Dynamics with Competing Neural Networks, *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, E78-A, No.10, 1306–1315.
6. Müller, K.-R., Kohlmorgen, J., Rittweger, J., Pawelzik, K. (1995). Analysing Physiological Data from the Wake-Sleep State Transition with Competing Predictors, in NOLTA'95: Symposium on Nonlinear Theory and its Applications, 223–226.
7. Packard, N.H., Crutchfield J.P., Farmer, J.D., Shaw, R.S. (1980). Geometry from a Time Series. *Physical Review Letters*, 45:712–716.
8. Pawelzik, K., Kohlmorgen, J., Müller, K.-R. (1996). Annealed Competition of Experts for a Segmentation and Classification of Switching Dynamics, *Neural Computation*, **8:2**, 342–358.
9. Takens, F. (1981). Detecting Strange Attractors in Turbulence. In: Rand, D., Young, L.-S., (Eds.), *Dynamical Systems and Turbulence*, Springer Lecture Notes in Mathematics, **898**, 366.
10. Weigend, A.S., Gershenfeld, N.A. (Eds.) (1994). *Time Series Prediction: Forecasting the Future and Understanding the Past*, Addison-Wesley.

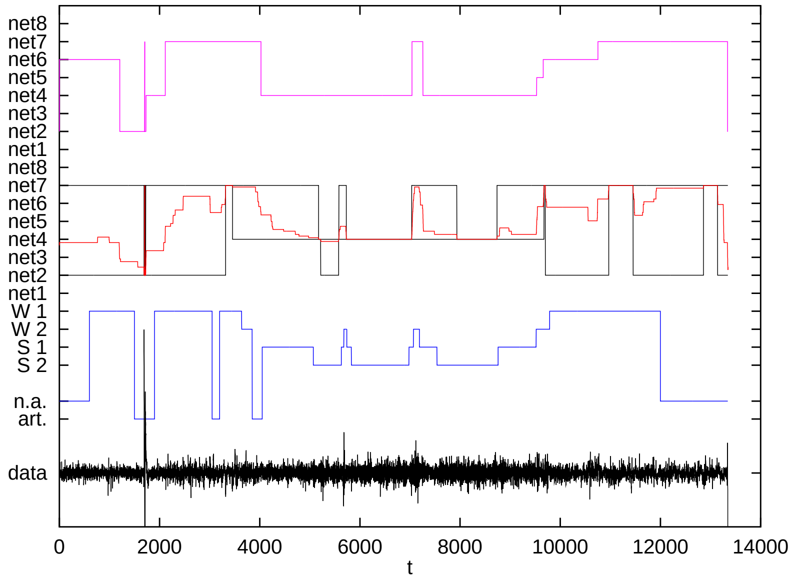

Fig. 1. *Comparison of hard segmentation (upper), drift segmentation (middle), and a manual segmentation by a medical expert (lower). Only a single-channel EEG recording (occipital-1,* 1400 *sec.) is given for the two algorithmic approaches.* $W1$ *and* $W2$ *indicate two wake-states (eyes open/closed) in the manual analysis,* $S1$ *and* $S2$ *indicate sleep stage I and II, respectively (*n.a.: *not considered,* art.: *artifacts). The dotted line in the drift segmentation indicates the evolution of the mixing coefficient* $a(t)$ *of the respective nets. For example, between* $t = 4000$ *and* 5000 *it denotes a drift from net* 7 *to net* 4*.*

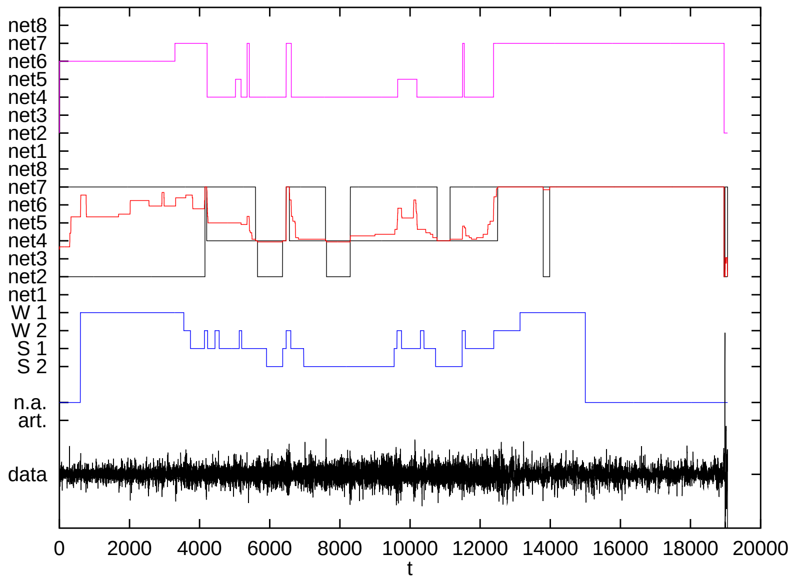

Fig. 2. *The segmentation plot for the test data: A second EEG recording (occipital-1,* 2000 *sec.) of the same person, analyzed by the ensemble already obtained in the first experiment. Compared to the manual segmentation, the EEG is properly segmented.*

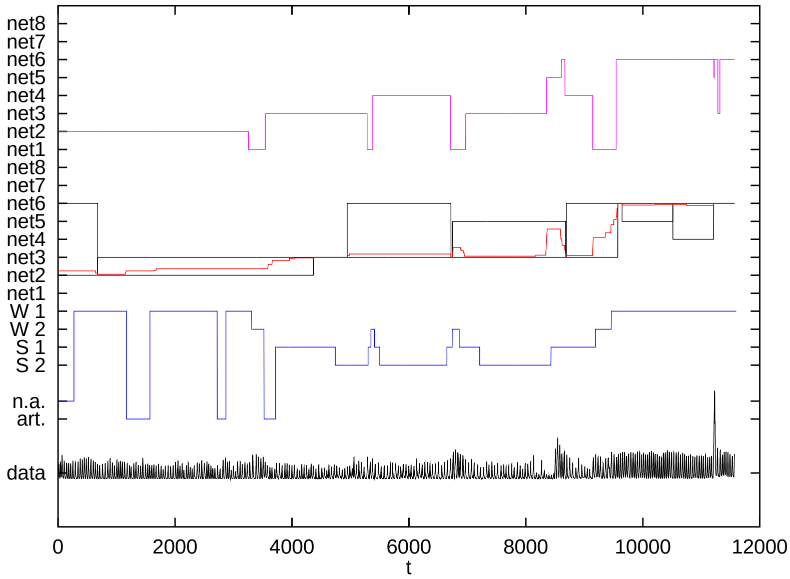

Fig. 3. *The respective segmentations of the respiration signal from the same experiment as the EEG analysis in Fig. 1. In this case, the hard segmentation (upper line) properly distinguishes between wake and sleep, whereas the drift segmentation (middle) does not indicate the intermediate arousals. However, the initial and final state transitions are modeled as drifts between prediction experts for wake and sleep.*

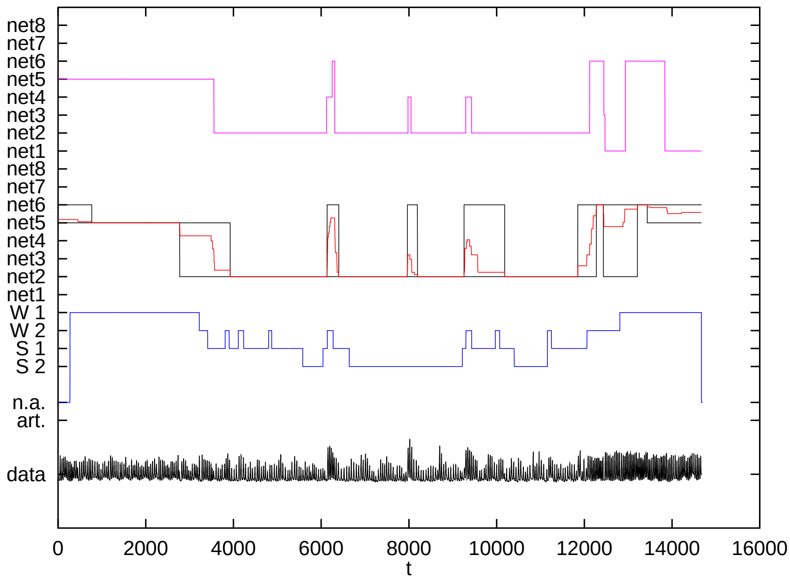

Fig. 4. *Analysis of the respiration signal from the experiment examined in Fig. 2. Here, the drift segmentation (middle) yields a similar result as the hard segmentation (upper). In both cases, the wake/sleep transitions are in good agreement with the manual segmentation (lower). Note that the short wake-up spikes in the manual segmentation were assessed with low confidence.*