Automated Selective REM Sleep Restriction Through Non-invasive Somatosensory Stimulation

Dillon Huffman[1],Asma'a Ajwad[1], Anuj Agarwal[2], Bruce O'Hara[2,3], Kevin Donohue[2,4], and Sridhar Sunderam[1]*

[1]F. Joseph Halcomb III, M.D. Department of Biomedical Engineering, University of Kentucky, Lexington, KY, [2]Signal Solutions, LLC, Lexington, KY

[3] Department of Biology, University of Kentucky, Lexington, KY, [4] Department of Electrical and Computer Engineering, University of Kentucky, Lexington, KY

* Address correspondence to ssu223@uky.edu

Introduction

The role of REM sleep in normal physiological function and health is widely studied, most commonly by selectively interrupting it when it occurs. Many protocols have been developed for experimental REM sleep restriction (RSR), but they are often stressful to the animal or alter its normal behavior. To address these limitations, we have developed an automated system that tracks sleep stage in real time and applies non-invasive vibrotactile stimulation to induce a state change. Stimulation can be tuned in frequency and amplitude to be more subtle or intense, making it more flexible than other systems. Here, we apply the system to the task of selective RSR, to assess its ability to yield an effect on REM sleep and explore the feasibility of producing graded effects by tuning the stimulation parameters (amplitude, frequency).

Experimental Methods

Six C57BL/6 mice (3m, 3f) were surgically implanted with EEG/EMG headmount electrodes. After recovery from surgery, animals were transferred to a cage where EEG, EMG, motion (piezo), and video were recorded. EEG/EMG signals were also decoded in real time by custom software, which classified signals as REM sleep, Non-REM sleep (NREM), or wakefulness (Wake) with 1-second resolution. During RSR, REM classification triggered vibro-tactile stimulation via a mechanical transducer under the floor of the cage. Each animal underwent 3 trials of RSR, each with a different combination of stimulation parameters. Following experiments, raw data was segmented into 4-second epochs and manually scored to assess the effects of RSR trials on sleep architecture.

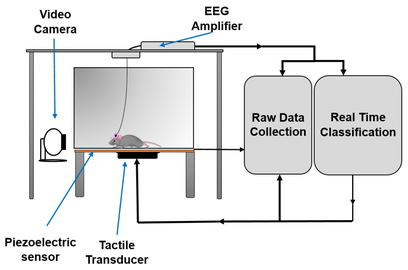

Figure 1: Experimental setup for sleep restriction experiments. EEG/EMG data were processed in real time by custom software, which stimulated the animal upon detection of REM via a tactile transducer mounted under the floor of the cage (*MouseQwake*, Signal Solutions, LLC). Annotations of REM detection and stimulation were stored in the data file.

Tracking Sleep in Real Time

In order to detect REM sleep in real time, features of the EEG/EMG data were computed at 1–second intervals for a 6-hour baseline recording, and used to build a Gaussian observation hidden Markov model (**Figure 2**). The model was then implemented for real-time use through a custom LabVIEW program, which reduced EEG/EMG data to a three-dimensional feature vector, that would be classified as either Wake, NREM, or REM, in a manner that closely approximates manual scoring (**Figure 3**).

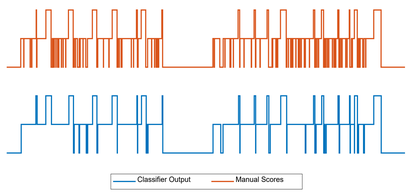

Figure 3: Comparison of classifier output and manual scores. While the model tends to underestimate sleep fragmentation (i.e. NREM to Wake transitions), it gives very reasonable approximations of REM sleep behavior.

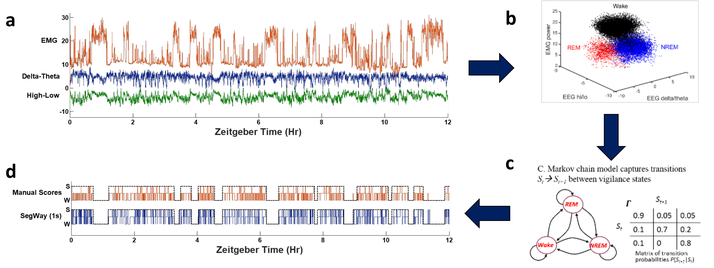

Figure 2: Process used to model sleep. A time series of EEG/EMG features are recorded (**a**), clustered (**b**), and modelled as a Markov chain (**c**). Data can then be passed through the model to obtain a time series of scores corresponding to sleep-wake states (**d**).

Efficacy of REM Sleep Restriction

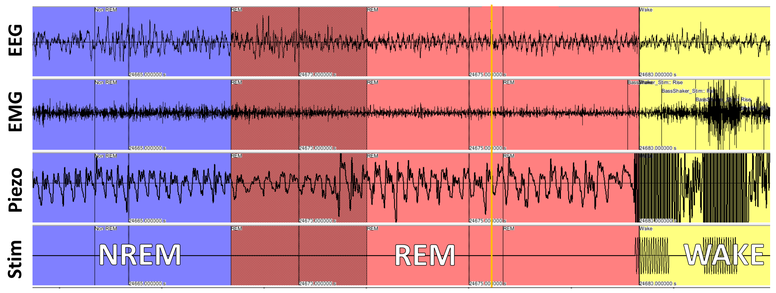

Figure 4: Example of REM sleep disruption. EEG/EMG data were processed in real time through custom LabVIEW software, which labeled incoming data as Wake, NREM, or REM sleep. Upon detecting REM, a sinusoidal waveform of pre-specified amplitude and frequency was generated to actuate the *MouseQwake* system to interrupt REM. A 5-second delay from detection to stimulation was programmed to avoid stimulating brief, probably erroneous REM classifications.

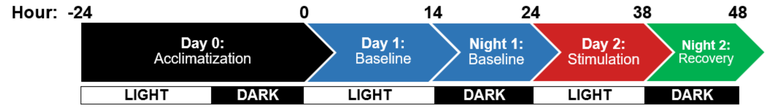

Figure 5: Timeline of RSR experiments. Each trial consisted of 48-hours, in which hours 0-14 served as a baseline for a time-locked RSR trial on day 2, allowing for direct comparison of sleep behavior to a recent baseline. In addition to the post-RSR recovery period, each animal was allowed to recover for an entire day before the next trial's baseline recording to prevent previous trials from affecting the outcomes.

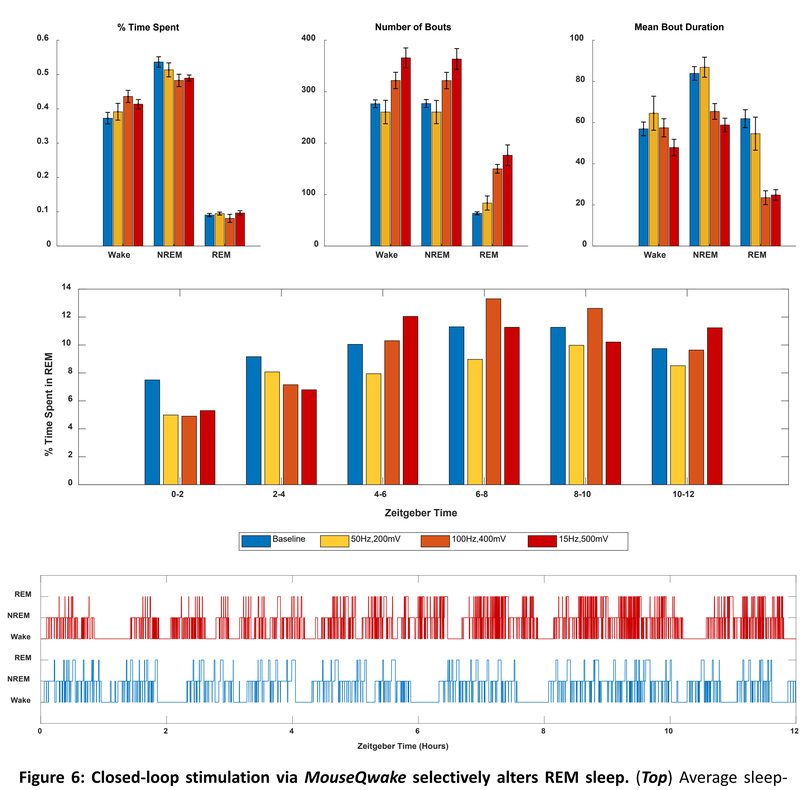

Figure 6: Closed-loop stimulation via *MouseQwake* selectively alters REM sleep. (***Top***) Average sleep-wake behavior over the stimulation period shows a decrease in mean REM bout duration with a corresponding increase in number of REM bouts that was proportional to stimulation intensity. However, homeostatic drive resulting from RSR yielded progressively more frequent transitions to REM sleep as a consequence, and less time spent in NREM sleep on the way (***bottom***), which translated to cumulative proportions of REM similar to that in the baseline. However, further dissection reveals that REM proportion was greatly reduced over the first 4 hours of the recording, before being later recovered – depending on stimulation parameters (***middle***).

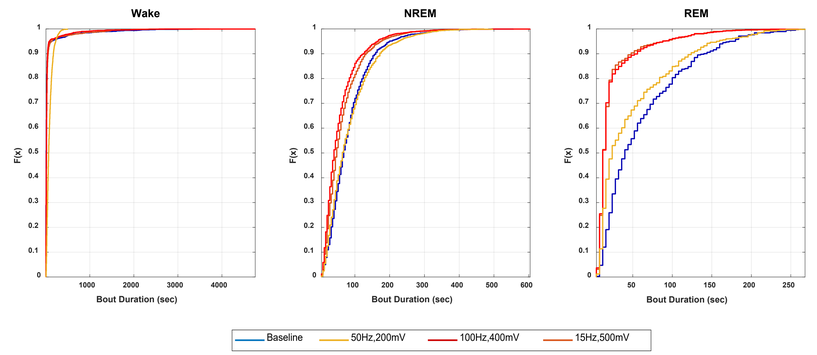

Figure 7: ECDFs of state bout durations resulting from each stimulation parameter set. RSR with high intensity parameters resulted in a dramatic decrease in the duration of REM bouts, with 90% of bouts being less than 50 seconds in duration. Lower intensity stimulation produced a more mild effect, but still diverged from that seen in the baseline.

Figure 8: Stimulation via MouseQwake causes sharp decline in REM probability. REM probabilities during RSR recordings were averaged across all REM onsets and compared to that seen in the baseline. REM detection latency (~10 sec), coupled with the programmed detection-to-stimulation delay (~5 sec) account for the time it takes for REM probability to diverge from the baseline trend (approximately 15 sec. after REM onset).

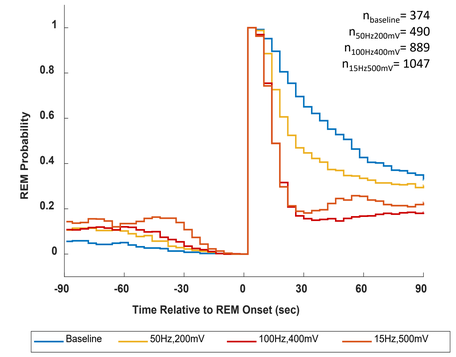

Conclusions and Future Directions

This initial investigation of *MouseQwake*'s use in selective sleep restriction showed promising results, drastically affecting the mean REM bout duration. However, after several hours of stimulation, it was difficult to combat the homeostatic drive using static, non-adaptive stimulation parameters. In light of this, future development will incorporate reinforcement learning for automatic selection and adaptation of stimulation parameters to overcome the circadian and homeostatic drives for REM sleep we have seen here. Furthermore, identification of sleep through the non-invasive piezoelectric motion sensor could alleviate the need for EEG implantation, resulting in a completely non-invasive system for sleep monitoring and perturbation.

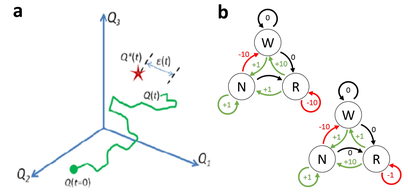

Figure 9: Application of reinforcement learning for stimulation adaptation. Stimulation parameters would be defined in a given space and the algorithm would search the space to learn which parameters work the best (***a***). The policy under which it would choose parameters can be tuned to favor either more intense (***b, left***) or subtle (***b, right***) stimulations.

Acknowledgment

This work was supported by National Institutes of Health grant NS083218.