10.3217/978-3-99161-093-9-019

Decoding Sleep and Wakefulness from Chronically Implanted Intracortical Signals in a Communication BCI User

Elena C. Offenberg[1], Zachary V. Freudenburg[1], Mariana P. Branco[1], Jonas B. Zimmermann[2], Janir R. da Cruz[2], Mariska J. Vansteensel[1], and Nick F. Ramsey[1,3]

[1] Brain Center, University Medical Center Utrecht, Utrecht, Netherlands
[2] Wyss Center, Geneva, Switzerland
[3] Donders Institute for Brain, Cognition and Behaviour, Radboud University, Nijmegen, Netherlands

E-mail: e.c.offenberg@umcutrecht.nl

ABSTRACT: As brain-computer interfaces (BCIs) move toward long-term home use, continuous operation across day and night introduces new challenges. Decoders trained during wakefulness may produce unintended activations during sleep due to large differences in neural activity between vigilance states. Users must be able to communicate at night, making reliable detection of sleep and wakefulness essential for safe and autonomous BCI operation.
We investigated whether awake-sleep stages can be inferred from chronically implanted intracortical neural signals in a communication BCI user with complete locked-in syndrome. Simultaneous non-invasive electroencephalography (EEG) and electrooculography (EOG) recordings provided reference sleep labels over a 44-hour period. A convolutional neural network was trained to distinguish between wakefulness and N2 sleep using intracortical data and achieved high accuracy (96.1±1.6%). The model was applied to six independent BCI spelling sessions, where more sleep-classified epochs coincided with reduced spelling performance. These results demonstrate the feasibility of decoding vigilance state from intracortical signals and support the development of state-aware BCIs for continuous home use.

INTRODUCTION

For implanted BCIs to be clinically viable for unsupervised home use, they must provide reliable and continuous operation over extended periods, ideally with 24/7 availability. This requirement is particularly critical for communication BCIs intended for individuals with severe motor impairments, such as those with amyotrophic lateral sclerosis or brainstem stroke, for whom the BCI may represent the primary or sole means of communication. For these individuals, the ability to signal caregivers during nighttime hours is essential for preserving autonomy and safety. At the same time, continuous operation introduces the risk of unintended activations when the user is asleep, or not intending to interact with the device, which may lead to false alarms, caregiver burden, and reduced trust in the system.
Robust identification of periods in which the user intends to engage with the BCI is thus a key requirement for safe and autonomous long-term deployment. One promising approach is the integration of an automated classifier of the user's vigilance state, specifically distinguishing between wakefulness and sleep. This can be used to infer whether BCI control is likely intended (during wakefulness) or clearly not intended and should be suppressed (during sleep), thereby reducing false positives and improving usability while maintaining continuous availability. Ideally, this classification should rely exclusively on neural signals already available from the implanted BCI system, avoiding the need for additional sensors of recording modalities that would complicate system design and limit clinical scalability.
Automatic detection of wakefulness, sleep, and sleep stages has been extensively studied using non-invasive electrophysiological recordings, including electroencephalography (EEG), electromyography (EMG), and electrooculography (EOG). Decades of work have established characteristic spectral and temporal features associated with vigilance, and numerous algorithms have achieved high accuracy in polysomnography-based sleep staging [1]. However, the BCIs most suitable for long-term at-home use have been based on chronically implanted electrodes [2, 3], which sample neural activity from a more spatially restricted cortical or subcortical region. The extent to which sleep-wake information can be reliably extracted from such localized intracranial recordings remains insufficiently explored.
Several studies have demonstrated the feasibility of discriminating wake and sleep states using temporary intracranial recordings, including electrocorticography [4], stereoelectroencephalography [5], and intracortical microwires [6], often reporting high classification performance. These investigations suggest that sleep-related neural signatures are detectable from implanted electrodes. However, most of this evidence derives from short-term clinical monitoring in epilepsy or research settings, from relatively distributed brain regions, rather than from spatially focused recordings from chronically implanted BCI users in naturalistic environments.
Observations in individuals with long-term implanted BCIs remain rare and reports on nocturnal BCI use, and

10.3217/978-3-99161-093-9-019

the effects of vigilance state highly scarce. In one report, a dedicated "night mode" was introduced to reduce unintended activations, relying on prolonged patterns of user behavior rather than direct neural markers of sleep [7]. Other work has identified neural correlates of drowsiness and linked them to fluctuations in BCI performance [3]. Despite these initial findings, systematic investigation of sleep-wake classification using chronic implanted BCI signals is still sparse, and the generalizability of existing results remains unclear.

MATERIALS AND METHODS

Participant and data acquisition: The data was collected from a participant with late-stage amyotrophic lateral sclerosis (ALS) who had been chronically implanted with two intracortical microelectrode arrays (8×8 electrodes per array, 1.5 mm electrode length, 0.4 mm inter-electrode pitch; Blackrock Microsystems LLC) in the left motor cortex in March 2019 (Figure 1). Neural signals were acquired using a digitizing headstage and a neural signal processor (CerePlex E and NSP, Blackrock Microsystems LLC) and sampled at 30kS/s. All procedures were approved by the Ethical Committee of the Medical Faculty of the Technische Universität München Rechts der Isar, and the participant and his legally responsible family members gave informed consent for the surgical procedure [8].

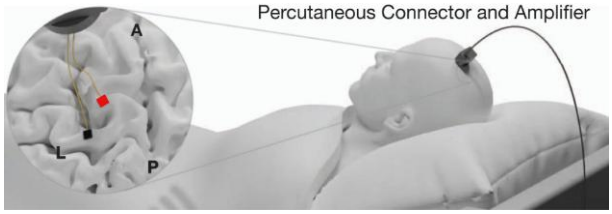

Figure 1: Experimental setup, adapted from [8]. Two microelectrode arrays were placed in the precentral gyrus and superior frontal gyrus. L marks the left central sulcus and A-P the midline from anterior to posterior. The Utah array used in this study, located in the Supplementary Motor Area (SMA), is marked in red.

A dedicated simultaneous recording session lasting over 44 hours (almost two days) was conducted, with the exact duration determined by setup constraints. During this session, non-invasive electrophysiological signals, including EOG and EEG data, were acquired along with the intracranial recordings. EOG data was recorded from two periocular electrodes positioned above the right eye and below the left eye, and EEG data was recorded from 12 scalp electrodes with a sampling rate of 500 Hz and using a 16 channel EEG amplifier (V-Amp DC, Brain Products, Germany) with gel-based Ag/AgCl electrodes. Electrode impedance was checked at the beginning of the recording to be below 5 kΩ. All channels were referenced to an electrode on the right mastoid and grounded to the electrode placed at the FPz location on the scalp.
The participant lacked reliable voluntary eye-movement control and was therefore unable to use eye-tracking-based communication systems, relying instead on an intracortical BCI spelling interface described previously [8]. Data used in the present study was recorded 459-461 days after implantation.

Sleep stage classification using non-invasive signals: EEG and EOG data recorded during the 44-hour period was preprocessed by band-pass filtering between 0.5 and 40 Hz and downsampling to 256 Hz to match the sampling rate used for preprocessed intracortical data. Outliers were removed using the interquartile range (IQR) criterion by clipping samples outside the range $[Q_1\text{-}5\cdot IQR, Q_3\text{+}5\cdot IQR]$.
Sleep staging was performed using the SleepyLand toolbox [9], which implements several state-of-the-art automated sleep stage classification algorithms using EEG and EOG data as input. Specifically, we applied DeepResNet [10], SleepTransformer [11] as well as USleep [12] to generate sleep stage predictions for non-overlapping 30-second epochs across the full recording period. Final sleep stage labels were obtained using an unweighted soft-voting ensemble of the three classifiers. The inclusion of both EEG and EOG signals within this multimodal ensemble implicitly accounted for ocular artifacts. These ensemble predictions served as reference labels for subsequent intracortical signal classification.

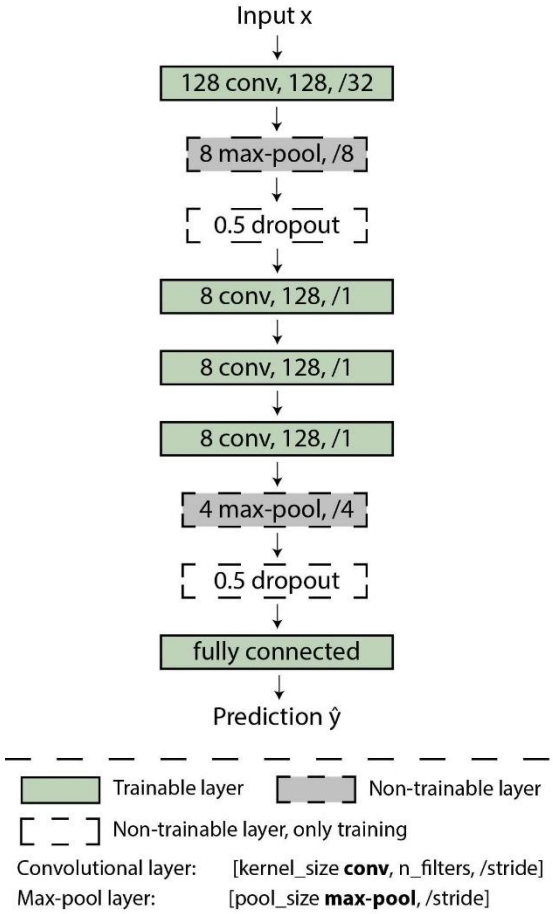

Figure 2: Architecture of the convolutional neural network used for awake-sleep classification. The model was adapted from TinySleepNet [13]. Rectangular blocks denote individual network layers, and arrows indicate the flow of data from 30-second neural signal epochs (x) to the binary sleep state prediction (ŷ).

10.3217/978-3-99161-093-9-019

Intracortical awake-sleep classification: Intracortical signals recorded from the Utah arrays were downsampled from 30 kHz to 256 Hz using polyphase filtering to prevent aliasing and notch-filtered at 50 Hz and its harmonics to remove line noise. Signals were averaged across electrodes from the array located in the SMA to obtain a single representative time series for classification. Averaging across electrodes was performed to capture global state-dependent changes in neural activity while reducing the influence of channel-specific artifacts and electrode instabilities, and to improve the signal-to-noise ratio of features relevant for vigilance state discrimination.

A convolutional neural network (CNN) was trained to discriminate between wakefulness and sleep (N2 sleep, see next paragraph) using 30-second intracortical signal epochs. The network architecture was based on the convolutional feature extraction component of TinySleepNet [13], excluding the recurrent neural network (RNN) module used for temporal sequence modeling. This architecture was chosen because it enables classification from individual 30-second epochs without relying on temporal context across multiple windows. The CNN comprised four convolutional layers interleaved with two max-pooling layers and two dropout layers to reduce dimensionality and mitigate overfitting. The complete architecture is shown in Figure 2 and was implemented using PyTorch [14].

For model training, all epochs labeled as N2 sleep by the SleepyLand ensemble classifier (n=743) were included, along with an equal number of randomly selected epochs labeled as awake to ensure class balance. The N2 sleep stage was selected because it was the most prevalent sleep stage in the dataset.

Each epoch was normalized by subtracting its mean and dividing by its standard deviation. Model performance was evaluated using 10-fold stratified cross-validation. The CNN was trained for 200 epochs using a cross-entropy loss function and the Adam optimizer with a learning rate of 0.001.

Frequency band contribution analysis: To assess the contribution of distinct frequency bands to awake-sleep classification, intracortical signals were decomposed into canonical frequency bands using bandpass filters following notch filtering at 50 Hz and its harmonics. For the spike band component (250-5000 Hz), bandpass filtering and power extraction were performed prior to downsampling by squaring the signal to obtain a slowly varying amplitude envelope and to prevent aliasing effects. The following bands were analyzed individually: delta (0.5-4 Hz), theta (4-8 Hz), alpha (8-13 Hz), beta (13-30 Hz), low gamma (30-50 Hz), high gamma (60-150 Hz) and spike band power (250-5000 Hz).

For each band-specific dataset, the convolutional neural network was retrained using the same preprocessing, training, and cross-validation procedures described above. Classification performance obtained using each isolated frequency band was then compared with performance obtained using the full broadband signal to quantify the relative contribution of each frequency range to awake-sleep discrimination.

Application to independent BCI spelling sessions: To assess generalizability, the trained classifier was applied to neural data from six independent BCI spelling sessions conducted in the afternoon and evening following the conclusion of the 44-hour continuous recording around noon. During each session, the participant was free to generate self-selected text and produced several sentences in both English and German using an independent spelling interface described in [8]. Spelling performance was evaluated by a fluent bilingual speaker, who assessed the percentage of output words that were comprehensible in either language.

Spelling session data was preprocessed identically to the training data, including downsampling, notch filtering, and normalization. The trained classifier was then applied to the first 20 minutes of each spelling session using sliding 30-second windows with a step size of 1 second. This fixed-duration segment was selected to enable direct comparison across sessions of unequal length. For each epoch, the model generated a binary prediction corresponding to either wakefulness or sleep (N2).

RESULTS

Sleep stage classification using non-invasive signals: Sleep stage predictions for each 30-second epoch are overlaid on the averaged EEG spectrogram across all 12

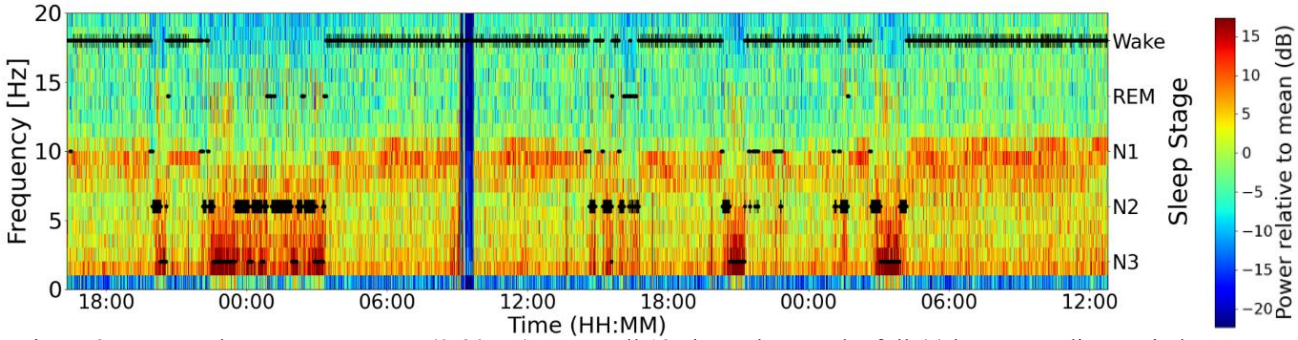

Figure 3: Averaged EEG spectrogram (0-20 Hz) across all 12 channels over the full 44-hour recording period. Predicted sleep stages by the SleepyLand classifier for each 30-second epoch are overlaid with black markers. Vertical lines indicate epochs used for training the intracortical classifier, consisting of all epochs labeled as N2 and an equal number of randomly selected epochs labeled as awake.

10.3217/978-3-99161-093-9-019

recorded channels in Figure 3. The resulting temporal and spectral distribution of vigilance states is consistent with known physiological patterns of human sleep. Periods classified as sleep occurred predominantly during nighttime hours, while wakefulness dominated daytime intervals, with the expectation of an afternoon nap on the second day of recording.
Furthermore, the sequence of predicted sleep stages followed the expected cyclic structure of human sleep architecture [15], providing qualitative validation of the automated sleep staging obtained from EEG and EOG signals.

Intracortical awake-sleep classification: Training and testing the convolutional neural network on the 44-hour intracortical recording yielded a mean classification score of 96.1±1.6% (Figure 4) of distinguishing wakefulness from N2 sleep, substantially exceeding chance level (practical chance level 48.7–51.3% [16]).

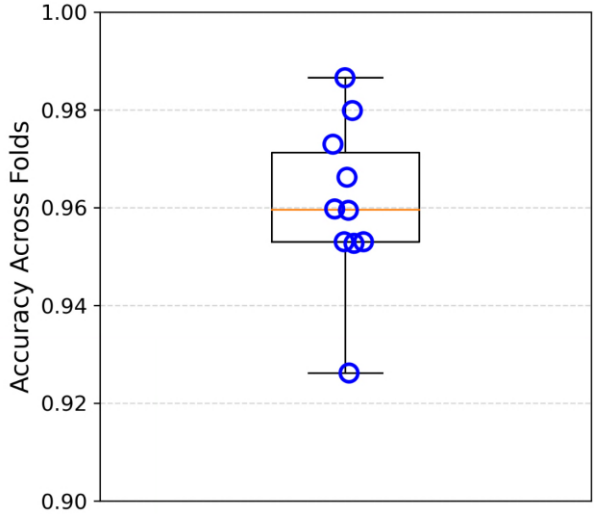

Figure 4: Classification accuracy across 10 cross-validation folds for intracortical awake-N2 sleep discrimination. Individual fold accuracies are overlaid on a boxplot showing the median and interquartile range (IQR); whiskers indicate the full range of observed values. Chance-level performance for this prediction is 48.7–51.3% [16].

Frequency band contribution analysis: Classification accuracies obtained when training the CNN on isolated frequency bands is shown in Figure 5. All individual frequency bands yielded significantly lower accuracy compared with the full broadband signal (Bonferroni corrected for number of bands, $p < 0.05$, Wilcoxon signed-rank test), indicating that no single frequency range alone was sufficient to achieve optimal discrimination between wakefulness and sleep. Additionally, only spike band power did not differ significantly from the practical chance-level range of 48.7–51.3% ($p = 0.77$ vs lower bound, $p = 0.49$ vs upper bound, Wilcoxon signed-rank test).

Application to independent BCI spelling sessions: The trained CNN was next applied to six independent BCI spelling sessions following the 44-hour recording period. For each session, overlapping 30-second epochs were extracted from the first 20 minutes of data, resulting in 1,172 epochs per session.

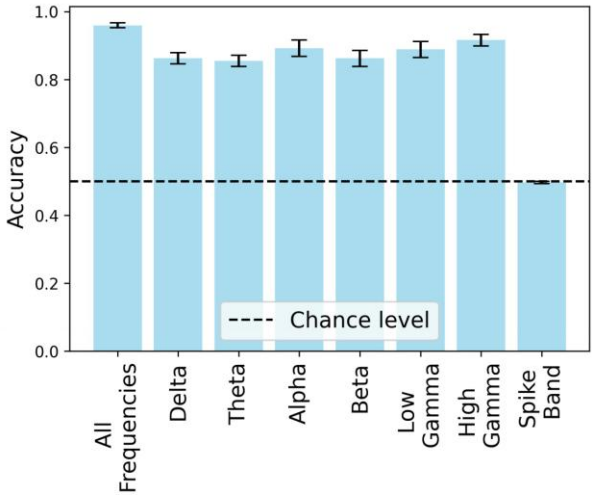

Figure 5: Classification accuracy for awake-sleep discrimination across cross-validation folds for individual frequency bands.

During the first two sessions (13:40 and 14:45), no epochs were classified as sleep by the model across any cross-validation fold. In the third session (16:02), only a single epoch was classified as sleep. The number of epochs predicted as sleep increased progressively in later sessions, reaching an average of 42 epochs in the final session at 22:20 (Figure 6).
Spelling performance declined in parallel with this increase in sleep-classified epochs. While 100% of words produced during the session at 20:37 were comprehensible, only 75% of produced words were comprehensible in the session at 21:51. In the final session, none of the produced words were comprehensible. This pattern suggests a relationship between reduced vigilance and impaired communication performance during prolonged BCI use.

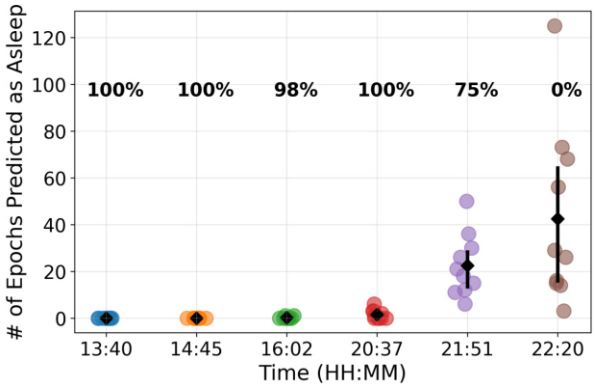

Figure 6: Number of epochs (out of 1,172 per session) classified as sleep by the trained CNN for each of six independent BCI spelling sessions recorded on the day following the 44-hour dataset. For each session, spelling performance is indicated as the percentage of comprehensible words produced during free spelling.

10.3217/978-3-99161-093-9-019

DISCUSSION

This study demonstrates the feasibility of detecting sleep and wakefulness from spatially focused, chronic neural signal recordings in a real-world BCI user, providing an important proof of concept for vigilance-aware autonomous BCI systems. Reliable identification of periods in which the user is asleep or drowsy is critical for reducing unintended activations, improving safety, and supporting continuous home use of communication BCIs. While the present findings are encouraging, substantial future work is required to extend and validate these findings.

The utilized reference sleep stage labels were derived from automated sleep staging of EEG and EOG signals without manual expert verification. While this approach enables scalable annotation, it may introduce some degree of label noise. Incorporating human scoring on a subset of the data could provide an estimate of labeling accuracy and help quantify the impact of potential misclassifications on model performance.

In this initial investigation, classification was restricted to discriminating between wakefulness and N2 sleep. We observed an increase in the number of epochs classified as sleep later in the day, which coincided with a decline in spelling performance. This binary framework could be expanded to include all sleep stages, enabling a more detailed characterization of vigilance state and potentially supporting adaptive BCI behavior across different levels of arousal. An alternative and complementary approach would be to estimate a continuous “drowsiness” or vigilance score rather than discrete sleep stages. Even during wakefulness, vigilance fluctuates considerably, and intermediate states of reduced alertness may already impair BCI performance. Such states may be detectable, as prior work has shown that increased low-frequency neural activity, including theta oscillations, is associated with fatigue and diminished task performance [17].

In the present study, non-invasive EEG and EOG recordings were used to generate reference sleep stage labels. While this approach enabled robust training and validation of the intracortical classifier, an important next step will be to identify sleep and vigilance signatures directly from intracortical data alone. Achieving reliable classification without simultaneous scalp recordings would simplify system design and facilitate long-term deployment in home environments. A related open question concerns the temporal stability of such classifiers: it remains unclear whether a model trained during a single calibration session would remain valid over weeks or months, or whether periodic recalibration using multimodal recordings would be required. Understanding this tradeoff will be essential for determining practical protocols for long-term clinical use.

A further limitation of this study is that it was conducted in a single participant. Future work must evaluate the generalizability of the proposed approach across individuals, electrode types, implantation sites, and BCI configurations. In addition, the participant had late-stage ALS, a condition known to alter sleep architecture, including reduced sleep efficiency and decreased proportions of N3 and REM sleep [18]. These disease-related factors may influence both the neural signatures of sleep and the performance of classification models. Validation in additional participants, including those with different etiologies and sleep profiles, will therefore be necessary.

Analysis of frequency band contributions indicated that the highest classification performance was obtained using alpha and gamma-band activity. Alpha-band activity encompasses sleep spindles, which are characteristic of N2 sleep, while gamma oscillations are prominent during active wakefulness and diminish as sleep onset approaches [19]. These findings are therefore consistent with established electrophysiological markers of vigilance state. In contrast, spike band power alone was insufficient to reliably discriminate between sleep and wakefulness in this dataset. This result may partly reflect limitations of the current network architecture, which includes max-pooling operations that may not optimally capture very high-frequency temporal patterns. In addition, spike band activity is known to vary substantially across individual electrodes [20], and averaging across channels may have attenuated informative features present at the single-channel level. Alternative architectures and channel-wise classification approaches may therefore be better suited to exploit spike band information.

Overall, this work provides an important first step toward vigilance-aware intracortical BCIs capable of continuous operation in naturalistic settings. By demonstrating that sleep-wake information can be extracted from chronically implanted electrodes and that these predictions relate meaningfully to communication performance, the results motivate further research into adaptive BCI systems that respond dynamically to user state. Future studies incorporating larger cohorts, longer timescales, and real-time implementation will be crucial for translating these findings into clinically robust and autonomous BCI technologies.

CONCLUSION

In this work, we investigated whether sleep and wakefulness can be inferred from chronically implanted intracortical BCI signals during naturalistic use. Using simultaneous non-invasive sleep staging as reference labels, we trained a convolutional neural network to distinguish between wakefulness and N2 sleep with high accuracy over a 44-hour recording period. We further demonstrated that the classifier generalized to independent BCI spelling sessions and that its predictions were associated with changes in communication performance. Together, these results establish the feasibility of intracortical vigilance state decoding and provide a foundation for future studies exploring state-aware BCI operation.

10.3217/978-3-99161-093-9-019

ACKNOWLEDGEMENTS

This publication is part of the project Dutch Brain Interface Initiative (DBI2) with project number 024.005.022 of the research programme Gravitation, which is financed by the Dutch Ministry of Education, Culture, and Science (OCW) via the Dutch Research Council (NWO). In addition, this project is funded by the European Union's HORIZON-EIC-2021-PATHFINDER CHALLENGES program under grant agreement No 101070939 (INTRECOM) and by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number 22.00198.

REFERENCES

[1] Gaiduk M, Serrano Alarcón Á, Seepold R, Martínez Madrid N. Current status and prospects of automatic sleep stages scoring: Review. Biomed. Eng. Lett. 2023;13(3):247–272

[2] Vansteensel MJ, Pels EGM, Bleichner MG, Branco MP Denison T, Freudenburg ZV et al. Fully Implanted Brain–Computer Interface in a Locked-In Patient with ALS. N. Engl. J. Med. 2016;375(21):2060–2066

[3] Card NS, Singer-Clark T, Peracha H, Iacobacci C, Hou X, Wairagkar M et al. Long-term independent use of an intracortical brain-computer interface for speech and cursor control. bioRxiv 2025.06.26.661591

[4] Pahwa M, Kusner M, Hacker CD, Bundy DT, Weinberger KQ, Leuthardt EC. Optimizing the Detection of Wakeful and Sleep-Like States for Future Electrocorticographic Brain Computer Interface Applications. PLOS ONE 2015;10(11):e0142947

[5] Von Ellenrieder N, Peter-Derex L, Gotman J, Frauscher B. SleepSEEG: automatic sleep scoring using intracranial EEG recordings only. J. Neural Eng. 2022;19(2):026057

[6] Lamsam L, Gu B, Liang M, Sun G, Khan KJ, Sheth KN et al. The human claustrum tracks slow waves during sleep. Nat. Commun. 2024;15(1):8964

[7] Leinders S, Arnoutse EJ, Brainco MP, Freudenburg ZV, Geukes SH, Schippers A et al. DO NOT LOSE SLEEP OVER IT: IMPLANTED BRAIN-COMPUTER INTERFACE FUNCTIONALITY DURING NIGHTTIME IN LATE-STAGE AMYOTROPHIC LATERAL SCLEROSIS. medRxiv 2024.10.11.24315027

[8] Chaudhary U, Vlachos I, Zimmermann JB, Espinosa A, Tonin A, Jaramillo-Gonzalez A et al. Spelling interface using intracortical signals in a completely locked-in patient enabled via auditory neurofeedback training. Nat. Commun. 2922;13(1):1236

[9] Rossi AD, Metaldi M, Bechny M, Filchenko I, van der Meer J, Schmidt MH et al. SLEEPYLAND: trust begins with fair evaluation of automatic sleep staging models. Npj Digit. Med. 2025;9(1):55

[10] Olesen AN, Jennum PJ, Mignot E, and Sorensen HBD. Automatic sleep stage classification with deep residual networks in a mixed-cohort setting. Sleep 2021;44(1):zsaa161

[11] Phan H, Mikkelsen K, Chen OY, Koch P, Mertins A, De Vos M. SleepTransformer: Automatic Sleep Staging With Interpretability and Uncertainty Quantification. IEEE Trans. Biomed. Eng. 2022;69(8):2456–2467

[12] Perslev M, Darkner S, Kempfner L, Nikolic M, Jennum PJ, Igel C. U-Sleep: resilient high-frequency sleep staging. Npj Digit. Med. 2021;4(1):72

[13] Supratak A, Guo Y. TinySleepNet: An Efficient Deep Learning Model for Sleep Stage Scoring based on Raw Single-Channel EEG, in Proc. 2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), Montreal, QC, Canada, 2020, 641–644

[14] Paszke A, Gross S, and Massa F. PyTorch: an imperative style, high-performance deep learning library, in Proc. 33rd Int. Conf. Neural Inf. Process. Syst., Vancouver, Canada, 2019, 8026–8037

[15] Feinberg I, Floyd TC. Systematic Trends Across the Night in Human Sleep Cycles. Psychophysiology 1979;16(3):283–291

[16] Mueller-Putz G, Scherer R, Brunner C, Leeb R, Pfurtscheller G. Better than random? A closer look on BCI results. International Journal of Bioelectromagnetism 2008;10:52-55.

[17] Hung CS, Sarasso S, Ferrarelli F, Riedner B, Ghilardi MF, Cirelli C et al. Local Experience-Dependent Changes in the Wake EEG after Prolonged Wakefulness. Sleep 2013;36(1):59–72

[18] Boentert M. Sleep disturbances in patients with amyotrophic lateral sclerosis: current perspectives. Nat. Sci. Sleep 2019;11:97–111

[19] Adamantidis AR, Gutierrez Herrera C, Gent TC. Oscillating circuitries in the sleeping brain. Nat. Rev. Neurosci. 2019;20(12):746–762

[20] Nason SR, Vaskov AK, Willsey MS, Welle EJ, An H, Vu PP, et al. A low-power band of neuronal spiking activity dominated by local single units improves the performance of brain–machine interfaces. Nat. Biomed. Eng. 2020;4(10):973–983