Predicting Cognitive Impairment from Polysomnograms with Block-Recurrent Transformers

Jad Haidamous*, Christoph Hoog Antink

KIS*MED - AI Systems in Medicine,
Technical University of Darmstadt, Darmstadt, Germany

Aims: Certain cognitive impairments are associated with sleep disorders that can be captured in polysomnography (PSG) recordings. However, PSG signals are high-dimensional, multimodal, and span several hours, making it challenging to efficiently model both local patterns and long-term dependencies. Furthermore, PSG signals are often incomplete due to missing modalities and measurement artifacts. Therefore, this work aimed to develop a robust and scalable deep learning approach to predict cognitive impairment from long-term incomplete PSG signals.

Methods: Our model consists of a multimodal transformer network with modality-specific embeddings. The model was pretrained on the 2026 George B. Moody PhysioNet Challenge dataset with 10 minute windows. The loss combines the DINOv3 masked self-supervised objective and an auxiliary classification loss to classify cognitive impairment, sleep stages, limb movements, and arousal and respiratory events. The pretraining encouraged the model to learn cross-modal and temporal features that are robust to missing signals. To model long-term dependencies, we extend the architecture using block-recurrent transformers. The PSG signal is split into consecutive 10 minute windows, and latent state tokens are propagated across windows to summarize past information, effectively serving as memory. The block-recurrent transformer was trained exclusively on the classification tasks. The gradient was propagated across multiple windows, allowing the model to learn long-term dependencies.

Results: The non-recurrent pretrained model achieved a score of 0.437 on the hidden test set in the unofficial phase. The block-recurrent extension still needs to be implemented.

Conclusion: We tackled the challenge by training a transformer network with the DINOv3 masked self-supervised objective and an auxiliary classification loss. The approach scored poorly in the unofficial phase with a challenge score of 0.437 (rank $213/243$, team *Inception vKISMED*), leaving room for improvement. By incorporating the proposed extensions, we aim to improve the model's performance on long-term recordings.