Heart Rate or Brain Waves? Sleep Signatures for Cognitive Impairment Detection

Marion Taconné[1,*], Stefano Magni[1], Cristian Drudi[1], Filippo Emanuele Colella[1,2], Anshum Patel[3], Joseph Cheung[3], Valentina DA Corino[1], Anna Maria Maddalena Bianchi[1], Pietro Cerveri[1], Riccardo Barbieri[1], Luca Mainardi[1] (Team name: Mind the Nap)

1. Department of Electronics, Information and Bioengineering, Politecnico di Milano, Milan, Italy

Context: Polysomnography (PSG) offers a promising alternative to neuropsychological testing for cognitive impairment screening, as sleep architecture is known to degrade with cognitive decline. While EEG has been extensively studied as a marker of brain health, ECG remains largely unexplored despite autonomic dysfunction being an established feature of neurodegeneration. In the context of the PhysioNet Challenge 2026, we investigate whether cardiac and cerebral PSG features can discriminate cognitive impairment.

Methods: We developed two feature-engineering pipelines using 622 labelled patients, both combined with 10 demographic features and classified using XGBoost, evaluated with 10-fold cross-validation. The first pipeline extracted 197 stage-aware ECG features: 18 HRV metrics computed separately for each CAISR-annotated sleep stage (N3, N2, N1, REM, Wake), and their cross-stage aggregations. The second pipeline augmented the ECG features with 156 stage-aware EEG features: per-stage spectral band powers, Hjorth parameters, spectral entropy, spectral edge frequency, spindle proxy, and 6 clinically motivated cross-stage ratios (e.g., N3 delta/REM theta, wake theta/alpha slowing index).

Results:

Submission	CV AUC	LB AUC
ECG	0.532 ± 0.044	-
EEG	0.587 ± 0.060	-
ECG + EEG	0.579 ± 0.055	-
Demo + ECG	0.822 ± 0.052	0.633
Demo + ECG + EEG	0.803 ± 0.021	0.590

Table 1: 10-fold cross-validation (CV) and leaderboard (LB) AUC for both submissions.

Conclusions: Both submissions show a gap between internal and leaderboard performance, suggesting overfitting or distribution shift across datasets. Notably, adding EEG features did not improve generalisation over ECG alone, likely due to the high dimensionality of the EEG feature space relative to the cohort size. Future work will focus on pre-trained deep learning encoders, in particular transformer-based EEG and ECG foundation models, to extract more robust and generalisable representations, alongside stronger regularisation and site-aware validation strategies.