

Abstract

Electroencephalogram (EEG) has become widely used across multiple application domains, including clinical neurology and brain-computer interaction. It provides a non-invasive and cost-effective means to monitor brain activity with high temporal resolution. Despite its wide use, automated EEG analysis remains difficult due to non-stationarity, joint time–frequency characteristics, and strong variability across subjects. Existing methods lack in addressing long temporal context, multi-view representation learning, and cross-subject adaptability. This thesis addresses three core problems shaped by these gaps: (i) multi-class harmful brain activity classification in critical care, (ii) sleep stage classification under subject and demographic variability, and (iii) motor imagery classification with limited labeled data. The objective is to design representation learning frameworks that improve discriminability, generalisation, and adaptability for clinically relevant EEG tasks.

For harmful brain activity, three frameworks have been developed. The first, *WaveEEG* models long-range temporal context through a convolutional–transformer pipeline (92.6%, $\kappa = 0.91$). The second, *BiSTFormer* factorises temporal and spectral streams with learnable gated fusion (94.2%, $\kappa = 0.93$). The third, *BiVISTA* fuses multi-scale waveforms, spectrograms, and montage-grouped region spectrograms in a clinically aligned design (96.0%, $\kappa = 0.952$).

For sleep stage classification, *XSleepFormer* (compact CNN–Transformer) and *NeuroWhisper* (speech-encoder transfer) are analyzed. NeuroWhisper reaches 86.4% on Sleep-EDF. Age-stratified subject-exclusive evaluation exposes an accuracy gap of about nine percentage points between the 25–65 and 80+ cohorts, exceeding the 2.5–3 percentage-point gain from pretraining.

For motor imagery, the *Meta-Adaptive Hilbert Framework* (MAHF) combines model-agnostic meta-learning, a trainable Hilbert layer for instantaneous spectral features, and scaling-shifting modules that limit catastrophic forgetting. Under the leave-one-subject-out evaluation on BCI Competition IV-2a, MAHF reaches $77.25\% \pm 3.25\%$, against 74.10% (RA-MDRM) and 72.50% (MUPS-EEG), retaining 79.03% on previously seen subjects.

The work treats EEG classification as a problem of representation design: long-range temporal context, spectro-temporal factorization, clinically aligned multi-view fusion, age-aware evaluation, and meta-learning with adaptive spectral filtering addresses a specific property of the EEG that limits classification. All results are on standard public benchmarks; external validation on independent clinical cohorts is a required next step.