

Automatic Sleep Arousal Identification from Physiological Waveforms using Deep Learning

Daniel R Miller, Andrew T Ward

Stanford University
Stanford, CA, United States

Aims: Feature extraction is one of the most time-intensive requirements in traditional machine learning. In most cases, extracting useful predictive features from dense data requires a great deal of domain expertise and experimentation. In contrast, deep learning techniques seek to avoid this restriction by automatically learning variable interactions such as those between pairs or groups of signals, and any relevant temporal dependencies. We seek to automatically extract sleep patterns from rich physiological time series using supervised deep learning algorithms.

Methods: 13 channels of annotated sleep waveforms were used as inputs to the model, including electroencephalography (EEG), electrooculography (EOG), electromyography (EMG), electrocardiology (EKG), and oxygen saturation (SaO₂). These waveforms were fed directly into a Convolutional Neural Network (CNN), which was trained to predict the probability of arousal at each time step. CNNs allow us to build a flexible model that may be applied to variable-length time series, automatically capturing interactions between temporally correlated physiological signals in the convolutional filter weights. Our model consisted of leading convolutional layers to capture cross-channel time dependencies, followed by a multi-layer perceptron output classifier. This model was trained on a classs-frequency-weighted binary cross-entropy loss function over the output probabilities, using an Adam optimizer.

Results: Our baseline model achieved an average training AUC of 0.70 on an 80% split of the data (795 samples), and an average validation AUC of 0.69 on the remaining 20% cross-validation split (199 samples). Our baseline model for the preliminary round achieved a test AUROC of 0.488, and a AUPRC of 0.051.

Conclusions: Although deep learning models require significant training and experimentation, they allow for the automatic extraction of complex interactions, alleviating the necessity of custom feature design.