

# Retrieving Filter Spectra in CNN for Explainable Sleep Stage Classification

S. Goerttler<sup>1,2</sup>, Y. Wang<sup>2</sup>, F. He<sup>1</sup> and M. Wu<sup>2</sup>

1. Centre for Computational Science and Mathematical Modelling, Coventry University, Coventry, UK

2. Institute for Infocomm Research, A\*STAR, Singapore

goerttlers@uni.coventry.ac.uk

**Abstract**— Despite significant advances in deep learning-based sleep stage classification, the clinical adoption of automatic classification models remains slow. One key challenge is the lack of explainability, as many models function as black boxes with millions of parameters. In response, recent work has increasingly focussed on enhancing model explainability. This study contributes to these efforts by introducing an explainability tool for spectral processing of individual EEG channels. Specifically, this tool retrieves the filter spectrum of low-level convolutional feature extraction and compares it with the classification-relevant spectral information in the data. We apply our tool on the EEGNet and MSA-CNN models using the ISRUC-S3 and Sleep-EDF-20 datasets. The tool reveals that spectral processing plays a significant role in the lower frequency bands. In addition, comparing the correlation between filter spectrum and data-derived spectral information with univariate performance indicates that the model naturally prioritises the most informative channels in a multimodal setting. We specify how these insights can be leveraged to enhance model performance. The code for the filter spectrum retrieval and its analysis is available at <https://github.com/sgoerttler/MSA-CNN>.

**Keywords**— *electroencephalogram, sleep stage classification, deep learning, explainable AI.*

## I. INTRODUCTION

Sleep stage classification is essential for assessing sleep and diagnosing sleep disorders using polysomnography [1]. With advances in machine learning, deep learning-based sleep stage classification has garnered significant attention [2]. However, despite promising results, its clinical adoption remains limited due to challenges in validation, professional integration, and explainability [3].

Explainability is particularly crucial for regulatory compliance and building trust among clinicians and patients while advancing model development [4]. As a consequence, explainability of sleep stage classification models has gained significant traction in recent years. For instance, gradient-weighted class activation mapping has been used for sample-specific post-hoc explanation of input importance [5, 6], as well as for highlighting weighted attention [7]. Furthermore, self-attention has been utilised to explain input importance [8] and even temporal interactions [9] in individual samples. While these approaches demonstrate sample-specific explanations, model explainability can also be global at the dataset level. For example, surrogate ablation

techniques have been used to assess the contribution of individual modalities globally [10]. Note that despite this abundance of available explainability techniques, each method generally captures only specific aspects of model behaviour and is often dependent on the model architecture.

This work aims to extend existing efforts by introducing a method for explaining spectral information processing globally in a multivariate context. Our approach can be applied to models that leverage temporal convolutions to extract spectral and morphological features from the raw signal, which are common in EEG signal classification [11]. In particular, our method retrieves the frequency spectrum of filters in the first convolutional layer. This layer may comprise several pathways to cover the multiple spectral scales relevant in neurophysiology [9, 12], in which case we employ a unification assignment matrix to integrate these pathways. Further, we correlate the extracted filter spectrum with the classification-relevant spectral information in the data, which provides insight into the extent of spectral information processing. Lastly, we compare these correlations with the single-channel performance.

We evaluate the proposed method on the EEGNet introduced by Lawhern *et al.* [13] as well as our previously introduced Multi-Scale and Attention Convolutional Neural Network (MSA-CNN), using two datasets. Both models are multivariate, CNN-based, and designed to capture a broad spectral range from 0 Hz to around 50 Hz.

## II. METHODOLOGY

### II-A. Method Overview

The principal idea of our proposed method is to explain channel-wise spectral processing within the first convolutional feature extraction layer. This layer can capture both spectral features (e.g., EEG oscillations) and morphological features (e.g., waveform shape). While this layer typically operates on a single scale, which is determined by the sampling rate, the layer may also operate on multiple scales, such as in the case of the MSA-CNN model. The filters of these convolutions capture distinct frequencies, which are determined by their Fourier transform and their scale. For multi-scale configurations, frequencies at different scales may overlap and thus require unification. The

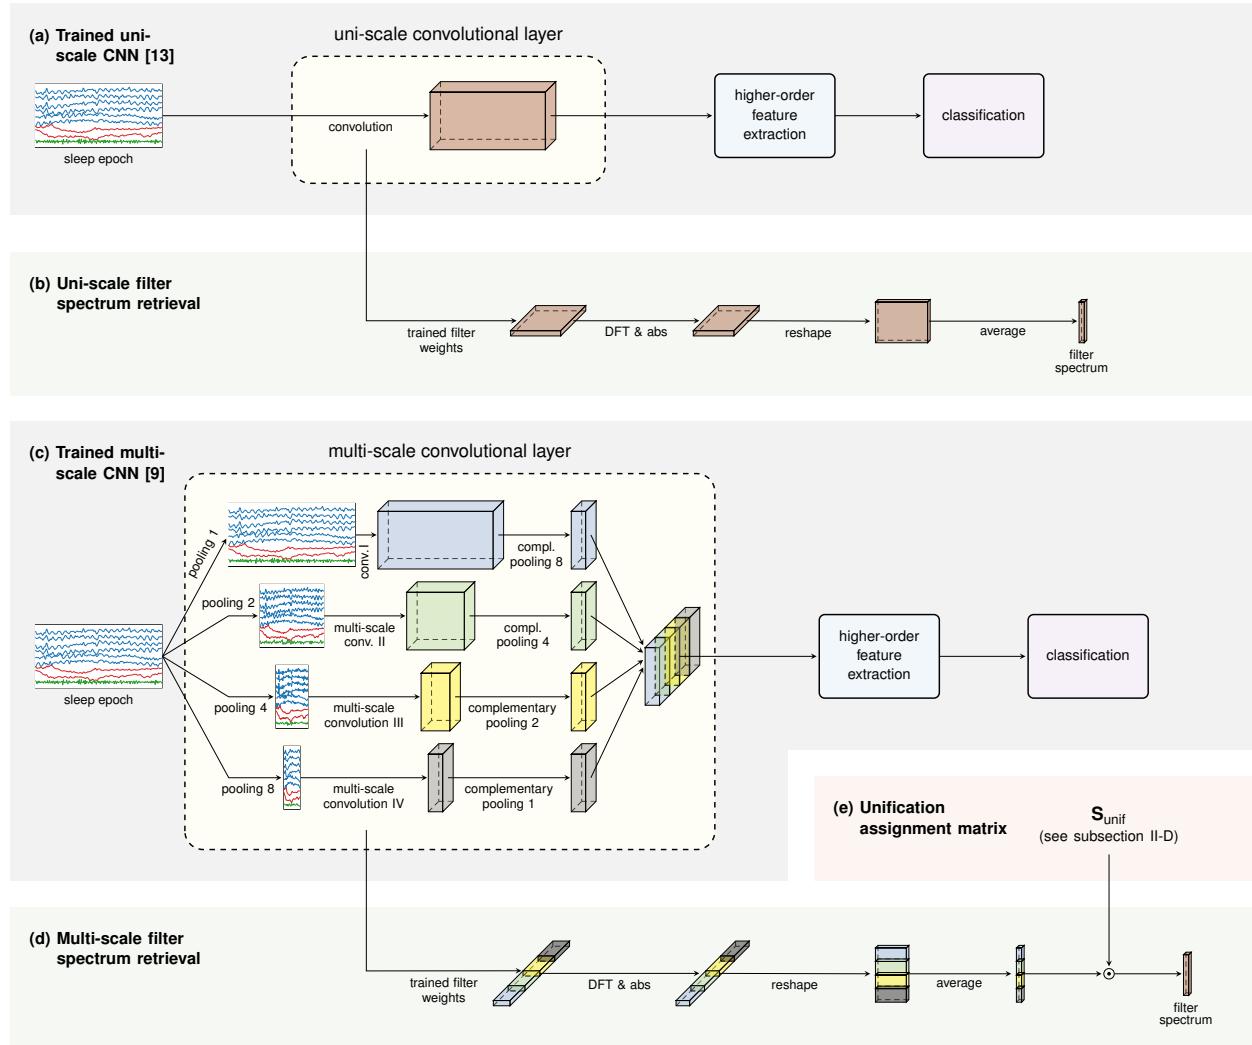


Figure 1. Overview of proposed filter spectrum retrieval model for uni- (a–b) and multi-scale (c–e) configurations. (a) Trained uni-scale CNN model, such as the EEGNet [13]. The first convolutional layer extracts temporal features from the raw signals, followed by higher-order modules. (b) Proposed uni-scale filter spectrum retrieval using the trained convolutional filter weights. (c) Trained multi-scale CNN model, such as the MSA-CNN [9]. The first layer comprises temporal convolutions on multiple scales, which are aggregated before being passed to the higher-order modules. (d) The multi-scale filter spectrum retrieval conducts the steps in b separately for each scale before the scales are averaged using the unification assignment matrix. (e) The unification assignment matrix  $S_{\text{unif}}$  transforms scale-based frequency to unified frequency.

amplitude values of the Fourier components are then averaged across filters to yield the final amplitude spectrum. This spectrum is then compared with the between-class spectral variation in the data, which serves as an indicator of classification-relevant spectral information. An overview of our proposed method for uni-scale (a–b) and multi-scale (c–e) configurations is presented in Figure 1.

#### II-B. EEGNet Architecture

The EEGNet is a lightweight CNN architecture introduced by Lawhern *et al.* [13]. The model comprises three convolutional layers arranged in two blocks as well as a classification layer. The first block consists of

a temporal convolutional layer and a depthwise spatial convolutional layer. The second block consists of a separable convolution, which comprises a depthwise temporal convolution and a point-wise convolution. Each convolutional layer is followed by batch normalisation, while each block ends with a sequence of exponential linear unit, average pooling and dropout. The classification layer consists of a fully connected layer and a softmax operation. Note that the spatial layer in the first block and the second block correspond to the higher-order feature extraction block in Figure 1(a). For consistency, the current study uses the same model parameters as the original study.

### II-C. Multi-Scale and Attention Convolutional Neural Network (MSA-CNN)

The MSA-CNN [9] is shown in Figure 1(c). To begin with, a multi-scale convolution and a scale-integration convolution extract spectral features for each channel. Subsequently, the channels are combined using a global spatial convolution. The time-resolved features are then passed to an attention-based Temporal Context Module (TCM), which captures long-range temporal dependencies. The TCM combines multi-head attention with a feed-forward neural network. Finally, the resulting features are averaged over time and classified using a fully connected layer.

### II-D. Unification Assignment Matrix

In many deep learning-based neurophysiology models, the first convolutional layer comprises convolutions on more than one scale [9, 12, 14–16]. This requires the use of a unification assignment matrix  $\mathbf{S}_{\text{unif}}$  for our explainability tool, which assigns the scale-based frequencies to unified frequencies on a single scale.

The scale-based frequencies are collected in the list  $\mathbf{f}_{\text{sc}} = (f_0^{(I)}, f_1^{(I)}, \dots, f_0^{(II)}, \dots, f_{\lfloor N_t/2 \rfloor}^{(IV)})$ , where the superscript denotes the scale and the subscript indexes the frequency within that scale. We determine the list of unique frequencies  $\mathbf{f}_{\text{uniq}}$  by sorting  $\mathbf{f}_{\text{sc}}$  by magnitude and removing duplicate frequencies. The unnormalised unification assignment matrix is then given by:

$$\widehat{\mathbf{S}}_{\text{unif},ij} = \begin{cases} 1, & \text{if } \mathbf{f}_{\text{uniq}}[i] = \mathbf{f}_{\text{sc}}[j], \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

The matrix is then normalised using the diagonal matrix  $\mathbf{D}_{ii} = \sum_j \widehat{\mathbf{S}}_{\text{unif},ij}$ . This yields the full frequency unification matrix  $\mathbf{S}_{\text{unif}} = \mathbf{D}^{-1} \widehat{\mathbf{S}}_{\text{unif}}$ , which transforms the scale-based frequencies into their unified representation as

$$\mathbf{f}_{\text{unif}} = \mathbf{S}_{\text{unif}} \mathbf{f}_{\text{sc}}. \quad (2)$$

### II-E. Filter Spectrum Retrieval

An overview of the uni- and multi-scale filter spectrum retrieval is given in Figure 1(b) and 1(d), respectively. The filter spectrum is retrieved from the filter weights  $\mathbf{W}$  of the trained model. In the case of unimodal filters, where filters are shared across all neurophysiological channels, the weight matrix has shape  $K \times L \times S$ , with  $K$  denoting the number of output channels,  $L$  the length of the filter, and  $S$  the number of scales. Subsequently, the discrete Fourier transform (DFT) is applied to the filter weights along the second, temporal dimension. Only the magnitudes of all positive frequencies are kept for further processing. The resulting feature maps are averaged across the  $K$  output channels, producing the filter spectrum in terms of frequency and scale. For multi-scale configurations, the scales are unified by

multiplying the filter spectrum by the assignment matrix  $\mathbf{S}_{\text{unif}}$ .

### II-F. Between-Class Spectral Variation

This section describes a measure that assesses the spectral information relevant for classification for each channel. To this end, we introduce the between-class spectral variation, which quantifies spectral density variation between class means normalised by within-class variation. Specifically, we compute the spectral density as the square root of Welch's power spectral density for every time series. The between-class variation is then retrieved by firstly averaging the spectral density across all samples for each class, and subsequently computing the standard deviation between the classes. On the other hand, the within-class variation is retrieved by computing the standard deviation for each class before averaging across classes. Lastly, we divide the between-class variation by the within-class variation to obtain our final measure of relevant spectral information.

## III. EXPERIMENTS

### III-A. Datasets

We use the publicly available datasets ISRUC-S3 and Sleep-EDF-20 in this study. The **ISRUC-S3** was recorded by Khalighi et al. from 10 healthy subjects during sleep [17]. The recordings were divided into 30-second epochs and labelled, resulting in a total of 8,589 annotated samples. For our experiment, we only use channels with sufficient spectral structure, which includes the six referenced electroencephalography (EEG) channels, the two electrooculography (EOG) channels, and the electromyogram (EMG) channel. The input signals are downsampled to 100 Hz and preprocessed with a fourth-order low-pass Butterworth filter at 40 Hz cutoff frequency.

The **Sleep-EDF-20** dataset is sourced from PhysioBank [18]. The dataset comprises 20 subjects and has overall 42,308 labelled, artefact-free 30-second samples with a sampling rate of 100 Hz. Channels with spectral structure include two referenced EEG channels (Fpz-Cz and Pz-Oz) and the EOG channel. Similar to the ISRUC-S3 dataset, we preprocessed the data with a 40 Hz low-pass filter.

### III-B. Model Specifications

We trained multivariate models for generating the weight matrices for the filter spectrum retrieval, as well as univariate models to assess the single channel performance. Multivariate models were trained on the entire dataset, while univariate models were trained and evaluated using 10-fold subject-wise cross-validation, with performance measured as mean accuracy.

We adopted the parameter settings for EEGNet as described in the original publication [13], with minor

Table 1. Model parameters for MSA-CNN small and MSA-CNN large, based on multivariate or univariate inputs. Kernel sizes are specified by their spatial and temporal dimensions

layer	hyperparameter	MSA-CNN (small)		MSA-CNN (large)	
		multivar.	univar.	multivar.	univar.
MSM I	# scales		4		4
	filter size		$1 \times 15$		$1 \times 15$
MSM II	# filters / scale		8		8
	filter size		$1 \times 5$		$1 \times 5$
spatial	# filters		16		32
	filter size	$N_{ch} \times 1$	$1 \times 5$	$N_{ch} \times 1$	$1 \times 5$
TCM	# filters		32		64
	embedding dim.		16		32
	# heads		2		4
	# layers		1		2

modifications. Specifically, we reduced the filter size from 64 to 50 due to the lower sampling rates in this study, such that the detectable spectral resolution of 2 Hz matches that in the original study. The same model parameters were used for both univariate and multivariate inputs across both datasets. To train the models, we followed the procedure outlined in the original publication.

The model parameters for the MSA-CNN configurations are listed in Table 1. Both temporal convolutional layers are set to unimodal, meaning that filters are shared across channels. We tailored the model size to each dataset, using a small configuration for ISRUC-S3 and a large configuration for Sleep-EDF-20. In addition, for the univariate configuration, we replaced the spatial convolution in the third layer with a temporal convolution. All models were trained for 100 epochs using the Adam optimizer [19], with a learning rate of 0.001 and a batch size of 64. Regularisation was applied using a dropout rate of 0.1 and weight decay of 0.0001.

### III-C. Filter Spectrum vs. Between-Class Variation

To test our method, we retrieve the filter spectra for the multivariate models trained on both datasets and compare it to the between-class variation of each channel. Given that filter settings are shared across channels, only one filter spectrum is retrieved for each dataset.

For the EEGNet, a filter size of 50 at a sampling rate of 100 Hz corresponds to a detectable frequency range of 0 Hz to 48 Hz, with a frequency spacing of 2 Hz. In the case of the MSA-CNN, the filters operate on multiple scales. To unify the scales, a unification assignment matrix was computed using the pooling settings shown in Figure 1(a). Based on the filter size setting of 15 and a sampling rate of 100 Hz, the unified frequencies span the range 0-46.7 Hz with a maximal frequency spacing of 6.7 Hz.

Figure 2(a-b) shows the filter spectrum for the EEGNet model for datasets ISRUC-S3 (a) and Sleep-EDF-20 (b), while Figure 2(c-d) shows the spectrum for the

MSA-CNN model for both datasets. The filter spectrum is overlaid with the between-class variation spectra of the EEG channels F3-A2 (ISRUC-S3) and Fpz-Cz (Sleep-EDF-20), which are commonly used in univariate configurations [15, 21]. All spectra were rescaled by dividing by the standard deviation in the frequency range 0.5-12 Hz, which corresponds to the combined lower EEG frequency bands  $\delta$  (0.5-4 Hz),  $\theta$  (4-8 Hz), and  $\alpha$  (8-12 Hz) [20]. For both models, the results indicate substantial similarity between filter spectrum and between-class variation in these lower frequency bands on both datasets. For example, both models capture a peak in between-class variation in the *alpha* band on the ISRUC-S3 dataset, whereas on the Sleep-EDF-20 dataset they capture a plateau. This result indicates that the model uses mainly lower frequencies for spectral information-based classification. On the other hand, the filter spectra align across datasets in the higher frequency bands  $\beta$  and  $\gamma$ , which suggests that higher frequencies may be used to construct complex wave patterns shared across EEG datasets.

### III-D. Spectral Information Extraction Relative to Modality

To understand the specific role of channel and modality in extracting spectral information, we compute the Pearson correlation of filter spectrum and between-class variation in the lower frequency bands ( $\delta$ ,  $\theta$  and  $\alpha$ ) across channels, as shown in Figure 3(a,b). We find that the correlation is highest for EEG, followed by EOG and EMG, irrespective of the model or dataset. This pattern is less pronounced for the ISRUC-S3 dataset, where EOG channels correlate almost as strongly with filter spectrum as EEG channels. There is also a within-modality discrepancy between frontal and occipital EEG channels: While for ISRUC-S3, frontal (blue) and occipital (light blue) channels correlate equally with filter spectrum, for Sleep-EDF-20 occipital channels correlate stronger than frontal channels.

We validate these results by comparing them to the univariate performance for each channel and modality. The univariate performances, shown in Figure 3(c,d), exhibit a similar pattern, with EEG outperforming EOG and EMG. This matching between correlation and performance indicates that both models are capable of prioritising more informative channels by optimising the retrieval of spectral information for these channels. Given the unimodal configuration of the first layer in both models, this prioritisation comes at the cost of neglecting less informative modalities.

## IV. CONCLUSION

This study introduced a tool to explain the role of spectral information processing in CNN-based classification models. We tested the tool on two multivariate sleep stage classification models across two datasets.

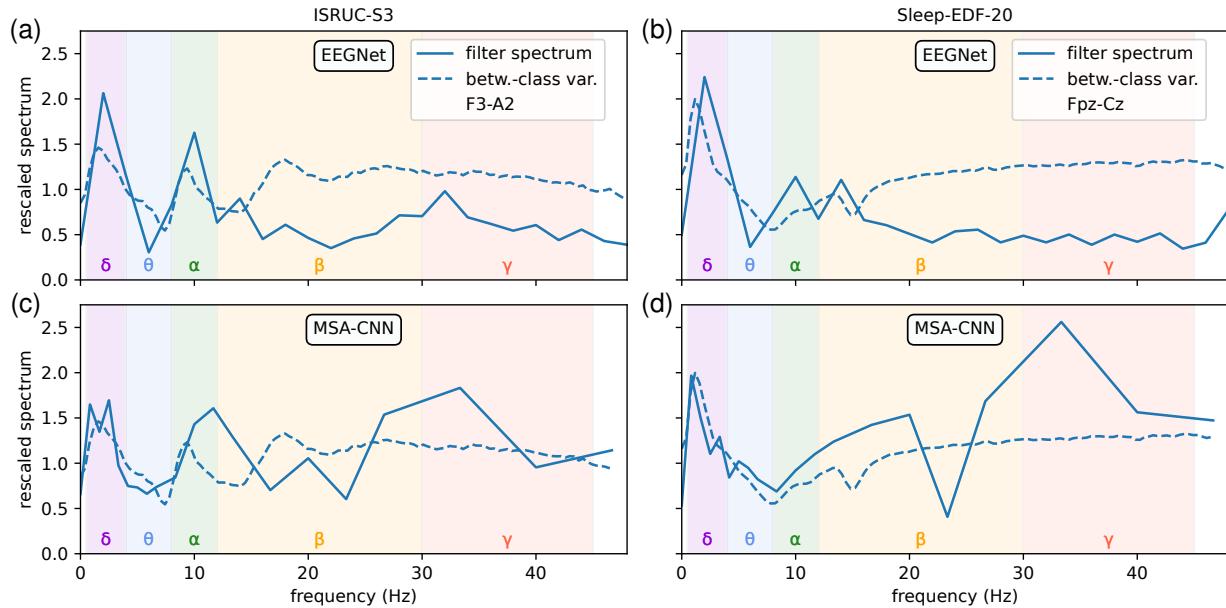


Figure 2. Rescaled filter spectrum (solid lines) for EEGNet (a,b) and MSA-CNN (c,d) on datasets ISRUC-S3 (a,c) and Sleep-EDF-20 (b,d). The spectra are overlaid with the rescaled between-class variation of the frontal channels F3–A2 (ISRUC-S3) and Fpz–Cz (Sleep-EDF-20), shown as dashed lines. Shaded regions along the frequency axis indicate the standard EEG frequency bands  $\delta$  (0.5–4 Hz),  $\theta$  (4–8 Hz),  $\alpha$  (8–12 Hz),  $\beta$  (12–30 Hz), and  $\gamma$  (30–45 Hz) [20]. At lower frequency bands ( $\delta$ ,  $\theta$  and  $\alpha$ ), filter spectra and between-class variations align. At higher frequency bands ( $\beta$  and  $\gamma$ ), filter spectra align across datasets.

For both models, we found that the convolutional filter spectrum aligns with the spectral information available in the data at lower frequencies, which indicates that the models optimise spectral information extraction at these frequencies. Note that the available spectral information depends on the dataset and application. For example, ISRUC-S3 shows greater between-class variation in the  $\alpha$  wave range than Sleep-EDF-20, likely reflecting differences in EEG channel configuration and system characteristics. The observed alignment highlights the way in which convolutional layers can function as spectral feature extractors [22].

At higher frequencies, filter spectrum and data-derived spectral information diverge, while filter spectra align across datasets. This suggests that the filters capture the general morphology of EEG waves rather than spectral information at these frequencies. The capability to capture both spectral and morphological features is in contrast to spectral density extractors such as Welch's method, which are limited to spectral features. Altogether, the explainability tool provided insight into the patterns of information processing relative to frequency, which were strikingly similar between the two models. In general, such similarity suggests that the models are functionally comparable.

In a second step, we used the explainability tool to analyse channel importance. We found that both models prioritise the more informative channels at the expense of less informative ones. Apart from providing insight

into model functioning, this finding may be leveraged for performance optimisation in the following ways: First, channel selection for single- and multi-channel applications can be performed with only a single trained multivariate model. Second, comparisons between filter spectrum and between-class variation can be used to balance the frequency scales for optimal extraction of spectral information. Third, large differences in between-class variation can signal the need to perform feature extraction separately for each modality. Together, these examples underscore the value of our method for model improvement specifically, and the value of explainability for healthcare applications more generally.

Nevertheless, several limitations should be noted. First, the correlation estimates may be unreliable due to the limited number of data points used in their computation. In addition, the comparison between correlation estimates and performance for channel importance analysis was only assessed visually. Lastly, the relatively small size of the datasets used in this study may limit the transferability of our model [23]. In future research, we will improve the channel importance analysis by incorporating channel ablation experiments. In addition, we plan to test our explainability tool on larger and more diverse datasets.

#### ACKNOWLEDGEMENTS

Stephan Goerttler was supported by the A\*STAR Research Attachment Program (ARAP).

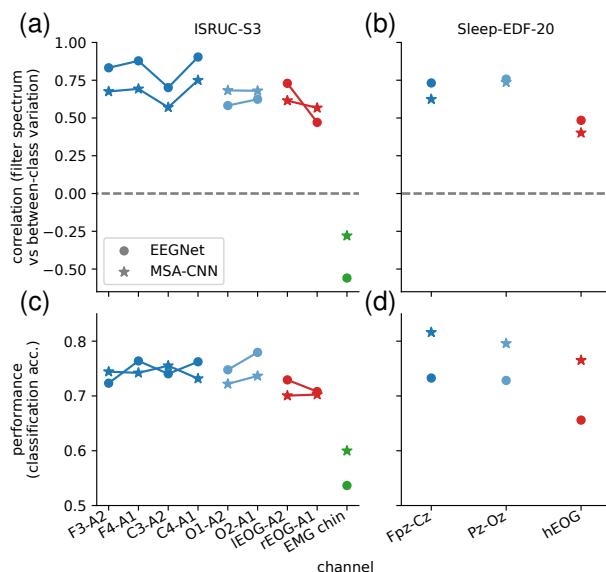


Figure 3. (a,b) Channel-wise Pearson correlation between spectral activation and between-class variation in the lower frequency bands on datasets ISRUC-S3 (a) and Sleep-EDF-20 (b). Generally, EEG channels (blue) exhibit the highest similarity, followed by EOG (red) and EMG (green). (c,d) Channel-wise classification performance for the two respective datasets. Channel performance approximately matches the correlations in a and b.

## REFERENCES

- J. V. Rundo and R. Downey III, "Polysomnography," *Handbook of Clinical Neurology*, vol. 160, pp. 381–392, 2019.
- X. Zhang, X. Zhang, Q. Huang, Y. Lv, and F. Chen, "A review of automated sleep stage based on EEG signals," *Biocybernetics and Biomedical Engineering*, 2024.
- H. Yue, Z. Chen, W. Guo, L. Sun, Y. Dai, Y. Wang, W. Ma, X. Fan, W. Wen, and W. Lei, "Research and application of deep learning-based sleep staging: Data, modeling, validation, and clinical practice," *Sleep Medicine Reviews*, p. 101897, 2024.
- K. Rasheed, A. Qayyum, M. Ghaly, A. Al-Fuqaha, A. Razi, and J. Qadir, "Explainable, trustworthy, and ethical machine learning for healthcare: A survey," *Computers in Biology and Medicine*, vol. 149, p. 106043, 2022.
- F. Vaquerizo-Villar, G. C. Gutiérrez-Tobal, E. Calvo, D. Álvarez, L. Kheirandish-Gozal, F. Del Campo, D. Gozal, and R. Hornero, "An explainable deep-learning model to stage sleep states in children and propose novel EEG-related patterns in sleep apnea," *Computers in Biology and Medicine*, vol. 165, p. 107419, 2023.
- M. Dutt, S. Redhu, M. Goodwin, and C. W. Omlin, "SleepXAI: An explainable deep learning approach for multi-class sleep stage identification," *Applied Intelligence*, vol. 53, no. 13, pp. 16 830–16 843, 2023.
- G. Liu, G. Wei, S. Sun, D. Mao, J. Zhang, D. Zhao, X. Tian, X. Wang, and N. Chen, "Micro SleepNet: Efficient deep learning model for mobile terminal real-time sleep staging," *Frontiers in Neuroscience*, vol. 17, p. 1218072, 2023.
- B. Adey, A. Habib, and C. Karmakar, "Exploration of an intrinsically explainable self-attention based model for prototype generation on single-channel EEG sleep stage classification," *Scientific Reports*, vol. 14, no. 1, p. 27612, 2024.
- S. Goerttler, Y. Wang, E. Eldele, M. Wu, and F. He, "MSA-CNN: A lightweight multi-scale CNN with attention for sleep stage classification," 2025. (available at: <https://arxiv.org/abs/2501.02949>).
- C. A. Ellis, R. Zhang, D. A. Carbajal, R. L. Miller, V. D. Calhoun, and M. D. Wang, "Explainable sleep stage classification with multimodal electrophysiology time-series," *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 2021, pp. 2363–2366.
- A. Craik, Y. He, and J. L. Contreras-Vidal, "Deep learning for electroencephalogram (EEG) classification tasks: a review," *Journal of neural engineering*, vol. 16, no. 3, p. 031001, 2019.
- A. Supratak, H. Dong, C. Wu, and Y. Guo, "DeepSleepNet: A model for automatic sleep stage scoring based on raw single-channel EEG," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 25, no. 11, pp. 1998–2008, 2017.
- V. J. Lawhern, A. J. Solon, N. R. Waytowich, S. M. Gordon, C. P. Hung, and B. J. Lance, "EEGNet: a compact convolutional neural network for EEG-based brain–computer interfaces," *Journal of Neural Engineering*, vol. 15, no. 5, p. 056013, 2018.
- Z. Jia, Y. Lin, J. Wang, R. Zhou, X. Ning, Y. He, and Y. Zhao, "Graphsleepnet: Adaptive spatial-temporal graph convolutional networks for sleep stage classification," *Ijcai*, vol. 2021, 2020, pp. 1324–1330.
- E. Eldele, Z. Chen, C. Liu, M. Wu, C.-K. Kwok, X. Li, and C. Guan, "An attention-based deep learning approach for sleep stage classification with single-channel EEG," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 29, pp. 809–818, 2021.
- F. Zucchi and T. Lampert, "Prism: Lightweight multivariate time-series classification through symmetric multi-resolution convolutional layers," *arXiv preprint arXiv:2508.04503*, 2025.
- S. Khalighi, T. Sousa, J. M. Santos, and U. Nunes, "ISRUC-Sleep: A comprehensive public dataset for sleep researchers," *Computer Methods and Programs in Biomedicine*, vol. 124, pp. 180–192, 2016.
- A. L. Goldberger, L. A. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley, "Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000.
- D. P. Kingma, "ADAM: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- E. Niedermeyer and F. L. da Silva, *Electroencephalography: basic principles, clinical applications, and related fields*. Lippincott Williams & Wilkins, 2005.
- D. T. Pham and R. Mouček, "Automatic sleep stage classification by cnn-transformer-lstm using single-channel eeg signal," *2023 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE, 2023, pp. 2559–2563.
- O. Rippel, J. Snoek, and R. P. Adams, "Spectral representations for convolutional neural networks," *Advances in Neural Information Processing Systems*, vol. 28, 2015.
- A. Supratak and P. Haddawy, "Quantifying the impact of data characteristics on the transferability of sleep stage scoring models," *Artificial Intelligence in Medicine*, vol. 139, p. 102540, 2023.