

# A Convolutional Network for Sleep Stages Classification

Isaac Fernández-Varela\*, Elena Hernández-Pereira\*, Diego Alvarez-Estevez<sup>†</sup> and Vicente Moret-Bonillo\*

\*CITIC

Universidade da Coruña

A Coruña, España

(isaac.fvarela, elena.hernandez, vicente.moret)@udc.es

<sup>†</sup>Sleep Center & Clinical Neurophysiology

Haaglanden Medisch Centrum

The Hague, The Netherlands

diego.alvarez@udc.es

**Abstract**—Sleep stages classification is a crucial task in the context of sleep studies. It involves the simultaneous analysis of multiple signals recorded during sleep. However, it is complex and tedious, and even the trained expert can spend several hours scoring a single night recording. Multiple automatic methods have tried to solve these problems in the past, most of them by classifying a feature vector that is engineered for a specific dataset. In this work, we avoid this bias using a deep learning model that learns relevant features without human intervention. Particularly, we propose an ensemble of 5 convolutional networks that achieves a kappa index of 0.83 when classifying a dataset of 500 sleep recordings.

**Index Terms**—convolutional network, sleep stages, classification

## I. INTRODUCTION

Sleep disorders affect a major part of the population. As an example, 20% of the Spanish adults suffer insomnia, and between 12% and 15% daytime sleepiness [1, 2]. Good sleep is essential for a healthy life, and the adverse consequences of restless nights have been extensively reported [3]. To evaluate the sleep function, and to help the diagnosis of sleep disorders, it is important to know the sequence of sleep stages that the patient goes through the night.

The most common technique to monitor the sleep function is the polysomnogram (PSG), which involves recording of the patient’s biosignals during sleep, including various pneumological, electrophysiological, and contextual information. This is an expensive test, uncomfortable for the patient, and for which interpretation of the results is difficult due to the complexity of the data involved. An usual way to summarize the sleep information contained in the PSG is the derivation of the hypnogram, an ordered representation of the sleep stages evolution.

The current gold standard for the building the hypnogram is the *American Academy of Sleep Medicine* (AASM) [4] guide for the identification of sleep stages and of their associated events (e.g. EEG arousals, limb movements, and cardiac or

respiratory events). This guide identifies five sleep stages: Awake (W), Rapid Eye Movements (REM), and 3 non-REM phases (N1, N2, and N3). Correct identification of the sleep stages and construction of the hypnogram is of fundamental importance to achieve a good diagnosis, allowing the clinician to focus efforts in the therapy. Such a task implies the analysis of huge amounts of data and expert knowledge [5]. Moreover, even following the guidelines, inter-expert agreement usually remains below the 90%. For example, Stepnowsky et al. [6] studied the agreement between two experts finding kappa index values between 0.48 and 0.89. Similarly, Wang et al. [7] found values between 0.72 and 0.85. Furthermore, agreement is worse for some particular stages, usually being stage N1 the one with the highest disagreement.

All given, automatic methods for sleep stages classification are needed. Most of these methods follow a two step approach. First, feature extraction takes place, usually with features hand tailored for a specific dataset. Then, feature vectors are built to train a classifier and predict the sleep stages. While some authors have used a single signal channel as reference (usually the EEG), other approaches have extracted features using several channels, building input vectors of various elements. At this respect usually features from the electrooculogram (EOG) or electromiogram (EMG) are added to those of the EEG, as recommended by the AASM guidelines. Often features are extracted either from the time or from the frequency domain.

Among the methods following this 2-step approach we find: Fraiwan et al. [8] use a random forest to classify features both from the time-frequency domain and Renyi’s entropy; Liang et al. [9] measure entropy with different scales obtaining autoregressive features which classify using a linear discriminant; Hassan and Bhuiyan [10], apply wavelet transformations for feature extraction and use a random forest technique for the classification step. Sharma et al. [11], compare several classifiers for iterative filters analysing a single EEG channel; Koley and Dey [12], train a support vector machine (SVM) with frequency, time and non-linear features extracted from a single EEG channel; Lajnef et al. [13], base their approach on multiple signals building a decision tree upon several SVMs;

\* This research was funded by the Xunta de Galicia (ED431G/01) and the European Union (ERDF).

Huang et al. [14], study power spectral density of 2 EEG channels classifying frequency features with a modified SVM; Finally, Günes et al. [15], also analyse power spectral density while classifying with a nearest neighbours algorithm.

The approach consisting in solving the sleep staging classification problem using handcrafted feature extraction induces biases due to the design of features based on one specific database. Thus, the aforementioned solutions usually do not generalize well, specially given the nature of PSG recordings, where variability effects are introduced due to several factors, including patient, hardware or scoring differences.

One alternative option to solve this problem is the use of methods than learn directly from the raw data, therefore avoiding the human bias. In this sense, deep learning represents a natural approach, as it demonstrated improvements against traditional methods in multiple general fields, including in particular, the medical diagnosis [16, 17].

Some works have already explored solutions with different deep learning models: Långkvist et al. [18], used deep belief networks learning a probabilistic representation of preprocessed signals from PSG inputs; Tsinalis et al. [19], still followed the 2 step approach, but with convolutional networks for classification. In other work, the same authors [20] relied on a stack of *sparse autoencoders*; Supratak et al. [21], performed classification from the raw signals with a bidirectional recurrent neural network; Biswal et al. [22], compared a recurrent network against different models, although all were trained with features instead of the raw signal; Finally, Sors et al. [23] also used a convolutional neural network using one single EEG channel as reference.

In this work we use deep learning to classify sleep stages with a convolutional neural network that learns the relevant features for each stage. Following the AASM guidelines we use multiple signals; namely, two EEG, one EMG, and two (left and right) EOG channels. Moreover, signals are filtered in the first place, to reduce noise and remove artifacts.

## II. MATERIALS

Design and analysis of the presented model was carried out using PSG recordings from real patients. These recordings belong to the Sleep Heart Health Study (SHHS) [24], a database offered by the Case Western University, originated from a cohort study involving multiple centers directed by the National Heart Lung and Blood Institute, with the goal of determining the cardiovascular consequences of respiratory related sleep disorders.

Each recording contains annotations for different events performed by clinical experts following the procedures described in [25]. All recordings were anonymized and blind scored. The montage for the signals acquisition included two EEG derivations (C4A2 and C4A1), left and right EOGs, chin EMG, and modified lead-II electrocardiogram (ECG). EEG, EOG, and EMG were sampled at 125 Hz whereas EOG were sampled at 50 Hz. All signals were filtered during acquisition with a high pass filter at 0.15 Hz.

From this database three different datasets were selected to train, validate and test our model. Training dataset included 400 recordings, validation 100, and test 500. The length of the training recordings is matched (limiting each to a total of 7 randomly selected hours) to facilitate the coding and the training of the algorithm. Finally, our training dataset contained 288.000 30-s epoch samples, the validation dataset 119.121 and the test dataset 606.981. Recordings were selected randomly, including those with high levels of noise or artifacts.

The distribution for the different classes, both for the complete dataset as for each individual recording is shown in Table I. This table shows how unbalanced the datasets are, being W the most represented class (about 38% of the samples), although with a similar proportion to N2 (around 36%). On the contrary, class N1 is only represented in 3% of the classes. It is also interesting to notice how some recordings do not contain samples for some of the classes, and how much the distribution differs between the recordings. For example, in the test dataset, whereas a particular recording contains a 7.10% of samples for class N2, another goes up to a 83.43%. Moreover, these are the two important problems when trying to develop an automatic sleep staging classifier: 1) the class unbalance and 2) the differences between individual recordings.

## III. METHODS

### A. Signal filtering

Signals are preprocessed to reduce noise and remove common artifacts. Both operations are typically applied in previous works before feature extraction.

The first of the two filters used to reduce noise is a Notch filter centered at 60 Hz to remove mains interference. This filter is applied to those signals with a sampling rate higher than 60 Hz: EEG and EMG. The second one removes DC component and frequencies not related with muscular movements from the EMG, applying a high pass at 15 Hz.

Regarding artifacts, most of them happen during particular short time periods, making it difficult even their detection. However, ECG artifacts, caused by the heart beat interference, are common and constant through the whole signals. We can remove this kind of artifact with an adaptive filter. To do so, we first obtained the beat series following a standard QRS detection algorithm [26]. Then, we studied the signal quality to assess which intervals could be safely included in the construction of the adaptive filter. Finally, during the intervals with enough signal quality, we applied and updated the filter template to remove the artifacts. More information about this process can be found in Fernández-Varela et al. [27].

### B. Convolutional network

Sleep stages classification is usually carried out with 30 s windows called epochs. Analyzing several features from each epoch, clinicians score the corresponding sleep stage.

A convolutional neural network [28] is a feedforward network solving the limitations of the multilayer perceptron with a weight sharing architecture. Basically, it applies a

TABLE I  
DISTRIBUTION OF THE DIFFERENT CLASSES IN THE TRAINING, VALIDATION, AND TEST DATASETS.

		W	N1	N2	N3	REM	Total
<i>Training dataset</i>	Total	187.513	17.283	172.451	44.454	62.168	483.869
	Proportion	38,75 %	3,57 %	35,64 %	9,19 %	12,85 %	100 %
	Min in single record	8,20 %	0,00 %	12,59 %	0,00 %	0,00 %	
	Max in single record	71,61 %	13,75 %	68,65 %	33,43 %	26,58 %	
<i>Validation dataset</i>	Total	43.742	3.963	43.510	12.900	15.006	119.121
	Proportion	36,72 %	3,33 %	36,53 %	10,83 %	12,60 %	100 %
	Min in single record	11,21 %	0,29 %	12,38 %	0,00 %	0,00 %	
	Max in single record	76,79 %	17,08 %	60,09 %	30,16 %	23,68 %	
<i>Test dataset</i>	Total	231.707	19.769	217.246	61.281	76.978	606.981
	Proportion	37,77 %	3,26 %	35,96 %	10,25 %	12,75 %	100 %
	Min in single dataset	7,75 %	0,00 %	7,10 %	0,00 %	0,00 %	
	Max in single dataset	76,53 %	16,93 %	83,43 %	43,82 %	31,11 %	




Fig. 1. Proposed convolutional neural network

convolution operation over the input, limiting the number of parameters. Thus, it allows the construction of deeper networks that are better at recognizing complex features. The proposed network is represented in Figure 1.

The input to the convolutional network is the set of signals (2 EEG channels, EMG, and both EOGs). Each input pattern corresponds to a 30 s epoch window. As the signals are sampled at different rates (aforementioned in Section II) we upsampled those with sampling rates lower than 125 Hz. We avoided downsampling to 50 Hz because it would mean losing high frequencies in the EEG that should contain important information from a clinical perspective. Moreover, we also discarded padding because the approach cannot be easily generalized to other datasets with different sampling rates. This way, each input to the network is a matrix with a dimension of  $3750 \times 5$ . Each signal was normalized with mean 0 and deviation 1, using the mean and deviation obtained from all the respective signals in the training dataset. When we tried other normalizations with lower granularity, our training did not converge. The convolutional block shown in Figure 1 is a stack of four layers including a 1D convolution that preserves the input dimension (with padding), a batch normalization layer [29] to improve regularization, ReLu [30] activation, and an average pool that reduces dimension by a factor of 2. By using 1D convolution we avoided imposing a spatial structure between our signals that is unknown a priori. This stack was repeated  $n$  times, being  $n$  an hyperparameter with a value selected during experimentation. All layers were configured with the same kernel size but the number of filters for layer  $i$  is twice the number of filters for layer  $i - 1$ . The selection value of  $n$ , the kernel size and the number of filters for the

first layers is explained in the following Section, together with the remaining hyperparameters.

The output of the last convolutional block, after adjusting dimensions with a global pooling and applying dropout, is used as input for a dense layer with a softmax activation. This layer returns the probability for each sleep stage given the initial input. As usual, the final predicted class is set to the output showing the highest probability.

To train the network we used Adam optimizer [31] and a batch size of 64. This batch size was limited by our hardware. The learning rate was configured whereas both betas are left with the default values. Training ends using early stopping by monitoring the validation loss with a patience of 10 epochs. To limit the impact of class unbalance, we used weighted cross entropy as the cost function, where weights were obtained using the training dataset.

### C. Hyperparameter optimization

A good selection of hyperparameters can mean the success of a deep learning model. The difficulty when selecting the best hyperparameters is not only to achieve the best performance, but doing it while at the same time minimizing the cost, either the economical or the computational cost.

In this work we relied on a Tree-structured Parzen Estimator (TPE) that has shown better performance than other methods [32, 33]. TPE is a sequential models based optimization. This kind of methods builds models sequentially to approximate the performance of hyperparameters selection based on historical results, and then chooses new hyperparameters that are checked with the model. Particularly, TPE uses two distributions  $P(x|y)$  and  $P(y)$  where  $x$  represents the hyperparameters and  $y$  the expected performance. The expected improvement (EI) is optimized according to the following equation:

$$EI_{y^*}(x) = \int_{-\infty}^{y^*} (y^* - y) \frac{P(x|y)P(y)}{P(x)}$$

where  $y^*$  is a quantil  $\gamma$  of the observed values  $y$  such as  $p(y < y^*) = \gamma$ .

We used TPE to select the best values for the following hyperparameters related with the convolutional network: the number of convolutional blocks, kernel size for the 1D convolutions, and the number of filters for the first convolutional block. Moreover, there is also a relationship between the number of blocks and the number of initial filters. Given our hardware restrictions, we did not add blocks that would have more than 1024 filters. We also used TPE to select the learning rate. The distributions for the random values of each of these hyperparameters are summarized in Table II.

TABLE II  
DISTRIBUTIONS FOR THE HYPERPARAMETERS

Hyperparameter	Distribution
Convolutional Blocks	Uniform between 1 and 10
Kernel Size	Uniform between 3 and 50
First Block Filters	Choice between 8, 16, 32 o 64
Learning Rate	Log-uniform between -10 and -1

To reduce the computational time for the hyperparameter selection we used a subset from the training set in order to train, validate, and test the different models. This subset contained 250 recordings where 20 were used for validation during training, and 50 to test each model. In total, we tried 50 different hyperparameter configurations, using the kappa index obtained with the test set as the criterion to select the best one.

#### D. Performance

The performance of the models was evaluated using the following metrics:

- **Precision**, the fraction between true positives and the predicted positives.
- **Sensitivity**, the fraction between true positives and the samples belonging to that class.
- **F1 score**, harmonic mean between precision and sensitivity.
- **Kappa**, agreement measure between two classifiers that takes into account the chances of random agreement. Perfect agreement gets a value of 1, and by chance a value of 0.

## IV. RESULTS

Before focusing on the results achieved with the final model, performance of the different models evaluated during the hyperparameters search is shown in Figure 2. Data in the figure suggest a clear trend toward low learning rates to ensure convergence.

To improve the results obtained by a single model we used an ensemble. Thereby, several models classify the same input, and the final decision is taken using the majority vote. In this case, we selected the 5 best models obtained during the hyperparameter selection. Values for the hyperparameters for each of those models are shown in Table II.

Results obtained with the ensemble using the test set are shown in Table IV. The best classification was achieved for

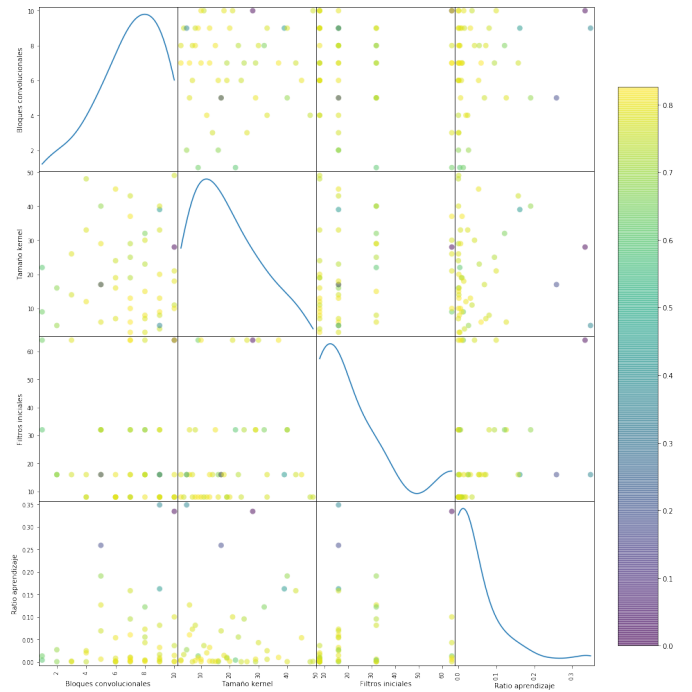


Fig. 2. Dispersion graph with the different configurations of hyperparameters. Each point color represents the kappa index for the model with the values for the hyperparameters represented in the axes. Diagonal represents the distribution for the values tried for a particular hyperparameter.

class W, with values near to 0.95 for the precision, sensitivity and F1 score; then, classes N2, N3, and REM showed similar results, specially if we compare the F1 score, although sensitivity for N3 was lower (thus, precision was higher). Lastly, results regarding the classification of class N1 were rather low, not even achieving a F1 score of 0.3. However, N1 is typically the most difficult class to predict, showing the highest disagreement also among trained experts.

The confusion matrix obtained with the ensemble is shown in Figure 3, where we can verify how most of the N1 samples are misclassified, specially towards class N2. Also, although in a smaller proportion, whenever there is a classification error it tends to be misclassifying as N2.

## V. DISCUSSION AND CONCLUSIONS

In this work we present an ensemble of convolutional networks for the classification of sleep stages. Sleep staging is a time consuming task, nevertheless critical for a good diagnosis of sleep disorders. Most of the automatic methods reported so far are based on human engineered features, designed for a particular dataset. Thus, it is difficult to find a method that generalizes correctly to other datasets. To solve this problem we propose the use of a convolutional network that self learns the relevant features for the classification, avoiding human biases.

An important aspect for the success or failure of convolutional methods is the correct choice of the hyperparameters. In this paper, we experimented with 4 hyperparameters, op-

TABLE III  
HYPERPARAMETERS FOR THE 5 MODELS WITH THE BEST KAPPA INDEX

Parameter	Model 1	Model 2	Model 3	Model 4	Model 5
Convolutional blocks	7	9	7	7	7
Kernel size	6	9	13	3	10
Initial filters	16	8	8	8	64
Learning rate	$5,99 \times 10^{-2}$	$9,00 \times 10^{-3}$	$1,45 \times 10^{-3}$	$1,91 \times 10^{-3}$	$5,49 \times 10^{-3}$

TABLE IV  
PERFORMANCE MEASURES FOR THE CLASSIFICATION OF THE TEST DATASET USING THE ENSEMBLE WITH THE 5 SELECTED MODELS.

Stage	Precision	Sensitivity	F1 score
W	0,94	0,96	0,95
N1	0,39	0,21	0,27
N2	0,87	0,89	0,88
N3	0,92	0,77	0,84
REM	0,82	0,90	0,86
<b>Average</b>	0,78	0,75	0,76




Fig. 3. Confusion matrix for the classification of the test dataset using the ensemble with the 5 selected models.

timizing their values with a tree-structured parzen estimator, trying 50 different configurations.

Our ensemble, built from the best 5 hyperparameters configurations, achieved an average precision, sensitivity, and F1 score of 0,78, 0,75 y 0,76 respectively, with a kappa index value of 0.83. Although globally our results are acceptable, our solution has shown problems for the classification of class N1. Also, in the event of misclassification, a trend has been noticed towards class N2.

Comparison of our results against similar works is difficult given the lack of standardization, both as with regard to the chosen datasets, as well as in the procedures for the evaluation process. In Table V we show results from previous works, limiting to those that report values separately for each class. As it can be seen, our kappa index is the highest, although it is not the case for the F1 score. According to the F1 score, and apart from class W, some works are able to achieve better classification for the remaining classes. However, the values that we obtained are competitive, excluding class N1, although

it is clear from all the results, that this is the most difficult class. Taking as reference the only work showing results with a similar dataset [23], our kappa index and F1 score for W class are higher, with similar values for N2, N3, and REM but lower for class N1.

Our results are promising and the chosen method should be easily adaptable to other datasets, specially if we can train the model for the different dataset. Moreover, training it with more than one dataset should improve generalization, avoiding biases for a single dataset.

To improve our result it is necessary to understand why and how the network is classifying. Also, it would be interesting to add memory to the model using recurrent networks, as the classification of some inputs, following the clinical definition, depends as well on the status of the neighbouring epochs.

## REFERENCES

- [1] M. M. Ohayon and T. Sagales, "Prevalence of insomnia and sleep characteristics in the general population of Spain." *Sleep medicine*, vol. 11, no. 10, pp. 1010–8, dec 2010.
- [2] J. Marin *et al.*, "Prevalence of sleep apnoea syndrome in the Spanish adult population," *International Journal of Epidemiology*, vol. 26, no. 2, pp. 381–386, apr 1997.
- [3] H. R. Colten and B. M. Altevogt, *Sleep Disorders and Sleep Deprivation*. Washington, D.C.: National Academies Press, sep 2006, vol. 6, no. 9.
- [4] R. B. Berry *et al.*, "AASM Scoring Manual Updates for 2017 (Version 2.4)." *Journal of clinical sleep medicine : JCSM : official publication of the American Academy of Sleep Medicine*, vol. 13, no. 5, pp. 665–666, may 2017.
- [5] Á. Fernández-Leal *et al.*, "A knowledge model for the development of a framework for hypnogram construction," *Knowledge-Based Systems*, vol. 118, pp. 140–151, 2017.
- [6] C. Stepnowsky *et al.*, "Scoring accuracy of automated sleep staging from a bipolar electrooculogram recording compared to manual scoring by multiple raters." *Sleep medicine*, vol. 14, no. 11, pp. 1199–207, nov 2013.
- [7] Y. Wang *et al.*, "Evaluation of an automated single-channel sleep staging algorithm." *Nature and science of sleep*, vol. 7, pp. 101–11, 2015.
- [8] L. Fraiwan *et al.*, "Automated sleep stage identification system based on time–frequency analysis of a single EEG channel and random forest classifier," *Computer Methods and Programs in Biomedicine*, vol. 108, no. 1, pp. 10–19, oct 2012.
- [9] J. Liang *et al.*, "Predicting seizures from electroencephalography recordings: A knowledge transfer strategy," in *Proceedings - 2016 IEEE International Conference on Healthcare Informatics, ICHI 2016*. IEEE, oct 2016, pp. 184–191.
- [10] A. R. Hassan and M. I. H. Bhuiyan, "A decision support system for automatic sleep staging from EEG signals using tunable Q-factor wavelet transform and spectral features," *Journal of Neuroscience Methods*, vol. 271, pp. 107–118, sep 2016.
- [11] R. Sharma, R. B. Pachori, and A. Upadhyay, "Automatic sleep stages classification based on iterative filtering of electroen-

TABLE V  
COMPARISON AGAINST PREVIOUS WORKS.

Work	Database	Kappa	F1 score				
			W	N1	N2	N3	REM
Biswal et al. [22]	Massachusetts General Hospital, 1000 recordings	0,77	0,81	<b>0,70</b>	0,77	0,83	<b>0,92</b>
Långkvist et al. [18]	St Vicent's University Hospital, 25 recordings	0,63	0,73	0,44	0,65	<b>0,86</b>	0,80
Sors et al. [23]	SHHS, 1730 recordings	0,81	0,91	0,43	0,88	0,85	0,85
Supratak et al. [21]	MASS dataset, 62 recordings	0,80	0,87	0,60	<b>0,90</b>	0,82	0,89
Supratak et al. [21]	SleepEDF, 20 recordings	0,76	0,85	0,47	0,86	0,85	0,82
Tsinalis et al. [19]	SleepEDF, 39 recordings	0,71	0,72	0,47	0,85	0,84	0,81
Tsinalis et al. [20]	SleepEDF, 39 recordings	0,66	0,67	0,44	0,81	0,85	0,76
This work	SHHS, 500 recordings	<b>0,83</b>	<b>0,95</b>	0,27	0,88	0,84	0,86

cephalogram signals,” *Neural Computing and Applications*, vol. 28, no. 10, pp. 2959–2978, oct 2017.

- [12] B. Koley and D. Dey, “An ensemble system for automatic sleep stage classification using single channel EEG signal,” *Computers in Biology and Medicine*, vol. 42, no. 12, pp. 1186–1195, 2012.
- [13] T. Lajnef et al., “Learning machines and sleeping brains: Automatic sleep stage classification using decision-tree multi-class support vector machines,” *Journal of Neuroscience Methods*, vol. 250, pp. 94–105, jul 2015.
- [14] C.-S. Huang et al., “Knowledge-based identification of sleep stages based on two forehead electroencephalogram channels,” *Frontiers in Neuroscience*, vol. 8, p. 263, sep 2014.
- [15] S. Günes, K. Polat, and S. Yosunkaya, “Efficient sleep stage recognition system based on EEG signal using k-means clustering based feature weighting,” *Expert Systems with Applications*, vol. 37, no. 12, pp. 7922–7928, dec 2010.
- [16] A. Esteva et al., “Dermatologist-level classification of skin cancer with deep neural networks,” *Nature*, vol. 542, no. 7639, pp. 115–118, feb 2017.
- [17] V. Gulshan et al., “Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs,” *JAMA*, vol. 316, no. 22, p. 2402, dec 2016.
- [18] M. Långkvist et al., “Sleep Stage Classification Using Unsupervised Feature Learning,” *Advances in Artificial Neural Systems*, vol. 2012, pp. 1–9, 2012.
- [19] O. Tsinalis et al., “Automatic sleep stage scoring with single-channel eeg using convolutional neural networks,” oct 2016.
- [20] O. Tsinalis, P. M. Matthews, and Y. Guo, “Automatic sleep stage scoring using time-frequency analysis and stacked sparse autoencoders,” *Annals of Biomedical Engineering*, vol. 44, no. 5, pp. 1587–1597, may 2016.
- [21] A. Supratak et al., “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, nov 2017.
- [22] S. Biswal et al., “Sleepnet: Automated sleep staging system via deep learning,” jul 2017.
- [23] A. Sors et al., “A convolutional neural network for sleep stage scoring from raw single-channel EEG,” *Biomedical Signal Processing and Control*, vol. 42, pp. 107–114, apr 2018.
- [24] S. F. Quan et al., “The sleep heart health study: Design, rationale, and methods,” *Sleep*, vol. 20, no. 12, pp. 1077–1085, dec 1997.
- [25] Case Western Reserve University, “Sleep Heart Health Study: reading center manual of operations,” Case Western Reserve University, Tech. Rep., 2002.
- [26] V. Afonso et al., “Ecg beat detection using filter banks,” *IEEE Transactions on Biomedical Engineering*, vol. 46, no. 2, pp. 192–202, 1999.
- [27] I. Fernández-Varela et al., “A simple and robust method for the automatic scoring of EEG arousals in polysomnographic recordings,” *Computers in Biology and Medicine*, vol. 87, pp. 77–86, aug 2017.
- [28] Y. Le Cun et al., “Handwritten digit recognition: applications of neural network chips and automatic learning,” *IEEE Communications Magazine*, vol. 27, no. 11, pp. 41–46, nov 1989.
- [29] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” 2015.
- [30] V. Nair and G. E. Hinton, “Rectified Linear Units Improve Restricted Boltzmann Machines,” *Proceedings of the 27th International Conference on Machine Learning*, no. 3, pp. 807–814, 2010.
- [31] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” dec 2014.
- [32] J. Bergstra et al., “Algorithms for hyper-parameter optimization,” in *NIPS*, 2011.
- [33] J. Bergstra, D. Yamins, and D. D. Cox, “Hyperopt: A python library for optimizing the hyperparameters of machine learning algorithms,” in *Proc. of the 12th python in science conf*, 2013.