Deep-Learning-Based Classification of Cyclic Alternating Pattern Sleep Phases

Yoav Kahana

Deep-Learning-Based Classification of Cyclic Alternating Pattern Sleep Phases

Research Thesis

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Electrical Engineering

Yoav Kahana

Submitted to the Senate
of the Technion — Israel Institute of Technology
Shevat 5785 Haifa February 2025

This research was carried out under the supervision of Prof. Israel Cohen and Dr. Alon Amar, in the Faculty of Electrical and Computer Engineering.

The results of Chapter 3 and 4 of this thesis have been published as an article during the course of the author's masters research period, the most up-to-date version of which is:

> Yoav Kahana, Aviad Aberdam, Alon Amar, and Israel Cohen. Deep-learning-based classification of cyclic-alternating-pattern sleep phases. *Entropy*, 25(10):1395, 2023.

The author of this thesis states that the research, including the collection, processing and presentation of data, addressing and comparing to previous research, etc., was done entirely in an honest way, as expected from scientific research that is conducted according to the ethical standards of the academic world. Also, reporting the research and its results in this thesis was done in an honest and complete manner, according to the same standards.

Acknowledgements

I would like to express my sincere gratitude to my supervisors, Prof. Israel Cohen and Dr. Alon Amar, for their guidance, wisdom, and support throughout my degree.

I am also deeply thankful to my parents for their love and the education they provided me.

Finally, to my dear wife Ronit, thank you for being by my side every step of the way.

Contents

List of Figures

Abstract

In this thesis we study the problem of classification of cyclic alternating pattern (CAP) phases in sleep. Determining the cyclic-alternating-pattern phases using electroencephalography (EEG) signals is crucial for assessing sleep quality. CAP is divided, generally, into two main phases by the distinction between cerebral activation (A-phase) and de-activation (B-phase) modes. Beyond being a physiological phenomenon, CAP is considered a reliable marker of sleep instability as it can correlate with several sleep-related pathologies. Consequently, accurate detection of the CAP phases has a crucial role in a sleep diagnosis. Nevertheless, most current methods for CAP classification primarily rely on classical machine learning techniques, with limited implementation of deep learning-based tools. Furthermore, these methods often require manual feature extraction. To bridge this gap, in this thesis we aim to leverage the capabilities of deep convolutional neural networks (CNN) for classifying CAP phases. Our approach is motivated by the remarkable strides that CNNs have made in recent years, achieving state-of-the-art performance in a wide range of tasks and applications.

For this goal we present an algorithm consisting of three main stages. Firstly, we analyze each second of the long EEG signal as a distinct prolonged 1D-EEG segment, considering the contextual information of the signal. Next, we transform each time segment into its time-frequency representation (TFR). This representation is highly suitable for classifying CAP phases due to the non-stationary nature of the signals, and it allows us to utilize a CNN-based architecture effectively. The TFR is treated as a 2D image and fed into a convolutional neural network that classifies it as either an A-phase or a B-phase sample. Our experimental results demonstrate that we achieve state-of-the-art performance even with the compact ResNet18 architecture, which can be efficiently run on-device.

Through our investigation, we explored using time-frequency analysis techniques and found that Wigner-based representations outperform the commonly used short-time Fourier transform for CAP classification. We conduct an ablation study to assess the contribution of each part of the proposed system. Evaluating the proposed method was conducted over the PhysioNet's public Cap Sleep Database (CAPSLPDB), considered a benchmark database for CAP identification and classification research. A thorough experimental study demonstrates that our algorithm surpasses existing machine learning-based methods. Overall, we conclude that the proposed algorithm exhibits

efficiency and scalability, making it suitable for on-device implementation to enhance CAP identification procedures.

Abbreviations

AASM	:	American Academic of Sleep Medicine
ACC	:	Accuracy
AF	:	Ambiguity function
AT	:	Auto-terms
AWGN	:	Additive white Gaussian noise
BOWFB	:	Biorthogonal wavelet filter bank
CAP	:	Cyclic-alternating-pattern
CAPSLPDB		Cap Sleep Database
CNN	:	Convolutional neural network
CT	:	Cross-terms
DL	:	Deep learning
EEG	:	Electroencephalography
EMG	:	Electromyogram
EOG	:	Electrooculogram
LDA	:	Linear discriminant analysis
LSTM	:	Long-short term memory
PRE	:	Precision
PSD	:	Power spectral density
PSG	:	Polysomnography
RE	:	Renyi entropy
REC	:	Recall
REM	:	Rapid-eye-movement
SGD	:	Stochastic gradient descent
SPE	:	Specificity
SPEC	:	Spectrogram
SPWVD	:	Smoothed pseudo Wigner-Ville distribution
STFT	:	Short-time-Fourier-transform
SVM	:	Support-vector-machines
TF	:	Time-frequency
TFR	:	Time-frequency representation
TPR	:	True-positive-rate
WVD	:	Wigner-Ville-distribution

Notations

$A_x(\nu, \tau)$	:	Ambiguity function of $x(t)$
$h(t)$	:	Window function
N	:	Number of samples in dataset
p	:	Probability
s_i	:	i-th element of s
$S_x(t, f)$	:	Spectrogram of signal $x(t)$
$W_x(t, f)$	:	Wigner-Ville distribution of signal $x(t)$
$W_x^{sp}(t, f)$	:	Smoothed-pseudo-Wigner-Ville distribution of signal $x(t)$
$x(t)$	:	Time domain EEG signal
x_i	:	i-th element of x
$\bar{x_i}$	:	Mean of data sample
$\tilde{x_i}$	:	Normalized data sample
$X(t, f)$	:	STFT domain of signal
X_i	:	i-th element of X
X_N	:	Set of data samples
$\Phi(t, f)$	:	Low-pass kernel function
σ_i	:	Standard deviation of data sample
$\bar{\sigma}$	:	Mean of standard deviations

Chapter 1

Introduction

1.1 Background and Motivation

Detecting sleep stages is essential for understanding and improving sleep quality and identifying and addressing many sleep-related pathologies. Especially the cyclic-alternating-pattern (CAP) is considered a key concept in evaluating the sleep process [1]. CAP is divided, generally, into two main phases by the distinction between cerebral activation (A-phase) and de-activation (B-phase) modes [2]. Beyond being a physiological phenomenon, CAP is considered a reliable marker of sleep instability [3] as it can correlate with several sleep-related pathologies [4]. Consequently, accurate detection of the CAP phases has a crucial role in a sleep diagnosis. Traditionally, sleep analysis relies on polysomnography (PSG) and is conducted by trained physicians and healthcare professionals in sleep laboratories. This approach poses significant challenges in terms of practicality and clinical applicability. The manual assessment process is labor-intensive, time-consuming, and susceptible to human fatigue, subjectivity, and potential errors. The development and implementation of such automated tools not only enhance the reliability and precision of sleep diagnosis but also have the potential to enhance clinical workflows, reduce healthcare costs, and ultimately improve patient care and outcomes.

To streamline and facilitate this vital process, various methods were proposed over the years to automate the detection of the CAP phases [5–16]. In many studies, the standard approach for CAP classification involves using feature extraction techniques to generate input data for a classifier, which aims to differentiate between the A and B phases. The feature extraction is typically based on prior scientific knowledge regarding the distinctions in energy and frequency content between the A and B phases (as detailed in 2.1). For instance, in [5,6], the EEG signal was divided into distinct frequency bands, and the power-spectral-density (PSD) was computed for each band separately. Subsequently, PSD-based features were extracted to feed various classifiers - [5] utilized a linear discriminant analysis (LDA) that assumes the data to be produced based on Gaussian distributions [17], while in [6] different supervised and unsupervised classifiers were evaluated including decision trees, support-vector-machines (SVM), k-means

clustering and others. Similarly, either [7–10] partitioned the EEG signals into different frequency bands, while in their studies, variance indices were utilized as features. As a classifier, a 3-layer neural network was employed in [7], while [8] used SVM and [9] utilized the LDA classifier. In [10], all these classifiers were compared to an adaptive boosting (AdaBoost) classifier, resulting in the superiority of the LDA classifier.

In several works [11–13, 18, 19], time-frequency transforms were utilized to address the pre-mentioned distinctions among the CAP phases. Particularly, [11, 13, 18, 19] employed the wavelet transform, while [12] used the Wigner-Ville distribution (WVD). Nevertheless, in all these works, the time-frequency transforms were used as a temporary representation for hand-crafted feature extraction, similar to the previous studies.

In recent research, there has been an emerging utilization of deep learning (DL) techniques for classifying CAP phases. Primarily, [15] achieved high performance ($82.4\% \pm 7.08\%$ accuracy) by employing a long-short term memory (LSTM) network. Nevertheless, it is worth noting that in this work, the neural network was fed by several hand-crafted features, and its outcomes were post-processed to improve performance by the CAP scoring guidelines outlined in [2].

In [14], a 1-dimensional convolutional neural network (1D-CNN) was suggested for both CAP classification and sleep macrostructure scoring task. Similarly, [20] employed a comparable 1D-CNN architecture to classify CAP phases of healthy and sleep-disordered individuals. The raw EEG signal was standardized in their works before feeding the 1D-CNN. The trained model was tested on both balanced and unbalanced test sets. At the same time, the outcomes indicated moderate performance when tested on an unbalanced dataset (52.99% in [14] and 60.59% in [20]). A 1D-CNN was also utilized in [16] but using a significantly more complex model based on the U-Net framework and a gated-transformer module to extract local features and global contexts.

As illustrated, a significant majority of the current methods rely on (1) hand-crafted feature extraction, which may not capture all relevant information of the data, and (2) traditional machine learning-based approaches. Yet, despite the increasing usage of deep learning tools in recent years [21] and its proven success in a broad spectrum of tasks [21–24], their application to classifying CAP phases remains limited. To bridge this gap, we aim to leverage the capabilities of deep convolutional neural networks (CNN) for classifying CAP phases. Our approach is motivated by the remarkable strides that CNNs have made in recent years [25], achieving state-of-the-art performance in a wide range of tasks and applications, including image classification [26], object tracking [27], text detection [28], speech and natural language processing [29] and others.

1.2 Main Contributions

The main contributions of the research in this thesis are:

- **Automatic Classifier Development:** A novel fully automatic classifier for cyclic alternating patterns (CAP) signals is introduced. The classifier is based on a computationally efficient neural network, which therefore can be implemented on-device, providing a practical and accessible solution for real-world implementation.

- **State-of-the-Art Performance:** Extensive experiments demonstrate state-of-the-art performance on a public CAP benchmark database, classifying its A and B phases using only a single EEG signal. Our classifier demonstrates superior accuracy and performance when applied to both balanced and unbalanced datasets which underscores its clinical relevance and robustness.

- **Ablation Study Insights:** A comprehensive ablation study to systematically assess the impact of various factors on the classifier's performance was conducted. This in-depth analysis includes the evaluation of different time-frequency representations, segment sizes and various of data augmentation techniques.

- **Bridging Deep Learning and CAP Research:** Leveraging a convolutional neural network (CNN)-based solution, our method effectively bridges the gap between deep learning capabilities and CAP research. This incorporation is distinctive, as it diverges from traditional methods in CAP studies, where the application of deep learning techniques has been infrequently explored.

1.3 Research Overview

In the first part of this research we propose an end-to-end fully automatic CNN-based method for classifying CAP phases. For this goal, we present an algorithm consisting of three main stages (Figure 1.1). Firstly, we analyze each second of the long EEG signal as a distinct prolonged 1D-EEG segment, s_i. Akin to human analysis, which considers the vicinity and context of the EEG signal, our method incorporates extended segments with long window sizes to extract and leverage contextual information from the signal. Next, due to the non-stationary nature of the signals we transform each time segment into its time-frequency representation (TFR), X_i. This representation is highly suitable for classifying CAP phases, and as a 2D representation it allows us to utilize a CNN-based architecture effectively afterward. We investigated the usage of several time-frequency transforms and show that the commonly-used spectrogram, which relies on the short-time-Fourier-transform (STFT), is inadequate for the CAP classification task, likely due to its limited time-frequency resolution [30]. Alternatively, we demonstrate that adopting Wigner-Ville-distribution (WVD)-based transformations, which in many cases capture the time-varying frequency content of the signals more accurately [31], significantly enhances the results. At the third stage of the algorithm, the obtained TFRs are treated as 2D images and fed into a convolutional neural net-

work that classifies them as either A-phase or B-phase samples. Our neural network is based on the common ResNet18 architecture [32], with the necessary adjustments for TFRs classification. Finally, to improve the generalization of our model and reduce overfitting [33–35], we propose a data augmentation composition specially designed to preserve the reasonable structure of the EEG time-frequency representation.

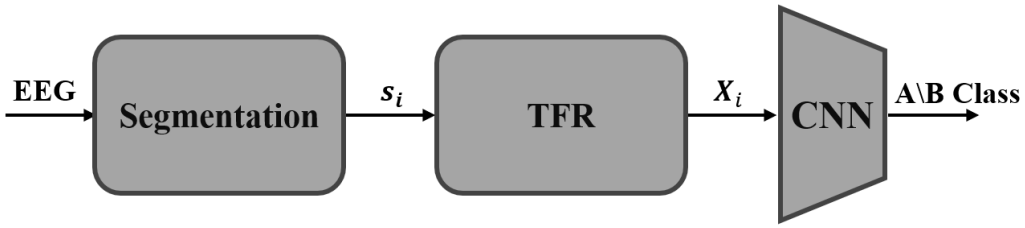

Figure 1.1: Proposed CAP Classification Workflow: The EEG signal is segmented, transformed into a 2D time-frequency representation (TFR), and fed into a convolutional neural network (CNN) architecture for A/B-phase classification.

The second part of this thesis focuses on a comprehensive evaluation of the proposed method along with an ablation study execution. Testing of the proposed method was conducted over the PhysioNet's public Cap Sleep Database (CAPSLPDB) [2, 36], considered a benchmark database for CAP identification and classification research. The performance evaluation was conducted both on balanced and unbalanced datasets, and a thorough experimental study shows our algorithm achieves state-of-the-art performance on the CAP-Sleep-Database, reaching an accuracy of 77.5% on a balanced test-set and 81.8% when evaluated on an unbalanced test-set. At the ablation study section we have conducted a series of experiments to assess the effect and contribution of the different components of the proposed algorithm. This research involved examining the effect of using various window sizes to improve outcomes and showed that involving the sequential information of the EEG signal is crucial to classify the CAP. Our experimental results demonstrate that we achieve state-of-the-art performance even with the compact ResNet18 architecture [32], which can be efficiently run on-device.

1.4 Organization

This research is structured as follows. In Chapter 2, we provide the scientific background and cover the database description, including insights into the dataset creation process. Chapter 3 introduces the primary contribution—our end-to-end CNN-based method for classifying CAP phases. This chapter thoroughly details the various stages of the algorithm. Chapter 4 focuses on a comprehensive evaluation and an ablation study, utilizing PhysioNet's CAP Sleep Database. We assess the algorithm's performance on balanced and unbalanced datasets and conduct experiments to understand

the contribution of different components. Chapter 5 summarizes findings, concludes the thesis, and suggests future directions for further research.

Chapter 2

Preliminaries

2.1 Cyclic Alternating Patterns

Sleep is as an essential part of life for many species including humans. According to the guidelines set by the American Academic of Sleep Medicine (AASM) [37], sleep is typically classified into five stages that characterize sleep's macrostructure. These stages include Wakefulness (W), Rapid-Eye-Movement (REM), and Non-Rapid-Eye-Movement(Non-REM), which consists of three interior stages (S1- S3). The alternation of Non-REM and REM sleep constitutes the sleep cycle, where each sleep cycle has a duration of approximately 90 minutes [38]. Stage S1 is commonly characterized by mixed frequency activity with low amplitude and constitutes between 2% and 5% of the total sleep time. The energy in the lower frequencies increases during the second sleep stage (S2) that occupies between 45% and 55% of the total sleep. High-amplitude waves with low frequency characterize stage S3. During the REM sleep, the cerebral activity increases, presenting mixed wave frequencies with low amplitude. Understanding the structure of sleep is essential, and this knowledge is often harnessed through polysomnography (PSG), a comprehensive sleep study and diagnostic tool in sleep medicine [39]. The PSG monitors many body functions, including brain activity, eye movements, muscle activity and heart rhythm. Specifically, the electroencephalogram (EEG) is one of the most commonly used techniques in this field for measuring transient and phasic events in the brain electrical activity. The EEG records the alternating electrical activity at the scalp surface using conductive media and metal electrodes [40]. The EEG power spectrum, calculated by the Fourier transform, is typically categorized in four bands - delta (0.5-4 Hz), theta (4-8 Hz), alpha (8-13 Hz) and beta (13-30 Hz) [41].

In 2001, the concept of cyclic alternating patterns (CAP) was introduced [2] to characterize the microstructure of sleep. CAP represents a periodic EEG activity that occurs during Non-REM sleep and is characterized by cyclic sequences of cerebral activation (A-phase) followed by periods of deactivation (B-phase). An A-phase period and the following B-phase period define a CAP cycle; at least two CAP cycles are required

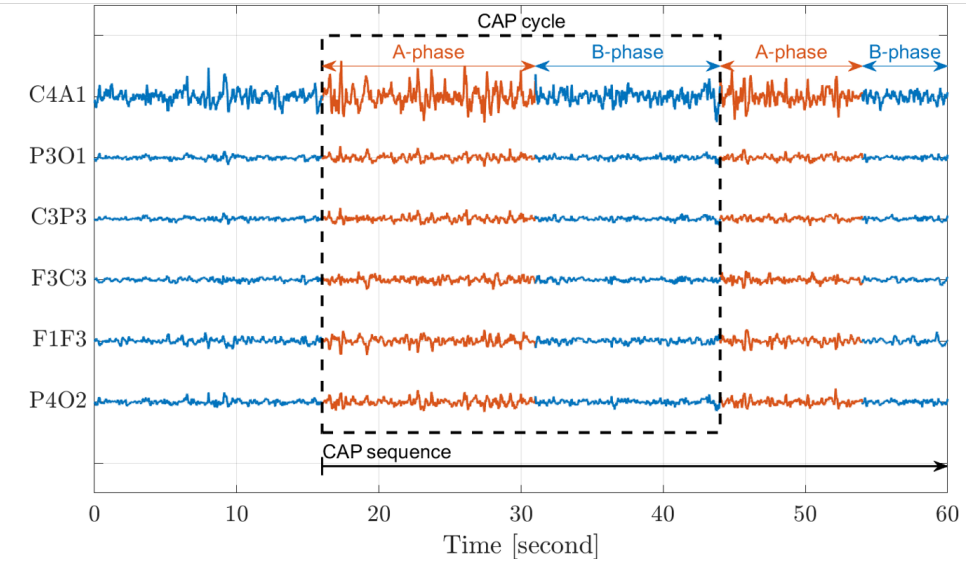

Figure 2.1: A demonstration of the cyclic alternating pattern (CAP) in sleep.

to form a CAP sequence (see Figure 2.1).

Research has shown that beyond being a physiological phenomenon, CAP has a pivotal role for sleep quality assessment [3]. In particular, the CAP rate is the most widely exploited microstructural parameter for clinical purposes. The CAP rate quantifies sleep instability and it is defined by the ratio of the total CAP time in Non-REM sleep, to total Non-REM sleep time. CAP rate increases when sleep is disturbed by internal or external factors, and its variations reflect the perception of sleep quality, with higher values of CAP rate related with poorer sleep quality. Hence, there is a evident correlation between autonomic functions, behavioural activities and CAP [38]. Since a poorer sleep quality is related to higher values of CAP rate, CAP is considered as the EEG marker of sleep instability [3]. Furthermore, CAP cycles have been associated with many sleep disorders such as sleep apnea, bruxism, insomnia, periodic limb movements, and nocturnal frontal lobe epilepsy [4]. Therefore, scoring CAP becomes crucial in characterizing and diagnosing these pathologies.

In the characterization of the CAP structure, the B-phase is recognized as the underlying background rhythm of the signal, contrasting with the A-phase that denotes transient and prominent events. The A-phase events are divided into three interior subtypes [42]:

- *A1* is dominated by slow varying waves (low frequencies, 0.5Hz - 4Hz) with a high amplitude about the typical background, B-phase.

- *A3* is characterized by increasing in frequency (8Hz - 12Hz) along with decreasing in the amplitude.

- *A2* is a combination of both *A1* and *A3*.

The A1 subtype is the most prevalent, constituting a range from 46.6% of A-phase events in elderly individuals to 85.5% in children [38].

This work focuses on the binary classification of Non-REM sleep into its A and B CAP phases.

2.2 Materials and methods

2.2.1 Database Description

The proposed method was developed and evaluated over the publicly available CAP sleep database (CAPSLPDB) [2, 36], considered a benchmark for CAP research. The database contains a collection of polysomnographic recordings registered at the Sleep Disorders Center of the Ospedale Maggiore of Parma, Italy. It includes data from a diverse group of 108 patients: healthy individuals and those with various pathological conditions, such as bruxism, insomnia, and others. Each record includes three or more EEG signals and a series of other indicators, such as electrooculogram (EOG), chin and tibial electromyogram (EMG), and ECG signals. Additionally, the database includes accurate CAP annotations corresponding to each second of the signals. The left side of Table 2.1 summarizes the sample composition per subject in the database. The database exhibits a highly imbalanced distribution, with a significantly higher occurrence of B-phase samples than A-phase events.

2.2.2 Performance Measures

To assess the classification performance under various configurations, we calculated several metrics. These included the number of correctly identified A-phase events (true positives, t_p), the number of correctly recognized B-phase samples (true negatives, t_n), as well as the count of samples incorrectly classified as A-phase (false positive, f_p) or as B-phase (false negative, f_n). Based on these metrics, we computed accuracy (ACC), precision (PRE), recall (REC), specificity (SPE) and F_1-score (F1) using the following expressions:

$$ACC = \frac{t_p + t_n}{t_p + t_n + f_p + f_n}\,, PRE = \frac{t_p}{t_p + f_p}\,. \tag{2.1}$$

$$REC = \frac{t_p}{t_p + f_n}\,, SPE = \frac{t_n}{t_n + f_p}\,. \tag{2.2}$$

$$F1 = \frac{2 \cdot t_p}{2 \cdot t_p + f_p + f_n}\,. \tag{2.3}$$

Regarding the data splitting for evaluation, in contrast to prior CAP studies [10–12, 15, 43] that employed K-fold cross-validation, we adopted the standard practice of dividing the data into three disjoint subsets: training, validation, and test sets as seen in various prominent works [23, 44–46]. The distribution was approximately 80% for training and

Table 2.1: Total number of samples per healthy subject in the original CAP sleep database (CAPSLPDB) and the corresponding number of samples selected for this study's dataset. The original CAPSLPDB shows a significantly higher prevalence of B-phase samples than A-phase samples. In contrast, the dataset utilized in this study exhibits a balanced distribution of both A-phase and B-phase classes.

Subject	**CAPSLPDB (Unbalanced)**					**Our Dataset (Balanced)**				
name	A_1	A_2	A_3	Total A	B	A_1	A_2	A_3	Total A	B
n1	2217	747	1122	**4086**	**21804**	2063	703	1046	**3812**	**3812**
n2	1115	590	783	**2488**	**12122**	1036	552	693	**2281**	**2281**
n3	611	597	891	**2099**	**15451**	550	556	830	**2281**	**2281**
n4	986	356	848	**2190**	**15030**	928	323	797	**2048**	**2048**
n5	2854	328	620	**3802**	**18158**	2673	314	586	**3573**	**3573**
n6	1871	970	1401	**4242**	**17268**	1723	905	1280	**3908**	**3908**
n7	1616	564	479	**2659**	**17501**	1508	525	438	**2471**	**2471**
n8	949	465	1868	**3282**	**17028**	914	421	1752	**3087**	**3087**
n9	1036	377	676	**2089**	**18341**	959	363	641	**1963**	**1963**
n10	1484	326	829	**2639**	**13351**	1385	282	785	**2452**	**2452**
n11	1724	583	796	**3103**	**15377**	1640	539	734	**2913**	**2913**
n12	1064	153	573	**1790**	**18040**	986	139	515	**1640**	**1640**
n13	1628	1037	1017	**3682**	**14078**	1532	985	955	**3472**	**3472**
n14	1035	1234	1209	**3478**	**15902**	950	1118	1126	**3194**	**3194**
n15	1449	1046	1244	**3739**	**18461**	1345	967	1159	**3471**	**3471**
n16	2247	1125	837	**4209**	**17841**	2110	1041	786	**3937**	**3937**

10% each for validation and test sets.

2.2.3 Dataset Creation

For this study, we built our dataset using the recordings of the sixteen normal (no pathology) patients, where a single EEG channel was utilized per participant (either the C4-A1 or the C3-A2 channel). Construction of the dataset from the long full-night EEG signals was done through several steps. To ensure the spread of the samples in the train, validation, and test sets throughout the entire recording, each full-night EEG signal was divided into non-overlapping 300-second segments. The first 240 seconds of each segment were assigned to the training set, the subsequent 30 seconds were allocated to the validation set, and the remaining 30 seconds were designated as the test set. Subsequently, since our proposed algorithm takes each second as a prolonged time window comprising contextual information, removing the seconds at the edges of the resulting segments is essential to achieve a complete separation between the train, validation, and test sets. Ultimately, an appropriate percentage of B-phase samples were randomly removed per recording to achieve a balanced dataset, i.e., the number of A-phase and B-phase samples is equal (see right side of Table 2.1). We acknowledge that our data-splitting approach results in different segments from the same recording

appearing in both the training and test sets, which could be perceived as a form of information leakage between the train and test sets. However, given the limited number of normal patients available (16 in total), this approach allows the model to learn from a diverse range of data while avoiding evaluation on an extremely small subset of subjects, as would happen with a complete separation. Nevertheless, if a significantly greater number of patient recordings were accessible, a stricter separation between training and testing would be preferable, ensuring test recordings remain entirely unseen during training.

Chapter 3

Classification of Cyclic Alternating Pattern using Convolutional Neural Networks

In this chapter, we propose a method for the classification of cyclic alternating pattern (CAP) using convolutional neural networks (CNNs). Our approach is driven by three key considerations: incorporating contextual information, utilizing distinct features of A-phase events, and leveraging the high-performance capabilities of deep learning through a CNN-based architecture. Our method comprises three main building blocks, each aligned with these factors. This chapter details each component, including data segmentation, downsampling, and the exploration of different time-frequency representations. Additionally, it discusses the deep learning architecture, normalization, and specialized time-frequency augmentations designed to enhance model generalization. The chapter provides a structured overview of the methodology, paving the way for subsequent discussions and evaluations.

In Section 3.1, we delve into a more detailed explanation of the proposed algorithm, clarifying the foundational considerations it is based on. The preprocessing stage is expanded in Section 3.2, while time-frequency analysis techniques that are relevant to the algorithm are presented and detailed in Section 3.3. Section 3.4 focuses on the neural network architecture and further associated processing. The chapter is summarized in Section 3.5.

3.1 Proposed Method

Our method consists of three key components, which are driven by three primary considerations:

- **Context** - Incorporating contextual information of the signal to make the prediction more analogous to the human diagnosis, which inherently involves close

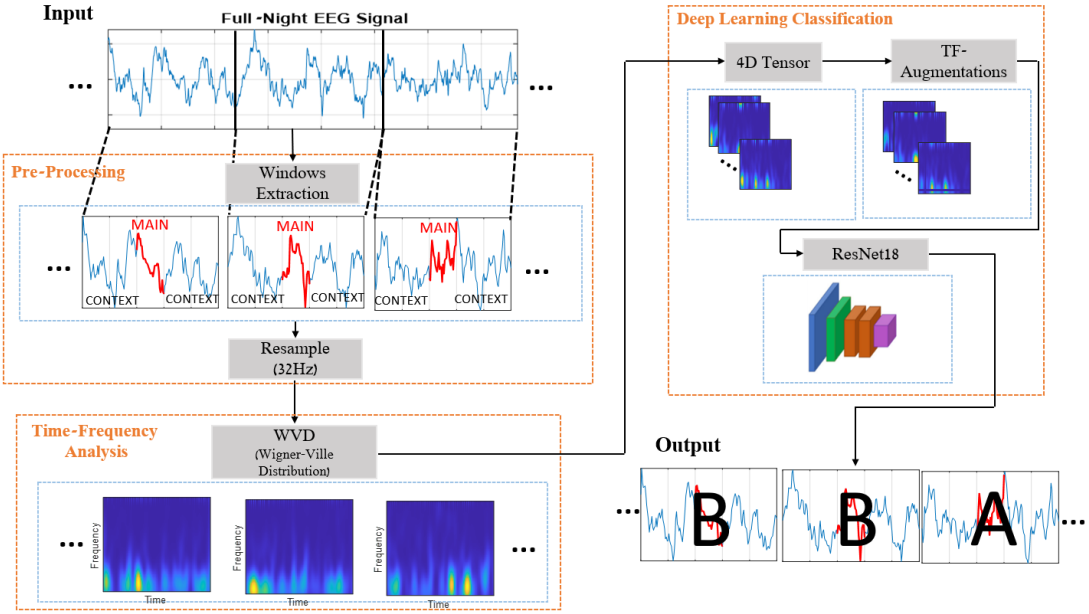

Figure 3.1: Proposed method scheme.

vicinity analysis.

- **CAP prior knowledge** - Utilizing the distinct features of the A-phase events, which are characterized by higher energy levels and high-frequency spectral content compared to the B-phase background.
- **Deep learning** - Employing a CNN-based architecture as a classifier to leverage the CNNs high-performance capabilities.

The proposed method consists of three main building blocks (Figure 3.1), which align with these factors. The first component involves pre-processing, where each analyzed 1-second EEG is treated as a more extended time window. This window contains the central part we want to classify, along with the near vicinity of the signal that is added to provide contextual information. Each 1D-EEG segment is transformed into a 2D time-frequency matrix in the second stage. This representation captures the non-stationary signal's energy and spectral content, which vary over time. From this point, the obtained 2D-time-frequency representations are treated as images, and the proposed method adopts a deep learning framework. In line with that, the received images are stacked into 4D tensors, normalized, and augmented to preserve their time-frequency structure. The processed images are finally fed into a CNN-based architecture for training in a supervised manner. Next, we detail each of these components.

3.2 Pre-processing

To address the extended length of full-night input signals, which typically last between 9 to 10 hours, we initiate the process with data segmentation. The annotations of the

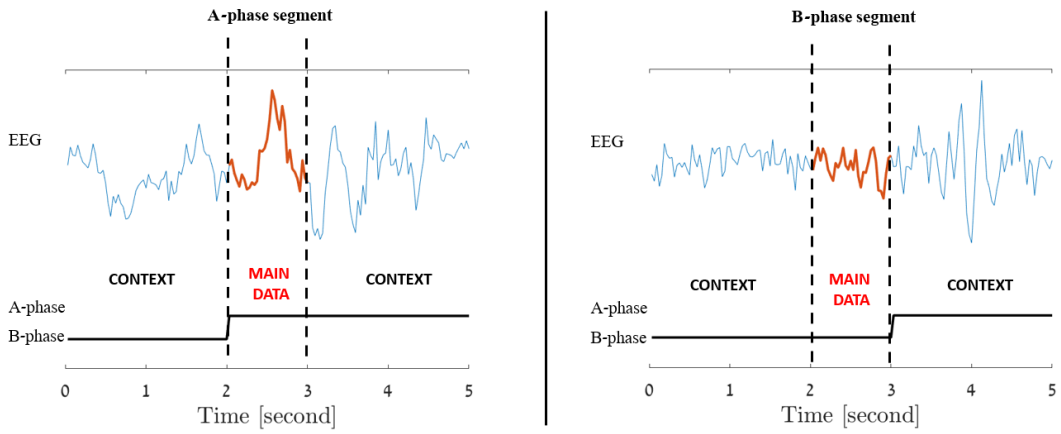

Figure 3.2: Incorporated Contextual Information: Each signal second extends to include preceding and subsequent seconds, labeled by its central (main data) second. The figure illustrates A-phase (left) and B-phase (right) 5 s data segments.

data pertain to each second of the signal; however, considering the sequential nature of EEG data, the analysis of each signal second involves the inclusion of preceding and subsequent seconds, resulting in prolonged EEG segments that incorporate the contextual information of the signal. In this work we evaluated different window lengths, ranging from the plain 1-second windows lacking contextual information to 11-second ones at most. Utilizing even longer windows appears to over-emphasize the contextual information and poses storage and computational efficiency challenges.

Since each EEG segment is an extension of its central second, for all window lengths, the label assigned to a segment is determined solely by the label of its central second, regardless of the labels of its other components. For instance, a 5-second segment composed of two seconds of B-phase followed by three seconds of A-phase would be labeled as an A-phase segment due to its central A-phase second, while on the other hand, a similar 5-second segment made up of three B-phase seconds at the beginning followed by two A-phase seconds, would be designated as a B-phase segment, due of its B-phase center. These two scenarios are illustrated in Figure 3.2.

Ultimately, due to the different sampling rates of the signals at CAPSLPDB, which vary across recordings between 100 Hz to 512 Hz, we downsampled each segment to 32 Hz as a significantly lower sampling rate that preserves the frequency content relevant to the CAP phases. Thus, an identical resolution at the analysis is obtained, and the complexity of calculations is significantly reduced.

3.3 Time-Frequency Analysis

Time-frequency analysis is applied to the signals to reveal and exploit both spectral structure and temporal changes of the EEG segments, which is essential for the distinction between A and B phases due to their non-stationary nature. Additionally,

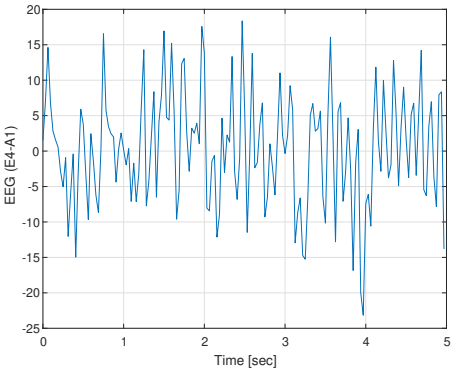

(a) 1D-EEG signal

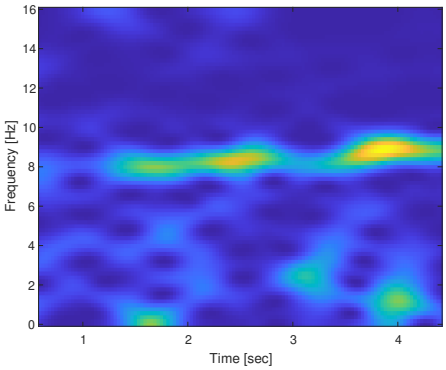

(b) Spectrogram

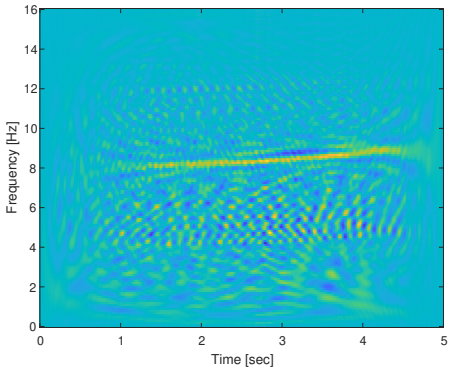

(c) Wigner-Ville distribution

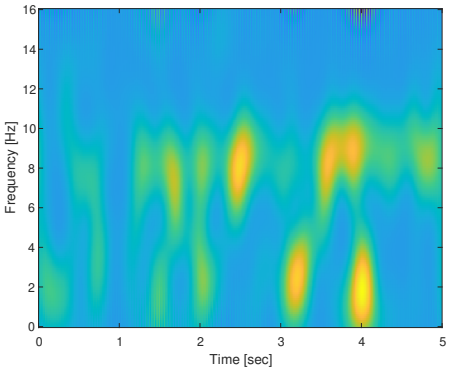

(d) Smoothed pseudo-Wigner-Ville distribution

Figure 3.3: Example of (a) 5 s 1D-EEG segment from channel E4-A1 and its corresponding time-frequency representations (TFRs): (b) spectrogram (SPEC), (c) Wigner–Ville distribution (WVD), and (d) smoothed pseudo-Wigner–Ville distribution (SPWVD). The WVD exhibits a distinct energy concentration when compared to SPEC, albeit with the tradeoff of noticeable cross-term patterns.

the transition of the 1D time segments to 2D time-frequency images allows using a CNN-based classifier. In this study, we explored the use of several time-frequency transformations, including Spectrogram (SPEC), Wigner-Ville Distribution (WVD), and Smooth Pseudo-Wigner-Ville Distribution (SPWVD). Definitions and additional details for these representations are given in the following paragraphs.

3.3.1 Spectrogram (SPEC)

The spectrogram is widely acknowledged as a prevalent method for analyzing time-varying and non-stationary signals. The spectrogram definition is based on the short-time Fourier transform (STFT), as for a signal, $x(t)$, the STFT is defined as:

$$X(t,f) = \int_{-\infty}^{\infty} x(t_1)h^*(t_1 - t)e^{-j2\pi f t_1}\, dt_1, \tag{3.1}$$

where $h(t)$ is a window function centered at time t. The window function cuts the signal just close to the time t, and the Fourier transform will be an estimate locally around this time instant.

The spectrogram, $S_x(t,f)$, is formulated as the squared magnitude of the STFT:

$$S_x(t,f) = |X(t,f)|^2\,. \tag{3.2}$$

The spectrogram is the most widely known and commonly used time-frequency transform [31]. It is well understood, easily interpretable, and has fast implementations, e.g., fast-Fourier-transform. However, its drawbacks are the limited and fixed resolution in time and frequency which is determined by the length of the window $h(t)$ [30].

3.3.2 Wigner-Ville Distribution (WVD)

The Wigner-Ville Distribution of a signal $x(t)$ is given by:

$$W_x(t,f) = \int_{-\infty}^{\infty} x(t + \frac{\tau}{2})x^*(t - \frac{\tau}{2})e^{-j2\pi f\tau} d\tau\,. \tag{3.3}$$

The WVD has the best possible concentration in the time-frequency domain [47], and in particular, can attain a perfect localization for pure frequency-modulated signals [48].

However, the notable drawback of the WVD is known as the cross-terms (CT) [47]. These artifacts arise when the signal contains a mixture of several signal components, which significantly reduces the readability of the time-frequency representation. The origin of CT lies in the non-linear nature of the WVD transform, which causes the superposition of several components to generate not only the desired auto-terms (AT) components but also CT. One of the methods to address this problem is the Smooth-Pseudo Wigner-Ville Distribution method, as explained immediately.

3.3.3 Smoothed-Pseudo Wigner-Ville Distribution (SPWVD)

The smooth-pseudo Wigner-Ville distribution of a signal $x(t)$ can be formulated as the 2-dimensional convolution of the Wigner-Ville distribution, $W_x(t,f)$, with a low-pass-nature kernel, $\Phi(t,f)$:

$$W_x^{sp}(t,f) = W_x(t,f) * *\Phi(t,f), \tag{3.4}$$

where $**$ represents a 2D convolution.

Equation (3.4) can be expressed explicitly:

$$\begin{aligned} W_x^{sp}(t,f) &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} x(u+\frac{\tau}{2})x^*(u-\frac{\tau}{2})\cdot ... \\ &\qquad ... \cdot \Phi(\nu,\tau)e^{j2\pi(\nu t - f\tau - \nu u)}dud\tau d\nu \\ &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} A_x(\nu,\tau)\Phi(\nu,\tau)e^{-j2\pi(f\tau-\nu t)}d\tau d\nu, \end{aligned} \tag{3.5}$$

where $A_x(\nu,\tau)$ is called the ambiguity function (AF) and is defined as:

$$A_x(\nu,\tau) = \int_{-\infty}^{\infty} x(t+\frac{\tau}{2})x^*(t-\frac{\tau}{2})e^{-j2\pi\nu t}dt\,. \tag{3.6}$$

Equation (3.4) formulates the SPWVD as a filtered version of the WVD, whereas the last term in (3.5) demonstrates that this filtering is achieved through the multiplication of the AF with the low-pass-nature kernel, $\Phi(\nu,\tau)$. This observation can be rationalized considering that the AF can be viewed as a time-frequency (TF) auto-correlation function of the original signal $x(t)$ [49]. As such, it exhibits most properties of a correlation function, including that its modulus is maximum at the origin [50]. As for the multi-component signal case, the total AF consists of both auto-terms neighboring the origin of the time-frequency plane ($\nu = 0$ and $\tau = 0$) and cross-terms which are mainly located at a time-frequency distance from the origin. This distance depends directly on the separation in time and frequency of the individual components of the signal. Considering this perspective, it is intelligible that low-pass filtering of the AF means suppressing the cross-terms alongside preserving the desired auto-terms. In that way, the SPWVD is an effective method to deal with the WVD cross-terms drawback.

Figure 3.3 demonstrates the abovementioned forms. It shows an example of a 5-second EEG signal taken from CAPSLPDB (patient 'n1') and its different time-frequency representations. The increased spectral content of the signal, which exists at the low frequencies (up to 2Hz) and 8Hz, is reflected from all the different representations. However, it can be seen that Wigner-based representations have a prominent better resolution in time and frequency compared to the spectrogram. Additionally, it presents the interference in visualization caused by the cross-terms when using WVD,

and the mitigation to that problem comes by employing an SPWVD technique.

3.4 Deep Learning Architecture

In this phase, the 2D time-frequency images generated in previous stages are stacked into tensors and employed as training data for a convolutional neural network. Table 3.1 outlines the hyperparameters employed in the training process. The chosen model details and further extensions that were executed, such as normalization and augmentation, will be discussed subsequently.

3.4.1 Model

We adopted the widely-used and well-established ResNet18 architecture [32], which is renowned for its effectiveness in visual classification tasks [51]- [52]. To align the ResNet18 model with our specific framework, we modified the first and last layers. These modifications were required to accommodate gray-scale images as input and to enable a binary classification at the output. Furthermore, we trained the model from scratch, considering the substantial disparities between the training data used to train

Hyperparameters	Value
Batch size	256
Loss functions	cross entropy
Optimizer	SGD
Learning rate	0.001
Momentum	0.9
Epochs	40
Dropout	No

Table 3.1: Hyperparameters used in the proposed framework.

the ResNet18 originally, ImageNet [23], which is composed of natural images, and our distinct time-frequency "images."

3.4.2 Normalization

Traditionally, the input data, x_i, of neural networks is normalized to be, $\tilde{x_i}$, with zero-mean and of unit standard deviation [53, 54], namely:

$$\tilde{x_i} = \frac{x_i - \bar{x_i}}{\sigma_i} \tag{3.7}$$

where $\bar{x_i}$ is the mean of x_i and σ_i is its standard deviation.

However, to preserve the difference in energy levels between A-phase and B-phase samples, in this work, we selected to divide each input sample constantly by the *mean*

standard deviation of the train set samples, namely:

$$\tilde{x_i} = \frac{x_i - \bar{x_i}}{\bar{\sigma}}, \bar{\sigma} = \frac{1}{N} \sum_{i \in X_N} \sigma_i \tag{3.8}$$

where X_N denotes the set of all train samples, and $\bar{\sigma}$ is the mean standard deviation of the this set.

3.4.3 Augmentations

To generalize the learned model [33] and to reduce over-fitting to train data [34], we employed a series of augmentations for every batch of data loaded. Initially, we evaluated several traditional augmentations commonly used in computer vision tasks, including color-jitter, rotation, and flips, as outlined in [35]. However, these augmentations led to highly inadequate performance, likely due to their strong correlation with natural images, different from the time-frequency "images" being analyzed [55]. Subsequently, we explored the usage of augmentations designed explicitly for our time-frequency image data. Ultimately, we investigated two main augmentation types:

1. Time-shifts: We employed random time-shifts by applying horizontal random cropping to the training data samples. The cropping was restricted to the horizontal axis, i.e., the time domain, to maintain the spectral information of the signals and preserve the distinction between the different phases of CAP, which differ significantly in their spectral characteristics."

2. Time-Frequency augmentations (TF-Aug.): A specialized selection of augmentations was utilized to characterize the time-frequency representations effectively. These augmentations were repeatedly applied to each data sample before inputting the neural network. The selected augmentations are:

    - **Noise**: Additive white Gaussian noise (AWGN) with a uniformly distributed standard deviation. Adding noise was specified in [56] as an appropriate and effective augmentation for EEG signals
    - **Gaussian blur:** The time-frequency images were blurred using a Gaussian kernel. This augmentation was randomly applied to the input samples with a probability of $p = 0.5$, meaning that approximately half of the images underwent blurring.
    - **SpecAugment** [57]: A commonly used method for augmenting spectrogram and other time-frequency representations, typically for speech recognition tasks. The augmentation is primarily based on applying random masks to certain frequency bands and time steps in the spectrogram. In this study, we randomly blocked bands up to 5% of image width for time and 3% of image height for frequency.

- **Crop and Resize:** To imitate extended temporal CAP events, we randomly cropped the images vertically and then resized them back to their original size, slightly stretching the temporal duration of CAP events.

3.5 Summary

This chapter presents a method for CAP classification using convolutional neural networks. The approach incorporates contextual information, leverages distinctive A-phase features, and employs a CNN-based ResNet18 architecture for training. It involves pre-processing, transforming 1-second EEG segments into 2D time-frequency matrices, and utilizing various time-frequency representations, such as spectrogram, Wigner-Ville distribution and smoothed pseudo-Wigner-Ville distribution. The deep learning architecture stacks these images into tensors and employs a modified ResNet18 model. The chapter covers normalization and specialized time-frequency augmentations, including noise, Gaussian blur, SpecAugment, and Crop and Resize, to enhance model generalization. By exploring different representations and optimizing the pre-processing data pipeline, we expect to achieve better performance in CAP classification.

Chapter 4

Experiments and Numerical Results

This chapter provides a comprehensive evaluation of the proposed algorithm's performance in classifying cyclic alternating pattern (CAP) phases. The study encompasses an in-depth analysis of various factors influencing the algorithm's accuracy, including the choice of time-frequency representations, the impact of contextual information incorporation, and the effectiveness of data augmentation strategies. Furthermore, we investigate the classification for each class separately, extracting insights into the inherent characteristics of the trained model. Additionally, this chapter introduces a comprehensive comparative analysis, comparing our approach against recent studies in the field.

Section 4.1 elaborates the experiments conducted for evaluating the proposed method. An ablation study to explore the impact of the different building blocks of the algorithm is provided in Sections 4.2, 4.3 and 4.4, while in Section 4.5 we delve into the A-phase inter-classes classification. In Section 4.6 we compare our method to other recent studies and summarize the whole chapter in Section 4.7.

4.1 Experiments

We evaluate the performance of the proposed algorithm on the dataset described above and compare it to the results of prior studies. An ablation study was also carried out through a series of experiments to investigate the following aspects:

1. The influence of utilizing various time-frequency representations as input to the classifiers.
2. The impact of incorporating the EEG signal context information by using segments with an increased duration.
3. The determination of appropriate data augmentations strategies for analyzing EEG signals within the proposed framework.

Table 4.1: Accuracy results (%) for different TFRs and segmentation lengths. In each cell, the upper result refers to the validation set performance, while the lower result refers to the test set performance. The highest results are highlighted within each column, demonstrating the superiority of Wigner-based representations over spectrogram. Additionally, the impact of increased window size is observable.

Method	**Window Size**					
	1 Sec	3 Sec	5 Sec	7 Sec	9 Sec	11 Sec
SPEC	65.65	70.97	71.87	73.07	74.39	75.84
	65.26	69.16	68.88	70.48	71.89	74.10
WVD	59.07	**74.53**	75.75	76.89	**78.50**	**78.46**
	57.93	**73.36**	73.35	74.85	**76.64**	**77.54**
SPWVD	**66.75**	72.92	**77.10**	**77.74**	78.26	77.38
	66.19	71.65	**74.63**	**75.24**	75.85	75.64

In all experiments, we trained the deep networks using Google-Colab. The training has lasted between a few hours to ∼20 hours, depending mainly on the chosen window size.

4.2 Time-Frequency Representations

In the first experiment, we compared the three time-frequency representations: spectrogram (SPEC), WVD, and SPWVD. In this experiment, only random time-shifts were applied without further augmentations. For the spectrogram, we used the Hanning window with support of 20% of single data sample length and maximal overlap between subsequent windows; namely, the window moves only 1 sample each time. As For the WVD and SPWVD, we used the built-in MATLAB function with its default parameters. Figure 4.1 presents the accuracy of the different representations using increasing window size from 1 second to 11 seconds (incremented by 2 seconds).

The results in Figure 4.1 and Table 4.1 clearly show that utilization of Wigner-based transforms (WVD and SPWVD) is much better compared to a Fourier-based spectrogram (SPEC). This is evident in the higher accuracy obtained by WVD and SPWVD for all window sizes greater than 1-second, with only SPEC achieving better accuracy compared to WVD for the 1-second window size (65.65% compared to 59.07% for WVD). Nevertheless, SPWVD still outperforms SPEC at the 1-second case, with an accuracy of 66.75%. As mentioned above, the prominence of Wigner-based transforms over the spectrogram can likely be attributed to the limitations of STFT in terms of time-frequency resolution. In contrast, WVD provides an optimal concentration in the time-frequency domain.

In general, the differences between WVD and SPWVD are negligible. This similarity reveals that cross-terms, which substantially hinder human interpretability, are considered tolerable by the trained model, which learns to deal with these components

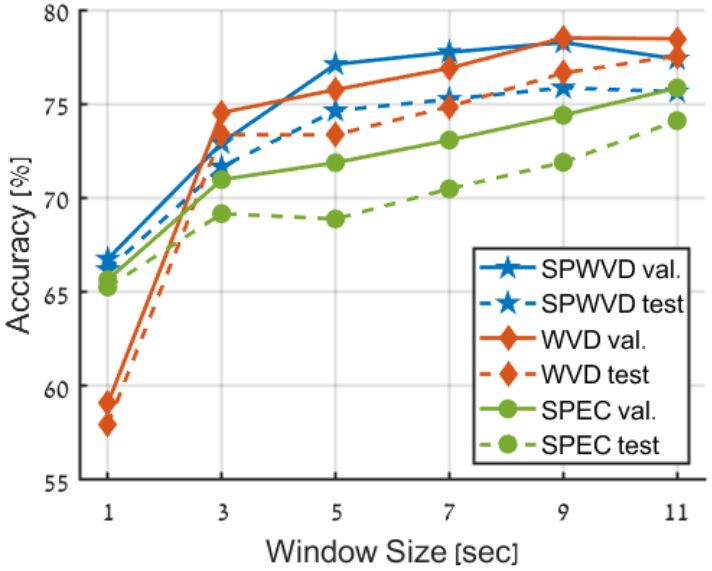

Figure 4.1: Comparison of performance achieved using different time-frequency representations (TFRs) and window sizes. The blue line corresponds to the SPWVD transform, the red line to the WVD, and the green line to the spectrogram (SPEC). The validation and test data are depicted as solid and dashed lines, respectively.

during the training process, and may even leverage them as supplementary features. Yet, the 1-second window case is an exception to this similarity, where WVD demonstrates a notably lower accuracy.

4.3 Contextual information

Additional valuable insight depicted from Figure 4.1 is that increasing the window size and the additional contextual information it provides significantly improves the accuracy across all time-frequency representations. The improvement in accuracy reaches approximately 10% for STFT and SPWVD to roughly 20% for WVD. Nevertheless, this trend seems to plateau at the increment from 9-second to 11-second window sizes for WVD and SPWVD, possibly due to an excessive proportion of contextual information about the central primary data. Further increasing the window size beyond 11 seconds was not explored in this study due to the required prolonged training time.

4.4 Data Augmentations

In this experiment, three data augmentation techniques were compared to determine an appropriate augmentation strategy for the proposed framework. Additionally, the no-augmentation case was tested as a benchmark. In all cases, the SPWVD was utilized as the time-frequency representation. The evaluated augmentation types were as follows:

Table 4.2: Accuracy results (%) for different data augmentations and extensions of the dataset. In each cell, the upper result refers to the validation set performance, while the lower result refers to the test set performance. In each column, the highest results are highlighted. The results demonstrate the effectiveness of integrating data augmentation. Notably, the highest accuracy (78.72%) was achieved using the proposed TF-augmentations with a 9-second window size.

Dataset's composition	**Window Size**					
	1 Sec	3 Sec	5 Sec	7 Sec	9 Sec	11 Sec
Basic dataset	66.70	**73.21**	75.48	76.49	77.63	77.22
	65.57	71.19	72.90	73.57	75.77	74.06
TF-augmentations	67.55	72.03	76.02	77.53	**78.72**	77.74
	67.39	70.08	73.82	74.73	75.76	74.24
Time-shifts	66.75	72.92	**77.10**	**77.74**	78.26	77.38
	66.19	**71.65**	**74.63**	**75.24**	**75.85**	**75.64**
TF-augmentations & time-shifts	**67.69**	72.10	76.52	77.49	78.08	**77.94**
	67.36	71.07	74.20	75.08	75.11	74.87

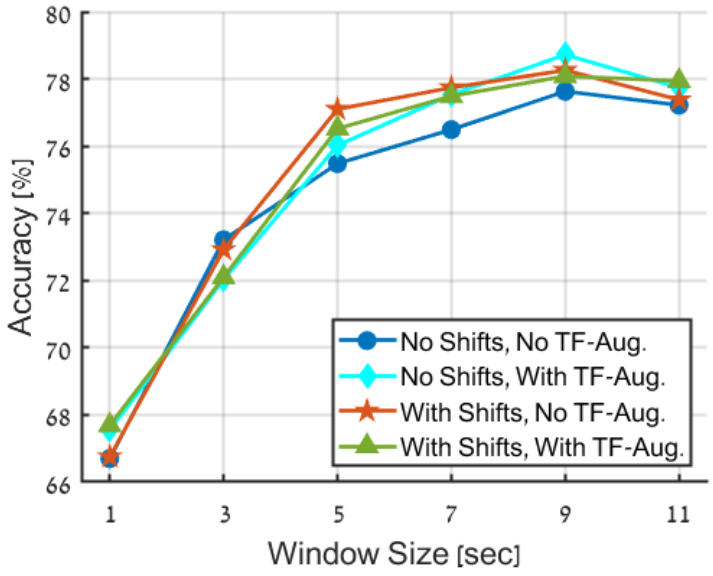

Figure 4.2: Accuracy comparison of various data augmentation techniques. The figure shows the performance of four strategies: no data augmentation (blue), proposed TF-augmentation only (cyan), random time-shifts only (red), and employment of both time-shifts and TF-augmentation (green).

Table 4.3: True-positive rate (%) per A-phase sub-types and B-phase.

Predicted	**True**			
	A1	A2	A3	B
A	1422	561	782	2914
B	343	99	400	693
TPR [%]	**80.6**	**85**	**66.2**	**80.8**

- *TF-augmentation*: TF-augmentations are applied solely. As described above, these augmentations are designed to maintain the time-frequency structure.
- *Random time-shifts*: In this case, the original dataset is augmented by incorporating random time shifts into its samples.
- *TF-augmentations and random time-shifts*: Both TF-augmentations and random time-shifts are applied to the dataset.

The results presented in Figure 4.2 and Table 4.2 highlight the positive effect of augmenting the primary dataset. The improvement in the accuracy of the trained model, with relation to the non-augmentation case, is observed for window sizes larger than 3 seconds and ranging from 1%-2%.

When considering the comparison between the different augmentation techniques, the disparities in accuracy results are inconsequential. However, from the perspective of overall system considerations, it is advantageous to utilize TF-augmentations instead of random time-shifts, since the former does not necessitate the production of extended signals, resulting in improved storage efficiency and reduced computational complexity. It is also noted that using both TF-augmentations and time-shifts does not result in further performance improvements and is, therefore, unnecessary.

4.5 A-Phase Detection

The confusion matrix of the resulting model (for a 9-second segment, SPWVD transform, and TF-augmentations) is shown in Figure 4.3, which reveals that the true-positive-rate (TPR) is similar for the A-phase and B-phase classes, with TPR values of 76.7% and 80.8%, respectively. Table 4.3 presents performance per A-phase subtypes. As the classification is binary, the TPR per subtype refers to the instances that were classified as A-phase generally. The results show that the TPR of A2 and A1 subtypes is significantly higher than the TPR of A3 (85% and 80.6% for A2 and A1, versus 66.2% for A3).

Considering the characteristics of the A-phase different subtypes detailed above, this finding may suggest that the learned model identifies A-phase events primarily as intense in power events (A1) rather than an elevation in the signal's spectral content

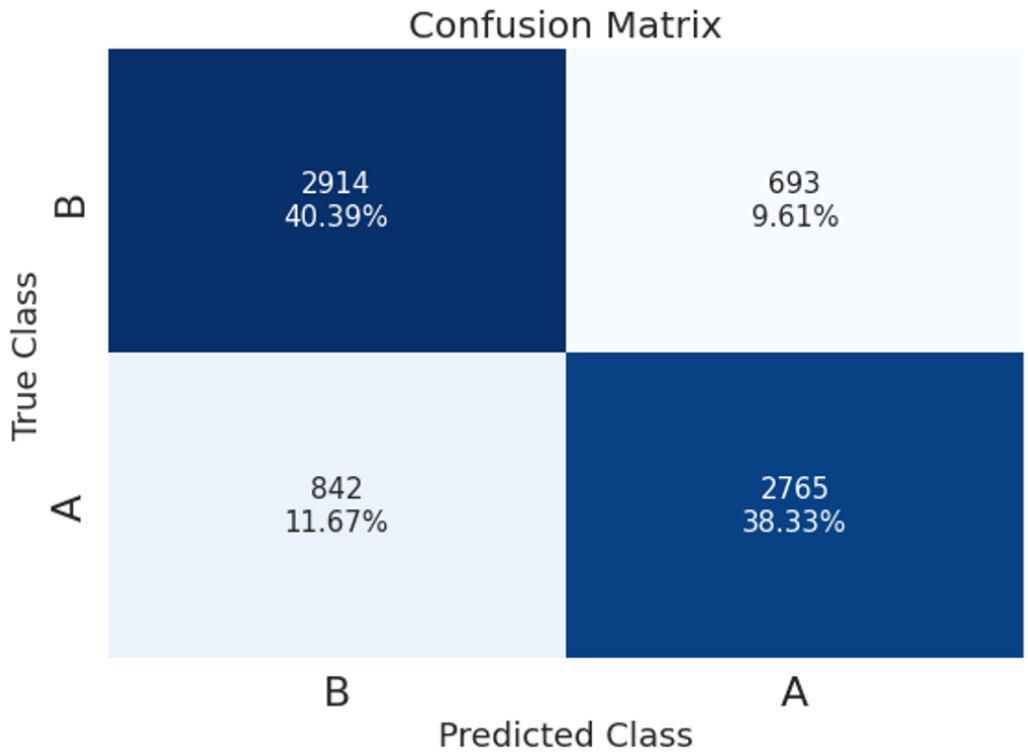

Figure 4.3: Confusion matrix of CAP detection using 9-second segment's length, SPWVD, and TF-augmentations.

(A3). In line with this, A2 events are best identified by the model since they exhibit an increase in both power and frequency.

4.6 Comparative Analysis of Methods

Table 4.4 presents a summary and comparison of recent studies evaluated on a balanced dataset derived from CAPSLPDB. The results of our method, which are indicated in the table were obtained using the smoothed pseudo Wigner-Ville distribution (SPWVD), a 9-second window size, and the proposed TF-augmentations. As shown in the table, our algorithm achieves 77.5% accuracy on a balanced test set and 81.8% on an unbalanced test set, matching and potentially surpassing recent state-of-the-art results. However, it should be noted that this comparison, which is also used in other papers [11, 12, 14], is not fully objective, since while all methods use the CAPSLPDB for evaluation, each constructs its dataset differently, influencing necessarily the reported results. A fair comparison would require all methods to be trained and tested on the exact same dataset, which remains difficult in practice due to the limited availability of other methods' open-source code needed for such standardized testing.

4.7 Summary

This chapter focuses on evaluating the proposed algorithm's performance and conducting an ablation study. The study assesses the impact of different time-frequency representations, the influence of contextual information, and the effectiveness of data augmentation strategies. Key findings are presented through accuracy results, tables, and figures. Notably, the Wigner-based transforms (WVD and SPWVD) outperform the spectrogram in accurately classifying of the CAP phases, and increased window sizes incorporating contextual information of the EEG signals enhance accuracy. Data augmentation, particularly the suggested TF-augmentations, proves effective in classification. The chapter concludes with an in-depth analysis of A-phase detection and a comparative analysis showing the algorithm's superior performance compared to recent studies, achieving 77.5% accuracy on a balanced test set and 81.8% on an unbalanced test set.

Table 4.4: Summary and comparison between recent studies.

Author	**Method**	**Segment Length [Sec]**	**Number of Subjects**	**Performance Parameter [%] on Validation Set**	**Performance Parameter [%] on Test Set**	**Accuracy [%] Evaluated on Unbalanced Test Set**
Dhok et al. (2020)	Wigner-Ville distribution (WVD), Renyi entropy (RE), support vector machine (SVM)	2	6 patients	ACC=72.3 PRE=64.1 REC=76.8 SPE=69.2 F1=69.9	-	-
Sharma et al. (2021)	Wavelet-based features, SVM	2	16 patients	ACC=75.7 PRE=75.0 REC=77.7 F1=76.0	-	-
Sharma et al. (2022)	Biorthogonal wavelet filter bank (BOWFB), ensemble bagged tree -	2 -	6 patients	ACC=74.4 REC=67.53 SPE=81.3		
Hartmann et al. (2019)	Hand–crafted features, long-short term memory (LSTM)	1-3	16 patients	ACC=82.4± 7.1 REC=75.3± 12 SPE= 83.9 ± 8.9 F1=$57.4 \pm$ 9.6	-	-
Loh et al. (2021)	1D-CNN	2	6 patients	ACC=74.4	ACC=73.6 PRE=71.0 REC=80.3 SPE= 67.0 F1=75.3	53.0
Murarka et al. (2022)	1D-CNN	2	6 patients	ACC=76.7	ACC=78.8 PRE=82.5 REC=73.4 SPE=84.3 F1=77.7	60.6
Our method	**Spectrogram, Wigner-based representations, ResNet18**	**1-11**	**16 patients**	**ACC=78.5 PRE=78.9 REC=77.8 SPE= 79.3 F1=78.4**	**ACC=77.5 PRE=78.4 REC=75.9 SPE= 79.1 F1=77.1**	**81.8**

Chapter 5

Conclusions

5.1 Summary

In this research we have focused on the classification of cyclic alternating pattern phases in sleep. The accurate identification of CAP phases is pivotal in assessing sleep quality and diagnosing sleep-related disorders. Presently, CAP classification predominantly relies on traditional machine learning techniques with manual feature extraction, lacking the potential of deep learning and automated feature extraction. This thesis aims to bridge this gap by harnessing the power of deep convolutional neural networks (CNNs) to classify CAP phases.

In section 3 we introduced a fully automatic classifier of cyclic alternating patterns signals, based on a computationally efficient neural network. The proposed method comprises three core components. The initial phase involves pre-processing, where each 1-second EEG segment is contextualized by extending it into a broader time window. Subsequently, these segments undergo transformation into 2D time-frequency matrices in the second stage. This transformation captures both the non-stationary energy and spectral characteristics of the signal, enabling the discrimination between distinct CAP phases. Finally, the resulting 2D time-frequency representations are treated as images. Following normalization and augmentation, they are input into a CNN-based architecture for supervised training. This holistic approach facilitates the neural network's ability to discern and classify CAP phases effectively.

In section 4, the effectiveness of the proposed method was evaluated using a balanced dataset sourced from the Cap Sleep Database (CAPSLPDB), a recognized benchmark for CAP identification and classification research. A thorough experimental study demonstrates that our algorithm surpasses existing machine learning-based methods, achieving an accuracy of 77.5% on a balanced test-set and 81.8% when evaluated on an unbalanced test-set. Furthermore, an ablation study was conducted to examine the individual contributions of the components comprising our approach. Importantly, this exploration has illuminated several pivotal insights - first we found that Wigner-based representations outperform the commonly used short-time Fourier transform for CAP

classification. Additionally, we validate and quantify the value of incorporating contextual information from EEG signals, revealing that 9-second segments yield optimal performance. Lastly, our comparative analysis of diverse data augmentation techniques reveals the most suitable augmentations for effective CAP classification.

5.2 Future Research

Throughout this thesis, we have proposed an end-to-end fully automatic CNN-based method for classifying CAP phases. While obtaining high performance and providing valuable insights into CAP classification details, some questions remain open which can provide a basis for future research. These include:

1. **Advanced deep-learning architecture:** Utilizing contrastive learning can enhance the model's ability to distinguish between similar EEG patterns by learning more discriminative feature representations. Additionally, exploring transformer-based architectures, which are well-suited for sequential data, may improve temporal modeling and offer advantages over traditional convolutional approaches.

2. **Classification noise:** Investigating methods to mitigate random classification noise is essential for improving model reliability, especially in cases where inconsistencies in labeling or inherent signal variability affect performance. Developing robust techniques to handle such noise can enhance the stability and generalizability of the model in real-world applications.

3. **Multi-modal data fusion:** Incorporating multiple data modalities, such as EEG signals along with other physiological signals like heart rate, eye movement, or respiratory patterns, could provide a more comprehensive understanding of sleep dynamics. Exploring methods to effectively fuse these diverse data sources could potentially lead to more robust CAP classification models.

4. **Out-of-distribution generalization** Expanding the dataset to include more patients with diverse clinical conditions can improve the model's ability to generalize across different populations. Additionally, incorporating data from multiple independent datasets can help reduce biases, improve robustness, and ensure the model performs well across a wide range of real-world scenarios.

Bibliography

[1] M. G. Terzano and L. Parrino, "The cyclic alternating pattern (CAP) in human sleep," in *Handbook of Clinical Neurophysiology*, vol. 6, pp. 79–93, Elsevier, 2005.

[2] M. G. Terzano, L. Parrino, A. Sherieri, R. Chervin, S. Chokroverty, C. Guilleminault, M. Hirshkowitz, M. Mahowald, H. Moldofsky, A. Rosa, *et al.*, "Atlas, rules, and recording techniques for the scoring of cyclic alternating pattern (CAP) in human sleep," *Sleep Medicine*, vol. 2, no. 6, pp. 537–554, 2001.

[3] L. Parrino, R. Ferri, O. Bruni, and M. G. Terzano, "Cyclic alternating pattern (CAP): the marker of sleep instability," *Sleep Medicine Reviews*, vol. 16, no. 1, pp. 27–45, 2012.

[4] M. G. Terzano and L. Parrino, "Clinical applications of cyclic alternating pattern," *Physiology & Behavior*, vol. 54, no. 4, pp. 807–813, 1993.

[5] F. Mendonca, A. Fred, S. Shanawaz Mostafa, F. Morgado-Dias, and A. G. Ravelo-García, "Automatic detection of a phases for CAP classification," in *Proc. 7th International Conference on Pattern Recognition Applications and Methods*, pp. 394–400, 2018.

[6] F. Mendonca, A. Fred, S. S. Mostafa, F. Morgado-Dias, and A. G. Ravelo-Garcia, "Automatic detection of cyclic alternating pattern," *Neural Computing and Applications*, vol. 34, no. 13, pp. 11097–11107, 2022.

[7] S. Mariani, A. M. Bianchi, E. Manfredini, V. Rosso, M. O. Mendez, L. Parrino, M. Matteucci, A. Grassi, S. Cerutti, and M. G. Terzano, "Automatic detection of a phases of the cyclic alternating pattern during sleep," in *Proc. Annual International Conference of the IEEE Engineering in Medicine and Biology*, pp. 5085–5088, 2010.

[8] S. Mariani, A. Grassi, M. O. Mendez, L. Parrino, M. G. Terzano, and A. M. Bianchi, "Automatic detection of CAP on central and fronto-central eeg leads via support vector machines," in *Proc. Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 1491–1494, 2011.

[9] S. Mariani, E. Manfredini, V. Rosso, A. Grassi, M. O. Mendez, A. Alba, M. Matteucci, L. Parrino, M. G. Terzano, S. Cerutti, *et al.*, "Efficient automatic classifiers

for the detection of a phases of the cyclic alternating pattern in sleep," *Medical & Biological Engineering & Computing*, vol. 50, pp. 359–372, 2012.

[10] S. Mariani, A. Grassi, M. O. Mendez, G. Milioli, L. Parrino, M. G. Terzano, and A. M. Bianchi, "EEG segmentation for improving automatic CAP detection," *Clinical Neurophysiology*, vol. 124, no. 9, pp. 1815–1823, 2013.

[11] M. Sharma, V. Patel, J. Tiwari, and U. R. Acharya, "Automated characterization of cyclic alternating pattern using wavelet-based features and ensemble learning techniques with eeg signals," *Diagnostics*, vol. 11, no. 8, p. 1380, 2021.

[12] S. Dhok, V. Pimpalkhute, A. Chandurkar, A. A. Bhurane, M. Sharma, and U. R. Acharya, "Automated phase classification in cyclic alternating patterns in sleep stages using wigner–ville distribution based features," *Computers in Biology and Medicine*, vol. 119, p. 103691, 2020.

[13] M. Sharma, A. A. Bhurane, and U. R. Acharya, "An expert system for automated classification of phases in cyclic alternating patterns of sleep using optimal wavelet-based entropy features," *Expert Systems*, p. e12939, 2022.

[14] H. W. Loh, C. P. Ooi, S. G. Dhok, M. Sharma, A. A. Bhurane, and U. R. Acharya, "Automated detection of cyclic alternating pattern and classification of sleep stages using deep neural network," *Applied Intelligence*, vol. 52, no. 3, pp. 2903–2917, 2022.

[15] S. Hartmann and M. Baumert, "Automatic a-phase detection of cyclic alternating patterns in sleep using dynamic temporal information," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 27, no. 9, pp. 1695–1703, 2019.

[16] J. You, Y. Ma, and Y. Wang, "GTransU-CAP: Automatic labeling for cyclic alternating patterns in sleep eeg using gated transformer-based u-net framework," *Computers in Biology and Medicine*, vol. 147, p. 105804, 2022.

[17] K. P. Murphy, *Machine Learning: a Probabilistic Perspective*. MIT press, USA, 2012.

[18] R. Largo, C. Munteanu, and A. Rosa, "Wavelet based CAP detector with ga tuning.," *WSEAS Transactions on Information Science and Applications*, vol. 2, no. 5, pp. 576–580, 2005.

[19] R. Largo, C. Munteanu, and A. Rosa, "CAP event detection by wavelets and ga tuning," in *Proc. IEEE International Workshop on Intelligent Signal Processing, 2005.*, pp. 44–48, 2005.

[20] S. Murarka, A. Wadichar, A. Bhurane, M. Sharma, and U. R. Acharya, "Automated classification of cyclic alternating pattern sleep phases in healthy and sleep-disordered subjects using convolutional neural network," *Computers in Biology and Medicine*, vol. 146, p. 105594, 2022.

[21] T. J. Sejnowski, *The Deep Learning Revolution*. MIT press, USA, 2018.

[22] A. Voulodimos, N. Doulamis, A. Doulamis, E. Protopapadakis, *et al.*, "Deep learning for computer vision: A brief review," *Computational Intelligence and Neuroscience*, vol. 2018, pp. 1–13, 2018.

[23] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Proc. IEEE conference on computer vision and pattern recognition*, pp. 248–255, 2009.

[24] I. Ariav and I. Cohen, "An end-to-end multimodal voice activity detection using wavenet encoder and residual networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 2, pp. 265–274, 2019.

[25] J. Gu, Z. Wang, J. Kuen, L. Ma, A. Shahroudy, B. Shuai, T. Liu, X. Wang, G. Wang, J. Cai, *et al.*, "Recent advances in convolutional neural networks," *Pattern Recognition*, vol. 77, pp. 354–377, 2018.

[26] R. Madan, D. Agrawal, S. Kowshik, H. Maheshwari, S. Agarwal, and D. Chakravarty, "Traffic sign classification using hybrid hog-surf features and convolutional neural networks.," in *Proc. ICPRAM*, pp. 613–620, 2019.

[27] Y. Chen, X. Yang, B. Zhong, S. Pan, D. Chen, and H. Zhang, "Cnntracker: Online discriminative object tracking via deep convolutional neural network," *Applied Soft Computing*, vol. 38, pp. 1088–1098, 2016.

[28] C. Zhang, C. Yao, B. Shi, and X. Bai, "Automatic discrimination of text and non-text natural images," in *Proc. 13th international conference on document analysis and recognition (icdar)*, pp. 886–890, 2015.

[29] N. Kalchbrenner, L. Espeholt, K. Simonyan, A. v. d. Oord, A. Graves, and K. Kavukcuoglu, "Neural machine translation in linear time," *arXiv preprint arXiv:1610.10099*, 2016.

[30] L. Xiang and A. Hu, "Comparison of methods for different time-frequency analysis of vibration signal," *Journal of Software*, vol. 7, no. 1, pp. 68–74, 2012.

[31] S. Scholl, "Fourier, gabor, morlet or wigner: Comparison of time-frequency transforms," *arXiv preprint arXiv:2101.06707*, 2021.

[32] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

[33] L. Taylor and G. Nitschke, "Improving deep learning with generic data augmentation," in *Proc. IEEE symposium series on computational intelligence (SSCI)*, pp. 1542–1547, 2018.

[34] L. Yaeger, R. Lyon, and B. Webb, "Effective training of a neural network character classifier for word recognition," *Advances in Neural Information Processing Systems*, vol. 9, pp. 807–813, 1996.

[35] A. Mikołajczyk and M. Grochowski, "Data augmentation for improving deep learning in image classification problem," in *Proc. international interdisciplinary PhD workshop (IIPhDW)*, pp. 117–122, 2018.

[36] 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.

[37] C. Iber, S. Ancoli-Israel, A. L. Chesson, and S. F. Quan, *The AASM Manual for the Scoring of Sleep and Associated Events: Rules, Terminology and Technical Specifications*, vol. 1. American Academy of Sleep Medicine, Westchester, IL, USA, 2007.

[38] L. Parrino, G. Milioli, A. Melpignano, and I. Trippi, "The cyclic alternating pattern and the brain-body-coupling during sleep," *Epileptologie*, vol. 33, no. 1, pp. 150–160, 2016.

[39] V. Ibáñez, J. Silva, and O. Cauli, "A survey on sleep assessment methods," *PeerJ*, vol. 6, p. e4849, 2018.

[40] E. Niedermeyer and F. L. da Silva, *Electroencephalography: basic principles, clinical applications, and related fields*. Lippincott Williams & Wilkins, 2005.

[41] M. Teplan *et al.*, "Fundamentals of eeg measurement," *Measurement science review*, vol. 2, no. 2, pp. 1–11, 2002.

[42] M. O. Mendez, A. Alba, I. Chouvarda, G. Milioli, A. Grassi, M. G. Terzano, and L. Parrino, "On separability of a-phases during the cyclic alternating pattern," in *Proc. 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 2253–2256, 2014.

[43] F. Machado, F. Sales, C. Santos, A. Dourado, and C. Teixeira, "A knowledge discovery methodology from eeg data for cyclic alternating pattern detection," *Biomedical Engineering Online*, vol. 17, pp. 1–23, 2018.

[44] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proc. of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.

[45] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.

[46] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.

[47] M. Sandsten, *Time-frequency analysis of time-varying signals and non-stationary processes*. PhD thesis, Lund, Sweden: Lund University, Centre for Mathematical Sciences, 2016.

[48] P. Flandrin, *Time-Frequency/Time-Scale Analysis*. Academic press, San Diego, CA, 1998.

[49] P. Flandrin and P. Borgnat, "Time-frequency energy distributions meet compressed sensing," *IEEE Transactions on Signal Processing*, vol. 58, no. 6, pp. 2974–2982, 2010.

[50] P. Flandrin, "Some features of time-frequency representations of multicomponent signals," in *Proc. ICASSP'84. IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 9, pp. 266–269, 1984.

[51] Y. Zhou, F. Ren, S. Nishide, and X. Kang, "Facial sentiment classification based on resnet-18 model," in *Proc. International Conference on electronic engineering and informatics (EEI)*, pp. 463–466, 2019.

[52] E. Jing, H. Zhang, Z. Li, Y. Liu, Z. Ji, and I. Ganchev, "Ecg heartbeat classification based on an improved resnet-18 model," *Computational and Mathematical Methods in Medicine*, vol. 2021, pp. 1–13, 2021.

[53] Y. LeCun, L. Bottou, G. B. Orr, and K.-R. Müller, "Efficient backprop," in *Neural Networks: Tricks of the Trade*, pp. 9–50, Springer, Germany, 2002.

[54] N. Bjorck, C. P. Gomes, B. Selman, and K. Q. Weinberger, "Understanding batch normalization," *Advances in Neural Information Processing Systems*, vol. 31, pp. 7694—7705, 2018.

[55] S. Kahl, T. Wilhelm-Stein, H. Hussein, H. Klinck, D. Kowerko, M. Ritter, and M. Eibl, "Large-scale bird sound classification using convolutional neural networks.," *CLEF (working notes)*, vol. 1866, 2017.

[56] E. Lashgari, D. Liang, and U. Maoz, "Data augmentation for deep-learning-based electroencephalography," *Journal of Neuroscience Methods*, vol. 346, p. 108885, 2020.

[57] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "Specaugment: A simple data augmentation method for automatic speech recognition," *arXiv preprint arXiv:1904.08779*, 2019.

בספרות, ולרוב בא לידי ביטוי באימון של רשת נוירונים באותות החד-מימדיים, זאת מבלי לכלול עיבוד מקדים של האותות. בתזה זו אנו שואפים לגשר על פערים אלו, ומציעים מסווג דפוסי שינה מחזוריים המבוסס על למידה עמוקה של התמרות זמן-תדר של אותות אא"ג. הפתרון המוצע מורכב ממספר שלבים, המותאמים לאופי בעיית הסיווג של דפוסי השינה המחזוריים. בהתחשב בכך שניתוח אותות אא"ג על פני פרקי זמן ממושכים עשוי לשפר את ההבנה של מבנה האות, נבחן שימוש בחלונות זמן בגדלים משתנים. כמו כן, מעבר לשימוש בהתמרת-פוריה-לזמן-קצר (ספקטרוגרמה) כאנליזת זמן-תדר סטנדרטית מקובלת, אנו בוחנים את השימוש בהתמרות זמן-תדר מבוססות התמרת ויגנר, ואף מראים כי ניצול של התמרות אלו משיג דיוק גבוה יותר בסיווג האותות. יתר על כן, בכדי להכליל את המודל הנלמד ולמנוע מצב של התאמת-יתר נעשה שימוש באוגמנטציות המתאימות למבנה אות אא"ג אופייני.

אנו בוחנים את השיטה המוצעת תוך שימוש במאגר נתונים ציבורי הנחשב למסד נתונים השוואתי מקובל להערכת ובחינת עבודות בתחום. המאגר כולל מספר רב של הקלטות שינה מלאות אשר תויגו ידנית על ידי מומחים. הערכת הביצועים נעשתה הן על פני סט מדידות מאוזן (מבחינת שכיחות הופעות של פאזה א' ופאזה ב') והן על פני סט מדידות שאינו מאוזן (הקלטת לילה מלאה). בפרט, אנו עורכים סדרה מקיפה של ניסויים לשם הערכת התרומה של כל חלק בשיטה המוצעת, כדוגמת משך חלון הזמן, סוג ההתמרה וסוגי האוגמנטציות של המדידות. ניתוח תוצאות הניסויים מלמד כי השיטה המוצעת משיגה תוצאות דיוק גבוהות ביחס לשיטות מבוססות למידת מכונה קלאסית קיימות, זאת גם תוך שימוש ברשת למידה עמוקה יעילה חישובית הניתנת למימוש על פני מכשור רפואי.

תקציר

תזה זו עוסקת בסיווג מבוסס למידה עמוקה של דפוסי שינה מחזוריים מתחלפים. מבחינה פיזיולוגית, השינה משתנה באופייה ואינה תהליך אחיד והומוגני והיא מתאפיינת לאורכה בשינויים בתגובות ובמדדים הפיזיולוגיים של הגוף. באופן כללי, מקובל לסווג את השינה לשני מצבים עיקריים - מצב שינה עם תנועת-עיניים-מהירה ומצב שינה ללא-תנועת-עיניים-מהירה, כאשר במהלך השינה הגוף נד לסירוגין בין שני מצבי שינה אלו. בתזה זו אנו מתמקדים בדפוסי שינה מחזוריים מתחלפים- מושג רב משמעות בחקר השינה המתייחס לשינויים בפעילות המוחית המתרחשים במהלך מצב השינה ללא-תנועת-עיניים-מהירה. נהוג לחלק את דפוסי השינה המחזוריים לשתי קבוצות עיקריות, כאשר ההבחנה היא בין מצב של פעילות מוחית מוגברת (פאזה א') לבין מצב של פעילות מוחית מוחלשת (פאזה ב'). מעבר להיותם תופעה פיזיולוגית, דפוסי שינה מחזוריים נחשבים למדד מהימן של אי-יציבות בשינה שכן קיימת התאמה בין הופעות לא סדירות של דפוסי שינה מחזוריים להופעה של מספר פתולוגיות הקשורות לשינה, כמו למשל, דום נשימה בשינה, חריקות שיניים נדודי שינה ועוד. כתוצאה מכך, לזיהוי של שלבי דפוסי שינה מחזוריים יש תפקיד מכריע באבחון והערכת איכות השינה. את הפעילות המוחית בזמן השינה מודדים מתוך ניתוח של אותות אלקטרואנצפלוגרם (אא"ג) המוקלטים בזמן השינה באמצעות אלקטרודות המוצמדות לקרקפת הנבדק. המדידות מוקלטות במהלך שנת לילה מלא של מטופל, מה שמוביל למדידות ארוכות בעלות אורך אופייני של 8-10 שעות כל אחת. כיום, ניתוח מדידות אא"ג וסיווג דפוסי השינה המחזורים המתחלפים לפאזה א' ולפאזה ב' נעשה בעיקר באופן ידני על ידי רופאים ובעלי מקצוע מוסמכים, מה שמוביל לתהליך המחייב השקעת זמן אנושי רב וכרוך בחסרונות נוספים כדוגמת סובייקטיביות בפענוח המדידות וטעויות אנוש הנובעות מעייפות וחוסר תשומת לב.

בכדי לייעל תהליך יקר זה הוצעו לאורך השנים שיטות שונות לפיענוח וסיווג אוטומטי של דפוסי השינה המחזוריים המתחלפים. רוב השיטות מתבססות על חילוץ של מאפיינים מתוך אותות האא"ג, ואימון בצורה מודרכת של מסווגים מבוססי שיטות למידת-מכונה קלאסיות. המאפיינים המחולצים נבחרים לרוב בצורה ידנית, על פי תכונות אופייניות למצבי דפוסי השינה המחזוריים הנבדלים זה מזה בעוצמה ובתדר. לחילופין, היות ואותות אא"ג הינם אי-סטציונרים, בחלק מהעבודות נעשה שימוש בהתמרות זמן-תדר של האותות לשם אפיון השינויים התדריים המתרחשים עם הזמן. אמנם, גם במרבית עבודות אלו ההתמרה משמשת, בסופו של דבר, כאמצעי לחילוץ ידני של מאפיינים ממישור הזמן-תדר של האותות.

בשנים האחרונות, אנו עדים לשימוש הולך וגובר בשיטות מבוססות למידה עמוקה לפתרון מגוון רב של בעיות, זאת לנוכח ההישגים הפנומנליים המתקבלים באמצעות השימוש בשיטות אלו. אמנם, יישום של טכניקות למידה עמוקה עבור בעיית סיווג דפוסי שינה מחזוריים מתחלפים מצומצם מאוד

המחקר בוצע בהנחייתם של פרופסור ישראל כהן וד"ר אלון עמר בפקולטה להנדסת חשמל ומחשבים.

התוצאות של חיבור זה פורסמו כמאמר מאת המחבר ושותפיו למחקר בכתב-עת במהלך תקופת מחקר המגיסטר של המחבר.

מחבר חיבור זה מצהיר כי המחקר, כולל איסוף הנתונים, עיבודם והצגתם, התייחסות והשוואה למחקרים קודמים וכו', נעשה כולו בצורה ישרה, כמצופה ממחקר מדעי המבוצע לפי אמות המידה האתיות של העולם האקדמי. כמו כן, הדיווח על המחקר ותוצאותיו בחיבור זה נעשה בצורה ישרה ומלאה, לפי אותן אמות מידה.

תודות

ברצוני להודות מקרב לב למנחים שלי, פרופ' ישראל כהן וד"ר אלון עמר, על שליוו אותי לאורך התואר והעניקו לי רבות מחוכמתם ומניסיונם.

תודה מעומק הלב להוריי, שגידלוני וחינכוני.

אחרונה חביבה- תודה לאשתי היקרה רונית, שצועדת איתי יד ביד לאורך כל הדרך.

סיווג מבוסס למידה עמוקה של דפוסי שינה מחזוריים מתחלפים

חיבור על מחקר

לשם מילוי חלקי של הדרישות לקבלת התואר
מגיסטר למדעים בהנדסת חשמל

יואב כהנא

הוגש לסנט הטכניון — מכון טכנולוגי לישראל
שבט התשפ"ה חיפה פברואר 2025

סיווג מבוסס למידה עמוקה של דפוסי שינה מחזוריים מתחלפים

יואב כהנא