Representing and Classifying EEG Signals in Mental Task Brain-Computer Interfaces

Research Exam
June 14, 2013

Elliott M. Forney
Colorado State University
Computer Science Department
idfah@cs.colostate.edu

ABSTRACT

Brain-Computer Interfaces (BCI) provide a mechanism for establishing a direct channel of communication between a user's brain and a computerized device. BCI are rapidly gaining acceptance as a form of assistive technology and may be used by people with motor impairments to communicate and use tools such as phones, computers, wheelchairs and prosthetics. Eventually, BCI might also become a commonplace form of human-computer interaction.

Although various types of BCI are currently being explored, approaches that combine electroencephalography (EEG) with imagined mental tasks (MT) appear to be promising because they are non-invasive, asynchronous and stimulus-free. However, such BCI have not yet reached a level of performance that is acceptable for use in practical applications. This may be partially due to limitations found in current methods for representing and classifying spontaneous EEG signals.

In this research exam we investigate various methods for representing and classifying EEG signals for this specific type of BCI. First, we explore the use of frequency-domain representations, including power spectral densities, wavelet transforms and phase locking value. These approaches do well at capturing oscillatory patterns in EEG signals and align well with the EEG analysis techniques used in other disciplines. However, they typically have a limited ability to capture some forms of spatial or temporal patterns.

We then explore the use of time-domain signal representations. Approaches that rely on time-delay embedding are able to capture both spatial and temporal patterns but at the expense of high dimensionality. Time-series modeling approaches based on Autoregressive Models and Artificial Neural Networks can be used to construct generative classifiers with fewer model parameters to optimize. It remains unclear, however, whether or not time-series models are able to capture patterns across various time scales. We then discuss some empirical results gathered from current literature and summarize our conclusions about each of these methods. Finally, we offer some potential directions for further research.

Contents

1 Introduction

THE SUCCESS of the human animal can be largely attributed to our capacity for language and communication and to our ability to construct and manipulate tools. Communication allows us to evade predators through collaboration and to avoid danger through warnings and advice. Exchanging information allows us to teach and learn complex ideas, express ourselves and share the human experience. Tools allow us to gather food through hunting and agriculture, protect ourselves using weapons, simplify repetitive tasks and attain things that would otherwise be inaccessible. When combined, tools and communication can also yield new forms of interaction, such as works of art, the written word, the radio and the internet. There is no doubt that these skills are extremely important to humankind and are responsible for many of our unique advantages.

Unfortunately, there are a number diseases that can limit an individual's ability to communicate or utilize tools. Such impairments are relatively commonplace and vary widely in form and severity. Spinal cord injury, traumatic brain injury, stroke and neurodegenerative diseases can all cause a wide variety of impairments to motor, speech and sensory function. Since all known biological means of communication rely on either locomotion or the release of hormones, with conscious communication relying primarily on the former, diseases that impair motor function often limit one's ability to communicate and use tools. In severe cases, a person may have very little or no ability to interact with the outside world, a condition known as Locked-In Syndrome (LIS), despite retaining cognitive function and an awareness of their surroundings [1, 2].

A number of assistive technologies (AT) have been developed to help restore function to those afflicted with motor impairments. Current AT range from wheelchairs and reaching tools to switches, eye-trackers and natural language processing systems. However, the effectiveness of these devices varies widely among users and their usefulness depends on the type and severity of impairment. Since current AT require some level of physical interaction, there are few, if any, options available to those with severe motor impairments or LIS.

Brain-Computer Interfaces (BCI) are a type of experimental AT that aim to restore an individual's ability to communicate and manipulate tools by establishing a direct channel of communication between the human brain and a computerized device. Since BCI bypass innate motor-based means of communication, they may provide a viable option for people who are unable to use other forms of AT. Eventually, BCI might also find their way into consumer devices and become an everyday form of human-computer interaction, replacing current electromechanical input devices in a number of circumstances.

Despite the exciting prospect of a new form of communication, BCI research is still in its infancy and current systems are not yet reliable enough or user-friendly enough to be adopted for wide-spread commercial use. The field has, however, seen increasing success and an explosion of research in recent years [3]. In fact, it appears that a BCI using functional magnetic resonance imaging (fMRI) has already been used to reestablish communication with at least some people that were previously diagnosed as being in a vegetative state [4]; although such BCI are not practical for everyday use.

In order to construct BCI that are more appropriate for use in commodity AT, a number of challenges must first be overcome. These challenges can be roughly divided among the three components that constitute a BCI. First, a sensitive method for monitoring brain activity is necessary to receive information from the user in the absence of physical movement. Second, an effective communication protocol must be established that allows the user to convey an intended instruction to the BCI. Third, signal processing and machine learning algorithms are required to represent and classify patterns in the user's brain activity in order to identify which messages from the communication protocol the user intends to send. Each of these components must be carefully designed to be reliable and robust while also striking an appropriate balance between accuracy and usability.

BCI that combine electroencephalography (EEG) for monitoring brain activity and imagined mental tasks (MT) for establishing a communication protocol appear to be particularly promising because they are non-invasive, asynchronous, do not require external stimuli and elicit activity across a wide variety of brain regions. However, representing and classifying EEG signals within this framework has proven to be extremely challenging because of noise, artifacts and the complexity of the human brain.

In this research exam, we explore state-of-the-art methods for representing and classifying spontaneous EEG signals within the context of this specific type of BCI. In the remainder of Section 1, we first offer some background information on BCI that use EEG and MT and explain their advantages over other approaches. Next, we describe the difficulties encountered in the signal processing and machine learning components of these BCI. We assert that finding signal representations and classification algorithms that perform well is difficult, especially since EEG signals tend to be noisy and non-stationary.

In Section 2, we continue by exploring several popular methods for representing EEG signals in the frequency-domain, including power spectral densities, wavelet transforms and phase locking values. In Section 3, we continue by investigating several promising time-domain signal representations. These approaches include time-delay embedding as well as linear and non-linear time-series models. For each of these signal representations, we also examine the classification algorithms that have been used and summarize the results obtained by various research groups. In Section 4, we conclude with some analysis of these approaches and by offering some potential directions for further research.

1.1 Non-Invasive BCI using Electroencephalography

Every BCI must have a method for monitoring changes in brain activity so that it can receive information from the user without physical interaction. The technologies that have been proposed for this role typically measure one of two types of physiological changes: metabolic or electromagnetic. Methods that monitor metabolic changes in the brain include functional magnetic resonance imaging (fMRI) and functional near-infrared spectroscopy (fNIRS) [4, 5]. These methods rely on the fact that active regions of the brain consume more oxygen than inactive regions. Unfortunately, these changes take place on the scale of decaseconds, which is likely too slow for real-time BCI. fMRI systems are also not practical for general-purpose BCI since they are expensive and too large to be portable. Although fNIRS systems are inexpensive and portable, they currently have poor spatial or temporal resolution and are only able to observe brain activity that takes place on the very surface of the cortex.

Methods that rely on electromagnetic changes include local field potentials (LFP), electrocorticography (ECoG) and electroencephalography (EEG) [6, 7, 8]. These methods rely on the fact that the depolarization of a neuron when it fires an action potential produces a small electromagnetic field. Among these methods, LFP and ECoG are invasive and require surgical implantation. Although invasive methods typically achieve high spatial and temporal resolution, they are often not an appealing option for BCI users because of the risks involved. Although some BCI users with severe impairments may be willing to accept these risks, medical ethics dictate that non-invasive approaches should be preferred when they are available and perform well.

EEG is another electrophysiological method for monitoring brain activity that is non-invasive and appears to have tremendous potential for use in BCI. EEG measures brain activity using an array of electrodes placed on the surface of a user's scalp. Electrical fields produced by the synchronized firing of action potentials in neurons near the cortical surface of the brain induce observable currents in the EEG electrodes [9]. Modern EEG equipment is small, portable and relatively inexpensive. EEG also has a high temporal resolution, on the scale of microseconds, allowing it to make observations at a rate that is acceptable for real-time BCI. There is also a long history of research surrounding EEG.

EEG does, however, involve a number of challenges that must be addressed in order for it to be used in an effective and practical BCI. First, EEG signals tend to have a very low signal strength because of the relatively small electrical potentials produced by action potentials and because the resulting electromagnetic fields must travel through a number of layers of tissue, including the meninges, skull and scalp. Ultimately measured on the scale of microvolts, EEG also suffers from artifacts and a low signal-to-noise ratio due to other sources of electrical interference at this scale. Sources of interference that originate from within the body include ocular movement, muscle contractions and cardiac cycles. Sources of external interference include computer peripherals, radio transmissions and alternating current power mains. EEG is also mostly limited to measuring electrical potentials generated by the synchronized firing of neurons in the cortex that have specific orientations and are near electrode sites [9]. This means that the firing of individual action potentials and activity that is located deep inside the brain cannot be measured.

Although EEG has some drawbacks, it appears to be one of the most viable options for constructing practical, non-invasive BCI. Some problems with interference and noise can also be mitigated by recent advances in EEG recording technology. For example, highly conductive electrodes with cable shielding, embedded preamplifiers and sophisticated common-mode rejection circuits can reduce interference from sources outside the brain [10, 11]. Filters can also be used to remove noise and artifacts from EEG signals. Since EEG appears to be extremely promising for use in non-invasive BCI, the remainder of this document will assume the use of EEG for monitoring brain activity unless otherwise noted.

1.2 Mental Task Communication Paradigm

In addition to a method for monitoring brain activity, BCI must have a communication protocol for exchanging information between the user and the device to be controlled. It is important to realize that BCI are not capable of directly extracting a user's thoughts and, in fact, are not designed to do so. Instead, BCI seek to identify a user's intent through voluntary changes in mental state. As such, a protocol must be established that associates specific mental states with potential instructions to be delivered to the BCI. An effective communication protocol should produce distinct changes in brain activity while also being easy to perform and minimally distracting the user.

A number of paradigms for establishing BCI communication protocols have been proposed. Notably, Farwell and Donchin proposed an approach known as The P300 Speller [12]. Using this paradigm, a user communicates with a BCI by focusing on a character within a grid shown on a computer screen, akin to a virtual keyboard. The rows and columns of the grid then flash while the user makes a mental note of when the character that they are attending to flashes. The BCI then attempts to identify the intended character by searching for a change in brain activity, known as a P300, that occurs shortly after a rare-but-expected stimuli [13].

Although it appears that P300 spellers can be used as an effective form of AT for some users with severe motor impairments [14, 15, 16], it suffers from a number of drawbacks. Since P300 spellers require the user to attend to external stimuli, a degree of control over eye-gaze is required [17]. Since some people with severe motor impairments have little or no control over eye-gaze, this may be prohibitive. Furthermore, attending to external stimuli prevents the user from focusing on their work and their surroundings, making it difficult to smoothly control devices such as electric wheelchairs or mouse cursors. Finally, P300 spellers and similar paradigms operate synchronously. In other words, the user must issue instructions in lockstep with the stimuli.

One paradigm for establishing a communication protocol that may avoid many of these limitations relies on spontaneously performed imagined motor movements, or motor imagery (MI). Extensively explored by Pfurtscheller and Neuper, et al., MI was inspired by the observation that activity in frequencies roughly between 9–14Hz and 18–26Hz, known as μ rhythms, tend to be suppressed in EEG signals over the motor cortex in the hemisphere that is contralateral to the side on which the motor task is being performed [18, 19, 20]. For example, if a user imagines moving their left-hand, a decrease in μ rhythms would be expected to be seen over the centro-parietal region of the right hemisphere of the brain. A BCI could then instruct a computer cursor to move to the left. Similarly, imagined right-hand movement could instruct the BCI to move a cursor to the right. Unlike P300 spellers, MI allows for asynchronous communication and spontaneous control.

MI does, however, suffer from a number of limitations. For instance, BCI that rely on MI typically have few degrees of freedom. In other words, it may be difficult when using MI to design a BCI that allows the user to issue more that one or two instructions. Although there are many potential motor tasks that may be used, the associated changes in μ rhythms are likely to originate from only two cortical regions, the left and right hemispheres of the motor cortex. Furthermore, there have been reports that not all people with motor impairments find it easy to perform imagined motor movements [21]. For those that have impaired motor function, possibly for a prolonged length of time or from birth, it may be difficult to imagine performing tasks that they cannot physically perform.

Including more general mental tasks may yield more flexible BCI communication protocols. In the mental task (MT) paradigm, a user delivers instructions to the BCI by performing one of several imagined mental activities. For example, a user might silently sing a song to move a computer cursor to the left or visualize a geometric figure to move it to the right. A countless number of mental tasks could potentially be used. MT was originally proposed by Keirn and Aunon and some early motivating works are briefly discussed in Section 2 [8].

Although MT appears to be used relatively rarely in current BCI literature, it has a number of advantages over other approaches. Since MT can be chosen so that they elicit brain activity in a variety of brain regions, it may be possible to achieve more degrees of freedom than with MI exclusively. This may also aid the machine learning component of a BCI by making the corresponding changes in brain activity more discriminable. Furthermore, MT may allow each user to choose mental tasks that work well for them and that they are comfortable performing repeatedly. This may help to provide a good user experience while maximizing attention and minimizing distraction. Users may also adapt the way that they perform each, potentially improving performance as their experience level increases. Finally, MT allows for asynchronous and stimulus-free control, potentially allowing fluid, second-nature control. For these reasons, the remainder of this document will focus on MT unless otherwise noted.

1.3 Capturing Patterns in Electroencephalographic Signals

In order to determine the instruction that a user intends to deliver, a BCI must have a method for representing EEG signals and classifying these representations according to the mental task that the user is performing. This is an important part of BCI because the effectiveness of this stage dictates the flexibility and speed allowed in the communication protocol as well as the level of noise and artifacts that can be tolerated. If EEG signals cannot be classified accurately, then a BCI will not reach satisfactory level of performance.

Representing and classifying EEG signals associated with spontaneous mental tasks is not trivial for a number of reasons. First, EEG signals are known to have a low signal-to-noise ratio, as discussed in Section 1.1. Second, the human brain is continually performing multiple conscious and unconscious tasks simultaneously. Patterns associated with relevant cognitive processing must be identified while irrelevant information should be ignored. EEG also produces a large amount of data, typically between 256–1024 observations per second for each of 8–64 channels. Since MT calibration procedures yield relatively little data, typically 5–10 segments each lasting 10-seconds per mental task, many EEG signal representations have high dimensionality relative to the number of training examples.

Given the fact that the human nervous system has a recurrent architecture consisting of billions of neurons that interact using sophisticated electrochemical mechanisms, it is not surprising that patterns found in EEG signals are often very complex. EEG signals contain patterns that are both temporal, i.e., across time, and spatial, i.e., across electrodes. EEG signals can also have components that are non-stationary, meaning that the characteristics of the signal change over the course of time; although there are studies demonstrating that MT can be reliably discriminated after days, weeks or even months [22, 23]. EEG signals can also change along with a user's mood, level of fatigue and attention.

In order to deal with these challenges effectively, EEG signal representations and classification algorithms should satisfy a number of conditions. First, representations should capture both spatial and temporal patterns. If this is not achieved, it is likely that important relationships, either across brain regions or through time, will be omitted. Second, the representation and classifier should be robust to noise, artifacts and background information. This may be achieved through the signal representation, by allowing unimportant aspects of the signal to be removed, or through the classification algorithm, via regularization. Third, classifiers must be capable of drawing class boundaries that adequately discriminate signal representations produced during different mental tasks. Finally, it should be possible to generate these representations and perform classification in real-time on portable hardware. The remainder of this research exam surveys a number of state-of-the-art methods for representing and classifying spontaneous EEG signals in BCI that use the MT communication paradigm.

2 Frequency-Domain Representations

Frequency-domain representations describe EEG signals in terms of oscillatory components and how those components change over time. Since each EEG electrode primarily measures the synchronized firing of local populations of neurons and because of the recurrent architecture of the central nervous system, it seems intuitive to look for patterns in the frequency domain. In fact, it is common in most neuroscience disciplines for EEG signals to be described in terms of activity in standard frequency bands: δ $(0-4\text{Hz})$, θ $(4-8\text{Hz})$, α $(8-13\text{Hz})$, β $(13-30)$ and γ $(\geq 30\text{Hz})$ [24]. Taking this a step further, the distribution of EEG signal power, energy or phase synchronization across EEG sensors can be used to make inferences about brain structure or activity and, of course, the mental state of a BCI user.

In a seminal work exploring the use of MT, Keirn and Aunon note that it may be possible to differentiate between various mental tasks by looking at the symmetry or asymmetry of EEG power in certain frequency bands across the two hemispheres of the brain [8]. This intuition is supported by a number of earlier EEG studies that demonstrate marked differences between activity over the two brain hemispheres during various mental tasks [25, 26, 27, 28]. In more recent works, changes in EEG power during various mental tasks are observed over a wide variety of brain regions and frequency bands [29, 30, 31, 32]. These changes in EEG signal power are often referred to as event-related de/synchronization (ERD/S).

Some of the most commonly observed forms of ERD/S are seen in the α and β bands of the EEG spectrum and are related to idling or the suppression of neural activity [31]. For example, a clear increase in α-band power can be seen over the occipital lobe, i.e., the visual cortex, when a subject closes his or her eyes. It is believed that neurons in the visual cortex tend to fire action potentials in a synchronized fashion resonating at roughly $8-12\text{Hz}$ in the absence of visual stimuli. Similarly, when an individual performs a motor task, a decrease in power in the α or β over the centro-parietal region of the cortex that is contralateral to the side of the movement can often be identified [19]. These changes are consistent with the notion that the motor cortex generates rhythmic activity when it is idle and that this activity desynchronizes when the appropriate region of the motor cortex becomes active. Known to as μ rhythms, changes in ERD/S over the motor cortex can be observed during both actual motor movements as well as imagined motor movements and have been the foundation for many BCI that rely on MI communication paradigms.

In light of this evidence, it appears that frequency-domain EEG representations may be well-suited for capturing patterns in EEG signals. In the remainder of Section 2 we examine several of the most common and effective frequency-domain signal representations used in MT-based BCI and the classification algorithms that work well with these representations. Through this analysis, we conclude that frequency-domain representations are good at capturing oscillatory components of EEG signals. There are, however, several types of spatial and temporal patterns that these representations are not able to capture.

2.1 Power Spectral Densities

Many of the most common methods for representing EEG in the frequency domain estimate power, typically measured in micro-volts squared per hertz ($\mu V^2 Hz^{-1}$), across the signal's spectrum. Known as Power Spectral Densities (PSD), these representations typically use a Discrete Fourier Transform (DFT) to represent a segment of an EEG signal as a linear combination of sinusoids in the complex plane. From there, it is relatively straightforward to determine the amplitude, and therefore power, of the signal at each frequency component of the DFT. Since, however, these segments are necessarily finite in length and since digital EEG is a discrete sample of the underlying signal, the resulting PSD is only an estimate of the true power spectrum.

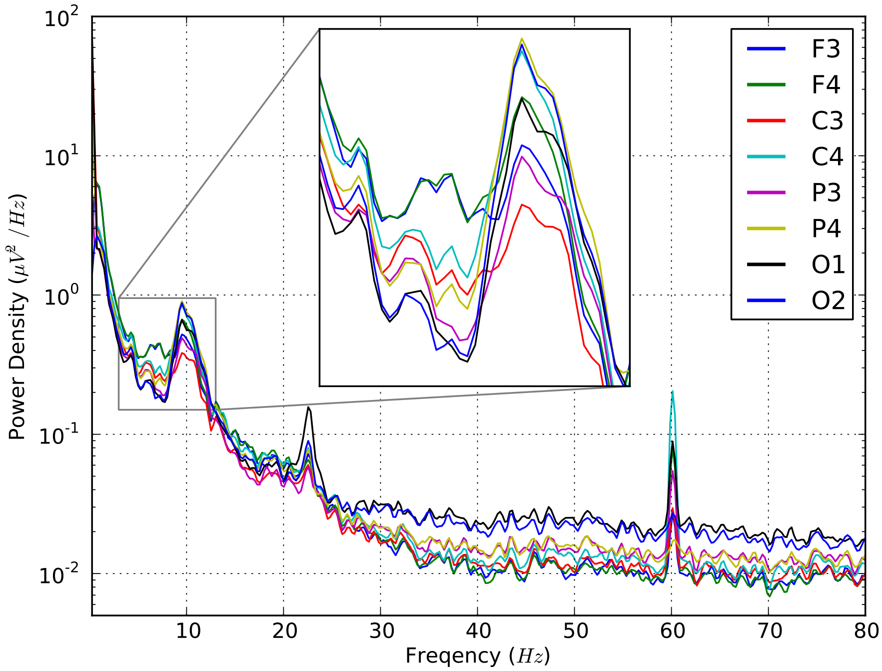

Figure 1: Example Welch PSD for resting-state EEG. The inset figure highlights differences in power in the α band across all 8 channels.

Known as raw periodograms, PSD that are estimated directly using the DFT of a short EEG segment tend to be extremely noisy. As such, a technique known as Welch's Method is commonly used to smooth PSD estimates. Welch's Method works by subdividing the EEG segment into smaller overlapping segments. These segments are then multiplied by a window function, typically a Hanning Window, in order to ensure equal representation of the overlapping portions of the segments. The raw periodogram is then computed for each segment and the result is averaged. This results in a significantly smoother PSD estimate, although resolution in frequency is also lost. Heinzel, et al., provide a detailed explanation of PSD estimates [33]. Figure 1 shows a sample Welch PSD estimated using three minutes of eight-channel resting-state EEG. The inset figure highlights differences in power in the α band across channels. Also, note the spike in power at 60Hz caused by alternating current interference from power mains.

A number of BCI studies have used MT in combination with features derived from PSD. In one of the earliest studies using MT, Keirn and Aunon suggested the use of what they refer to as power asymmetry ratios to capture differences in brain activity across the two hemispheres of the brain. These asymmetry ratios have the form

$$\frac{R_{f,g} - L_{f,g}}{R_{f,g} + L_{f,g}} \tag{1}$$

where $R_{f,g}$ and $L_{f,g}$ are the average power over a given frequency band f in the channels over the left and right hemispheres of the brain for a given brain region g.

Although this approach has the advantage of reducing the dimensionality of the PSD by combining the left and right channels and by averaging the power over the frequency range defined by f, the number of features is still likely to exceed the number of training examples. Keirn and Aunon suggest further reducing the dimensionality of this feature representation by applying forward selection; i.e., repeatedly inserting the remaining feature that improves validation performance until little or no improvement is seen. They then suggest the use of Quadratic Discriminant Analysis (QDA) in order to classify the resulting features. QDA is a Bayesian classifier that finds a quadratic class boundary by fitting sample data from each class with a Gaussian using the sample means and covariance matrices [34].

Although Keirn and Aunon's approach delivers promising results, achieving classification accuracies between 76–84% correct for various two-task problems at two-second intervals, it also suffers from a number of clear drawbacks. First, their use of paired left and right channels eliminates some forms of patterns that may occur. For example, if the power in a given frequency band increases in both the left and right channels simultaneously, as has been observed during some mental tasks, then the denominator in equation (1) will scale down the resulting feature so that the change is not represented [27]. This scheme also places an emphasis on changes in power across the left and right hemispheres while neglecting changes in power that occur in various regions of the same hemisphere of the brain. Furthermore, it remains unclear if the use of forward selection to further reduce the number of features is effective and whether or not these power asymmetry ratios are normally distributed, a prerequisite for QDA.

Following the work of Keirn and Aunon, a number of research groups have shown encouraging results with PSD features. Notably, Millán, et al., have progressively refined a BCI that uses spontaneous mental tasks and PSD features [22, 35, 36, 37]. Early in their work Millán, et al., outline a technique for classifying EEG using PSD that they describe as a Local Neural Classifier [22]. In this approach, a committee of networks is used to assign class labels using the output of Welch's Method. These networks consist of two layers, the first of which contains a Gaussian transfer function while the second layer contains a linear transfer function. In order to train the networks, a Self Organizing Map (SOM) is first used to cluster the PSD features in an unsupervised fashion. The means and covariance matrices of these clusters are then used to initialize the Gaussian units in the first layer of the network. Gradient descent is then used to further update the network weights until peak validation performance is achieved. In the more recent works by Millán, et al., this approach has been somewhat simplified into what is essentially a Mixture-of-Gaussians Classifier that is initialized using SOM and updated using gradient descent [36, 37]. The intuition behind this approach is based on the observation that the ERD/S associated with a given mental task may manifest itself in several different ways as the user repeats the task over time. Millán, et al., achieve good results with offline classification accuracies greater than 70% for three mental tasks at $1/2$ second intervals. They were also able to achieve limited control of mobile robots and electric wheelchairs in online settings [35, 37, 36].

Although the above two groups have shown that PSD features combined with Bayesian classifiers can be used to construct BCI, the performance of these systems does not appear to be high enough for use in a practical AT. One possible reason for this may be that the classifiers used rely on sample estimates of the covariance matrix for the PSD features. Since the number of PSD features used typically exceeds the number of training observations by an order of magnitude, these covariance estimates are likely poor. One potential solution to this problem may be to regularize the covariance estimates using shrinkage [38]. Furthermore, it is unclear whether or not these classifiers are able to adequately fit the class boundaries, suggesting that other non-linear classifiers should be explored.

PSD representations also suffer from several inherent problems that may limit their ability to capture some forms of patterns in EEG. Since PSD do not typically include phase information, they have a limited ability to express spatial patterns. For example, consider a two-channel signal where both channels oscillate at 10Hz. A PSD of this signal is unable to distinguish the case where the signals in both channels have identical phase from the case where both signals have a constant phase difference. PSD also have a limited ability to represent the order of frequency changes within the window over which the PSD was estimated. For example, the PSD for a signal that moves from a low frequency to a high frequency is indistinguishable from a signal that moves from a high frequency to a low frequency. This is a result of the fact that DFT represent a signal in terms of sinusoids. Since sinusoids oscillate indefinitely, they are not well suited for representing non-stationary signals.

In order to characterize non-stationary patterns in EEG signals and assign successive class labels, PSD are typically generated using a number of sequential, potentially overlapping, windows. However, this leads to a trade-off between frequency resolution, temporal resolution, dimensionality and the quality of the power estimates. For example if the width between successive PSD estimates is large, then a higher frequency resolution is obtained at the expense of temporal resolution. If the width between successive PSD is small, then higher temporal resolution is achieved but the resulting power estimate may be poor.

It is also important to note that Welch's method leads to similar trade-offs. For example, if a small window width is used in Welch's method, then a very smooth PSD will be produced but frequency resolution will be lost. On the other hand, if a wide window width is used, the resulting PSD will have high frequency resolution but may yield a poor estimate of power. Unfortunately, none of the studies reviewed in this section appear to have examined these trade-offs.

2.2 Continuous Wavelet Transforms

Continuous Wavelet Transforms (CWT) are another technique for generating frequency-domain representations of a signal that have been successfully used in BCI. CWT represent a time varying signal in terms of a function known as a wavelet or mother wavelet. Simply stated, a wavelet is a function of time with a localized response centered at zero. A wavelet that is admissible for use in CWT must be square integrable with finite, non-zero energy. In other words, a wavelet $\Psi(t,\tau,s)$ that is a function of time t, centered at time τ and with scale s, has a total energy

$$E = \int_{-\infty}^{\infty} |\Psi(t,\tau,s)|^2 dt \tag{2}$$

which must converge to a finite, non-zero, constant. An example of an admissible wavelet, known as the Ricker Wavelet, is shown in Figure 2.

It follows that the convolution of a wavelet with a digital signal is a finite impulse response filter, causing certain frequency ranges to be attenuated while leaving others unchanged. Since the characteristics of the wavelet are known, the frequencies and energy that are allowed to pass

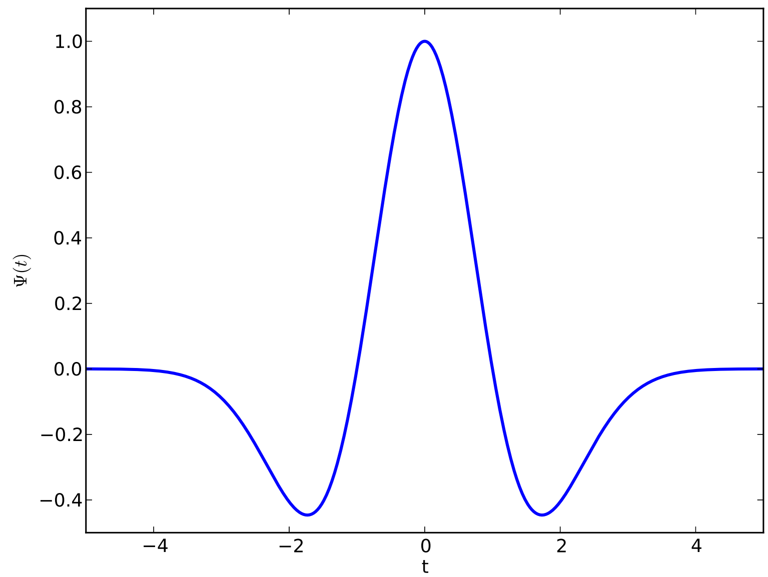

Figure 2: A Ricker Wavelet is admissible for use in CWT.

through the filter can be determined with a high degree of accuracy. In practice, wavelet coefficients for a discrete signal are defined as

$$C(\tau,s) = \sum_{t=0}^{n} x(t)\Psi(t,\tau,s), \tag{3}$$

where $C(\tau,s)$ is the wavelet coefficient centered at time τ, with scale s and $x(t)$ is the sample signal at time t and n is the length of the EEG segment. The energy content of an EEG signal, typically measured in micro-volt squared seconds per hertz ($\mu V^2 \cdot s \cdot Hz^{-1}$), can be estimated across time using wavelet coefficients at various times and scales. Figure 3 shown an example of a spectrogram estimating the energy content of an EEG segment. Addison provides a formal and thorough description of CWT [39].

Since a CWT represents a signal in terms of wavelets, which have a localized response, they are better suited for representing non-stationary signals. In other words, they typically achieve a better combination of time and frequency resolution than can be obtained with a series of PSD estimated using DFT. Furthermore, many potential wavelets can be used and various combinations of time and scale parameters can be explored. This means that CWT can be tuned to have very specific properties and to estimate the energy content of a signal at a wide range of resolutions.

Zhiwei and Minfen have explored the use of representations derived using the Shannon Entropy of the energy in the δ, θ, α and β bands computed using CWT [40]. Support Vector Machines (SVM) with linear, polynomial and radial basis function kernels were then used to classify the resulting features, with linear kernels typically achieving the best results. Zhiwei and Minfen achieved classification accuracies between 67–100% for two task problems and 65–92% for four task problems using one-second intervals overlapping by 0.8 seconds for two subjects in the dataset recorded by Keirn and Aunon [8].

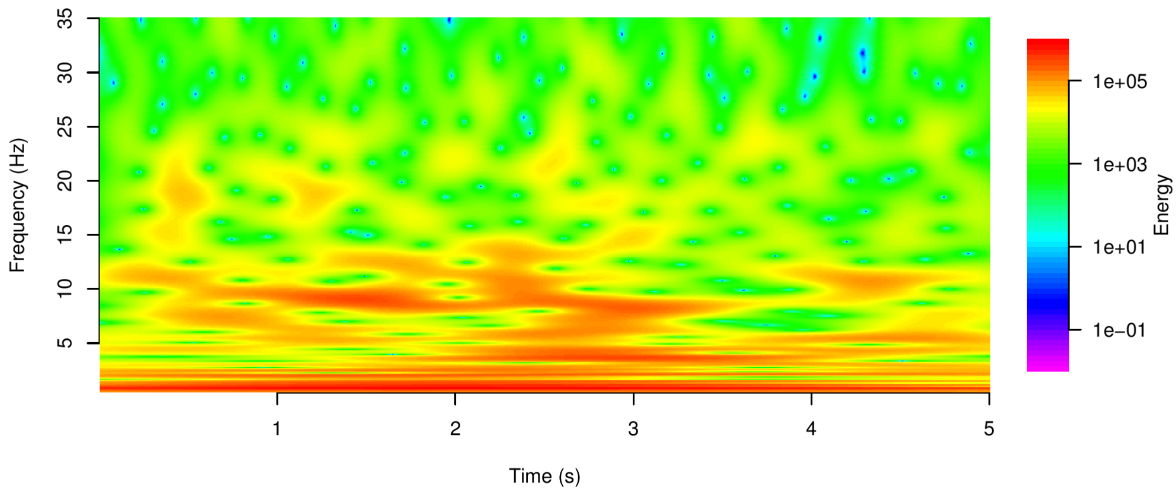

Figure 3: An example of a spectrogram of an EEG segment generated using a CWT (generated using code provided by Charles Anderson).

Iáñez, et al., also explore the use of CWT to represent EEG produced during imagined left and right motor imagery along with resting state [41]. In this work, raw CWT coefficients are classified using a number of pair-wise Linear Discriminant Analysis (LDA) classifiers, similar to the QDA classifier described in the previous section but with a pooled covariance estimate [34]. Final class labels are assigned by applying a scoring system to the output of the pair-wise LDA outputs. Iáñez, et al., achieve classification accuracies in the 20–60% range at one-second intervals and demonstrate a limited ability for participants to remotely control a robotic arm.

Although representations that use CWT may be better able to handle non-stationary patterns EEG signals than methods that rely on DFT, they still suffer from a number of limitations. First, there still exists a trade-off between dimensionality and resolution. CWT with high resolution may yield as many features as the raw EEG signal while CWT with low resolution may not sufficiently capture differences in EEG signals produced during various mental tasks. Classifiers that require accurate covariance estimates, such as LDA, may be especially sensitive to this problem. It is possible that Iáñez, et al., may achieve better classification results if they use a regularized version of LDA or another classifier altogether [38]. Alternately, examining the trade-off between resolution and dimensionality may yield better results. The works described in this section also suffer from the same problems that a lack of phase related information may cause in PSD representations. Although CWT could potentially be used to extract phase information, there does not appear to be any research that currently explores this possibility in the context of MT-based BCI.

2.3 Phase Locking Value

Although EEG signal representations that rely on estimates of power or energy often disregard phase information, a number of approaches have been explored for quantifying cross-channel phase synchronization. One such metric, originally proposed for use in EEG by Lachaux, et al., and by Mormann, et al., is known as Phase Locking Value (PLV) [42, 43]. PLV measures the amount of change in instantaneous phase between two signals over short window. Let $x_m(t)$ be a univariate signal sampled from channel m at time t. Then the instantaneous phase of $x_m(t)$, i.e., the angle through the signal's instantaneous period at time t, can be written as

$$\Phi_m(t) = tan^{-1}\left(\frac{\hat{x}_m(t)}{x_m(t)}\right) \tag{4}$$

where $\hat{x}_m(t)$ is the Discrete Hilbert Transform of $x_m(t)$, which yields a signal with a $\pi/2$ shift in instantaneous phase. The difference in instantaneous phase between two signals, $x_1(t)$ and $x_2(t)$, can then be expressed as

$$\Delta\Phi(t) = \Phi_2(t) - \Phi_1(t). \tag{5}$$

Finally, the PLV over a window of width w is

$$P = |\langle e^{i\Delta\Phi(t)} \rangle_w| \tag{6}$$

where $\langle\cdot\rangle_w$ denotes component-wise averaging over a window of width w and $|\cdot|$ denotes the complex modulus. Equation (6) expresses $\Delta\Phi(t)$ as a vector with unit length on the complex plane and then averages its vector components over a short window. This means that if $x_1(t)$ and $x_2(t)$ have a constant phase difference, then the PLV will be one. If, on the other hand, the instantaneous phase difference between $x_1(t)$ and $x_2(t)$ is randomly distributed, then the PLV will be near zero. Figure 4 shows a plot of two EEG signals filtered between $30-50$Hz, the instantaneous phase for each signal and the PLV between the two signals with $w = 0.5$ seconds. Notice that the signals tend to have a relatively high PLV even though the signal amplitudes are quite different.

PLV representations include information that PSD and CWT cannot express in the absence of phase information. It is important to note, however, that PLV do not express amplitude information in themselves. PLV also cannot represent all forms of spatial patterns, even when combined with amplitude information. This is because PLV measures changes in phase over a window instead of exact phase differences. For example, if two signals have a constant phase shift of $\pi/4$ during one mental task and $\pi/6$ during another mental task, the PLV will be one in both cases. Furthermore, since PLV measures phase information between pairs of channels, there are $\binom{c}{2}$ PLV for an EEG montage consisting of c channels, yielding relatively high dimensionality.

Although it does not seem likely that PLV alone will be able to represent the necessary information present in EEG signals, combining PLV with amplitude features may, at least partially, overcome some of the limitations found in the previous two sections. Gysels and Celka have investigated this possibility for use in MT-based BCI [44]. In this work, a number of combinations of PLV and PSD features were combined to represent EEG signals recorded during three mental tasks, including two motor imagery tasks and a word generation task. PLV dimensionality was reduced by averaging the features for nearby channels over several brain regions. PSD dimensionality was reduced by focusing solely on α and β frequency bands. Gysels and Celka performed classification using linear SVM applied to sliding one-second windows with majority voting to assign class labels after 8 successive windows. This work suggests that combining PLV and PSD representations may yield better performance than the use of PLV or PSD alone. However, only modest classification results, in the 60–75% range, were achieved.

Several other works have explored the use of PLV in motor imagery BCI with somewhat conflicting results. Krusienski, et al., have investigated the performance of a two-class motor imagery

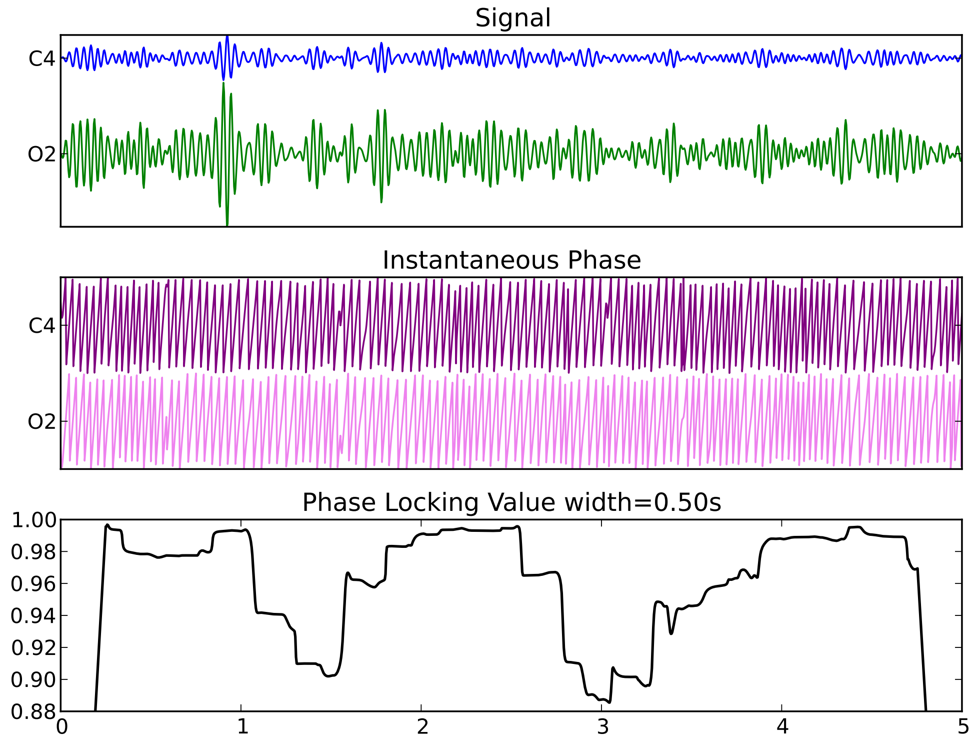

Figure 4: Phase Locking Value between sites C4 and O2 during five-seconds of β-band EEG shown along with instantaneous phase and the raw EEG signal.

BCI when using PLV representations with and without amplitude information from PSD [45]. In order to classify the resulting representations, Step-Wise Linear Discriminant Analysis (SWLDA) was used. SWLDA is a variant of LDA that uses an iterative feature selection procedure [46]. The results obtained by Krusienski, et al., appear to suggest that PLV performs worse than PSD alone and that combining the two representations does not yield a significant improvement in performance. On the other hand, Hamner, et al., have concluded that using instantaneous phase, as described in equation (4), can significantly increase classification performance for a two-class motor imagery BCI when combined with amplitude features and a Naive Bayes classifier. This work achieved classification accuracies between 84–99% correct for one-second sliding windows [47]. In any case, performance in motor imagery BCI may not be representative of performance in the more general MT-based BCI. Clearly, more research is required in order to determine the effectiveness including PLV or other types of phase information in frequency-domain representations.

3 Time-Domain Representations

A number of options also exist for representing EEG signals in the time-domain. In fact, raw EEG signals are recorded in the time-domain in terms of signal voltages, typically measured in microvolts ($\mu \cdot V$), across EEG sensors and through time. Raw EEG may be thought of as a multivariate time-series written as

$$\mathbf{X} = \{\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_n\} \tag{7}$$

where $\mathbf{X}_t$ is a column vector with length equal to the number of EEG channels, c, and where n is the total length of the EEG segment. Figure 5 shows five seconds of an eight-channel raw EEG segment in the time-domain.

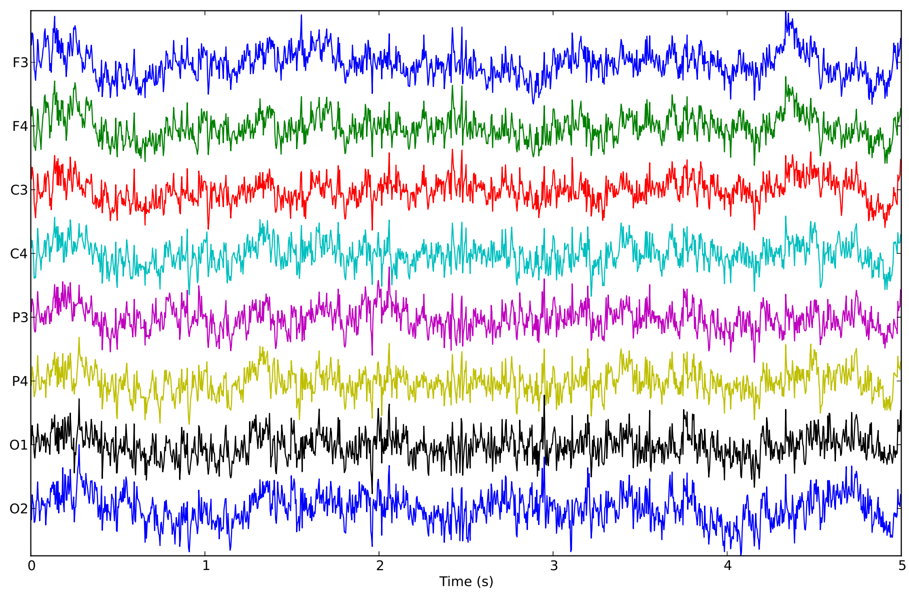

Figure 5: Raw EEG trace in the time-domain.

In contrast to frequency-domain representations, it is relatively uncommon to represent spontaneous EEG signals in the time-domain. Nevertheless, time-domain representations have a number of practical advantages that may make them successful in the context of asynchronous BCI. For instance, time-domain representations may be well-suited for representing spatial patterns across EEG electrodes because they utilize the voltages measured at each site rather than an estimate of power or energy over a period of time. Time-domain representations also do not draw a distinction between amplitude and phase across channels. Furthermore, time-domain representations do not necessarily rely on assumptions of periodicity. Although oscillatory components of EEG signals are clearly related to mental state in a number of ways, there is no evidence that EEG signals do not contain patterns that are not well represented by sinusoids or wavelets. When smoothing occurs in the frequency-domain, some of these patterns may be lost.

On the other hand, it can be difficult to interpret and visualize time-domain representations of spontaneous EEG signals, leading to difficulties determining the biological correlates of the underlying mental state. The ease of interpreting frequency-domain representations is likely one of the reasons for their popularity. It can also be more challenging to incorporate temporal information into time-domain representations. When done properly, however, time-domain representations can precisely represent temporal patterns on short time scales, a challenge in the frequency-domain.

In the remainder of Section 3, we examine several approaches for representing EEG signals in the time-domain. We also discuss the classification algorithms that have been used in combination with these representations. It appears that time-domain representations are promising because of their ability to capture precise patterns. However, these representations often yield high-dimensionality and it can be challenging to regularize classifiers in this setting or to reduce the dimensionality in a meaningful way. It can also be difficult to attain temporal invariance with time-domain representations and fully capture patterns across a range of time scales.

3.1 Time-Delay Embedding

One method for incorporating temporal information into time-domain representations is to embed past signal values into the raw EEG time-series. This technique is known as Time-Delay Embedding (TDE). Constructing a TDE representation requires two parameters. First, the embedding dimension determines the number of past values to be included and the length of temporal patterns than can be captured. Second, the amount of overlap between adjacent windows offers a trade-off between the number of training examples and the amount of redundant information.

In general, assuming that the length of the raw EEG segment is a multiple of the number of windows, a TDE representation can be written as

$$\mathbf{Y}^{d,l} = \left\{ \begin{bmatrix} \mathbf{X}_d \\ \vdots \\ \mathbf{X}_2 \\ \mathbf{X}_1 \end{bmatrix}, \begin{bmatrix} \mathbf{X}_{2d-l} \\ \vdots \\ \mathbf{X}_{d-l+2} \\ \mathbf{X}_{d-l+1} \end{bmatrix}, \begin{bmatrix} \mathbf{X}_{3d-2l} \\ \vdots \\ \mathbf{X}_{2d-2l+2} \\ \mathbf{X}_{2d-2l+1} \end{bmatrix}, \ldots, \begin{bmatrix} \mathbf{X}_n \\ \vdots \\ \mathbf{X}_{n-d+1} \\ \mathbf{X}_{n-d} \end{bmatrix} \right\} \tag{8}$$

where $\mathbf{Y}^{d,l}$ is the resulting TDE time-series with embedding dimension d, window overlap l and where $\mathbf{X}$ is the non-embedded time-series, as described in equation (7).

Although TDE is capable of capturing both spatial and temporal patterns, this approach leads to high-dimensional representations. TDE representations also lack temporal invariance. In other words, every dimension of the column vectors of the TDE representation, Y, must be exposed to many possible signal values in order to fully characterize temporal relationships in the signal. As a simple example, consider a sinusoid that is oscillating at a constant frequency. If a classifier is trained using TDE windows that always start at the beginning of the signal's period, then poor performance may be achieved when windows are encountered that start at a different time in the signal's period. This is in contrast to PSD and CWT representations which describe the signal in terms of periodic components. Furthermore, the length of a temporal pattern that can be explicitly expressed using TDE is limited by the embedding dimension, d. Combined, this means that large amounts of data must be processed by the classifier and that steps must be taken to prevent overfitting.

Nevertheless, the flexibility of TDE may yield good results. Anderson, et al., assert that TDE representations perform well when combined with non-linear classifiers and sequential evidence accumulation [48]. In this work, a two-layer feed-forward Artificial Neural Network (ANN) with a hyperbolic tangent transfer in the first layer and a multinomial logistic regression in the second layer was compared with LDA and QDA for classifying TDE representations. Final class labels were assigned using an M-ary Sequential Probability Ratio Test (MSPRT) to accumulate evidence across a number of TDE windows [49]. This procedure yields a variable decision rate with the BCI only issuing instructions after it has become sufficiently confident in the user's intent. In offline experiments using a single subject and four mental tasks, Anderson, et al., estimate classification performance to be roughly three seconds per correct decision, with ANN significantly outperforming LDA and QDA.

Presumably, ANN are able to outperform QDA and LDA and achieve good results using TDE for two reasons. First, ANN are able to find complex non-linear patterns in EEG signals that are not well represented using linear or quadratic methods. Second, the ANN used in this case had far fewer parameters to optimize, yielding a regularized classifier.

Despite the encouraging results achieved by Anderson, et al., it seems likely that large amounts of training data and careful regularization may be required to achieve performance that is acceptable for practical BCI. Larger studies are required in order to make firm conclusions about the performance that can be expected from this approach. Furthermore, other methods for representing EEG signals in the time-domain that yield lower-dimensionality and a greater degree of temporal invariance should be explored.

3.2 TDE with Signal Transformations

One possible method for addressing the problems with TDE involves the use of linear transformations to separate the signal into a number of components that each have different characteristics. Components that are deemed to be unnecessary, either because they contain noise, artifacts, redundant information or irrelevant information, can then be dropped to yield a lower-dimensional representation. A linear signal transformation can be written as

$$\hat{\mathbf{Y}}_t = \mathbf{T}\mathbf{Y}_t \tag{9}$$

where $\mathbf{Y}_t$ is the TDE representation at time t, $\mathbf{T}$ is a linear transformation matrix and $\hat{\mathbf{Y}}_t$ is the transformed signal at time t. The rows of $\hat{\mathbf{Y}}_t$ then contain the components of the transformed signal.

Many methods for constructing $\mathbf{T}$ have been explored and signal transformations constitute an entire field of research in themselves. As a result, a thorough description of these techniques is largely beyond the scope of this research exam. However, signal transformations have been successfully used in BCI and it is worth briefly discussing these techniques and the results they achieve.

In several studies, Anderson, et al., have proposed the use of Maximum Signal Fraction (MSF) and Principal Components Analysis (PCA) in combination with TDE [50, 51, 52]. MSF is a signal transformation that separates a signal into components ordered by signal to noise ratio, measured as the change in signal value between each time-step. Figure 6 shows the MNF components of a five-second EEG segment with eight channels. PCA is a linear signal transformation that separates the signal into components that are ordered by variance. These works suggest that MSF can be used to produce filters for removing artifacts and that PCA can be used to reduce dimensionality while improving classification accuracy. Using LDA and committees of decision trees, classification accuracies in the 68–78% range were achieved for five mental tasks with class labels assigned roughly between 1-2 seconds.

Zhang, et al., suggest the use of Independent Components Analysis (ICA) for eliminating EEG artifacts while preserving γ-band information [53]. ICA is a signal transformation that attempts to maximize the statistical independence between the resulting components. After filtering in the time-domain, Zhang, et al., generated PSD features and performed classification using LDA. The

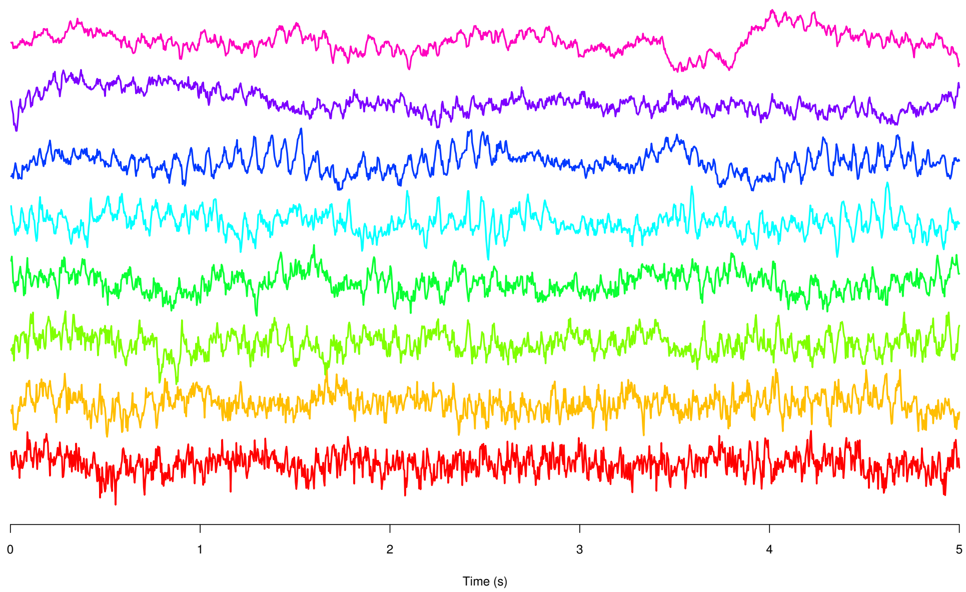

Figure 6: MSF components of a five-second, eight-channel EEG segment (generated using code provided by Charles Anderson).

resulting classification accuracies ranged from 38–69% for a five-class problem and 62–86% for a two-class problem at one-second intervals.

Friedrich, et al., have explored the use of Common-Spatial Patterns (CSP) for reducing the dimensionality of TDE representations [54, 23, 55]. CSP is a supervised method for generating a linear transformation that separates a signal into components that are ordered by the ratio of the variance between classes. Components that appear to have little or no between-class variance can then be removed. CSP has also been explored in the context of motor imagery BCI [56, 57]. Using LDA and CSP, Friedrich, et al., achieved a mean classification accuracy of 50% for four mental tasks across 14 participants with decisions made in real-time every two seconds. The resulting classifiers appear to be stable for as long as 10 weeks for some users.

Although signal transformations have been successfully used to reduce the dimensionality of TDE representations, these approaches have a number of drawbacks. First, signal transformations do not address the lack of temporal invariance associated with TDE. Each dimension of the representation must encounter a variety of signal values in order for the dynamics of the signal to be fully captured. Second, it can be extremely difficult to know which method should be used and which components should be removed. Finally, these methods often involve a large amount of manual tuning. Determining the components to be removed through trial-and-error is not acceptable for BCI that are to be used by non-experts.

3.3 Autoregressive Models

Linear Autoregressive (AR) models are an alternative to TDE for representing EEG signals in the time-domain. AR models represent a time-series as a linear process. In other words, each value in a time-series is modeled as a linear combination of previous values plus a residual term. If an EEG signals are viewed as time-series, as in equation (7), then AR models can be used to fit the signal. Formally, a multivariate AR model can be written as

$$\mathbf{A}_t = \sum_{i=0}^{p} \mathbf{C}_i \mathbf{A}_{t-i} + \mathbf{R}_t \tag{10}$$

where $\mathbf{A}_t$ is the value of the model time-series at time t, $\mathbf{C}_i$ is a row vector of model coefficients at i time-steps back in time, p is the order of the model and $\mathbf{R}$ is the residual term. A number of techniques can be used to estimate AR coefficients [58].

AR models have been studied extensively for modeling and forecasting time-series and have been successfully applied to a number of real-world problems [58]. AR may also yield more temporal invariance than TDE since AR models describe the time-series in terms of past values. Additionally, AR may lead to lower dimensional representations than TDE because only the AR coefficients and residuals are required to describe the model. AR models can also act as a filter, removing aspects of the time-series that are not linearly predictable. For these reasons, a number of research groups have explored the use of AR models for representing EEG signals in the time-domain.

Keirn and Aunon have proposed that AR coefficients may be used to represent EEG signals directly [8]. In this approach, univariate AR models are computed for two-second EEG segments for each EEG channel and then classification is performed using the AR coefficients and QDA. Since the AR coefficients determine the AR model, except for the residual term, they yield a filtered and lower dimensional representation of the EEG signal. Keirn and Aunon also note that AR coefficients can be used to estimate the PSD of an EEG signal, implying that these representations contain similar information. This method was compared side-by-side to the PSD asymmetry ratios described in Section 2.1. Keirn and Aunon achieved classification accuracies between 78–88% correct at two-second intervals for various pairs of imagined mental tasks. In these experiments, AR models achieved a roughly 3.6% improvement in classification accuracy over the PSD representations.

Anderson, et al., extended the work of Keirn and Aunon by using feedforward ANN trained using backpropagation and early-stopping for classifying AR coefficients [59, 60]. Multivariate AR models were compared to univariate models and PCA was explored for reducing dimensionality. Two imagined mental tasks were explored and class labels were assigned at one-second intervals. These works achieved classification accuracies in the 88–96% range. Multivariate AR models produced better results on average, although only slightly, and PCA did not typically improve performance. Although the ANN used by Anderson, et al., achieved higher classification accuracies than the QDA models used by Keirn and Aunon, very few hidden units were required. This suggests that the class boundaries may have few non-linear regions.

Curran, et al., proposed the use of reflection coefficients, which are related to the change in model error as the order of an AR model increases [61]. In this work, Logistic Regression was compared with a non-linear Bayesian classifier trained using Variational Learning. An analysis was performed on EEG recorded during four imagined mental tasks. The electrodes used were hand picked and classification was only attempted for two-task combinations using one-second windows. Ultimately, this approach achieved modest classification accuracies, between 63–74%. Curran, et al., observed that their non-linear Bayesian classifier achieved an average of about 2% higher classification accuracy than logistic regression.

Despite the fact that encouraging results have been achieved on two-class problems using AR models, it remains unclear whether or not the assumption of linearity is beneficial or harmful. On one hand, AR models may be viewed as a filter by assuming that non-linear relationships in the signal are largely noise. On the other hand, non-linear patterns may be important for discriminating between mental tasks and representing an EEG signal as a linear process may eliminate some of this information. Furthermore, the order of AR models may limit their ability to capture long-term patterns in EEG signals. Although, the works discussed in this section conclude that orders larger than 10 do not typically improve performance, there may be a trade-off between model order and other forms of regularization.

3.4 Time-Series Modeling with Artificial Neural Networks

If non-linear patterns in EEG signals are important for discriminating between mental tasks, then linear AR models may not be a sufficient representation. Since the extent to which non-linear patterns are useful in BCI is still largely unknown, empirical evidence gathered by experimenting with non-linear representations and classifiers is important [62]. Artificial Neural Networks (ANN) appear to be ideal candidates for constructing non-linear time-series models because the are well known for their ability to solve non-linear regression problems [63].

Performing classification using these models can, however, be challenging. Unlike AR models, the weights of an ANN trained to forecast an EEG signal are not a unique representation of the model obtained. This is because ANN contain many symmetries; i.e., there are many weight combinations that can yield an equivalent network. Furthermore, the process of optimizing the network weights is typically stochastic and non-convex. Gupta, Oeda, Coyle and Forney have all proposed that time-series may be classified using forecasting errors instead [64, 65, 66, 67]. In this approach, a separate network is trained to model example time-series data from each class, yielding one model per class. Each model is then viewed as an expert at forecasting the type of time-series over which it was trained. When novel data is encountered, each model is applied and the resulting residuals are extracted. These residuals are then used to determine which model fits the data best and, ultimately, assign class labels.

One method for modeling EEG signals using ANN involves the use of TDE. In this approach, a feedforward network is trained to continually forecast the signal using a window of past signal values embedded into the network inputs. Coyle, et al., have explored such models for use in two-class motor imagery BCI with favorable results [66]. In this work, two-layer ANN trained using backpropagation and early-stopping were used to model the EEG signals and LDA was subsequently used to classify the resulting model residuals. This method resulted in classification accuracies between 85–91% correct at five-second intervals. Additional experiments suggest that this approach outperforms linear ANN and AR models; although there are likely too few subjects to draw firm conclusions. In later works, Coyle, et al., have proposed non-linear time-series models using ANN for filtering EEG signals, with the intuition being that aspects of the signal that are not predictable are likely noise [68, 69, 70, 71].

Although the approach used by Coyle, et al., yields non-linear time-series models, the use of TDE may result in some of the same problems discussed in Section 3.1. Forney and Anderson suggest that Artificial Recurrent Neural Networks (RNN) may be well-suited for generating non-linear time-series models while avoiding TDE [67, 72, 73]. RNN are a type of ANN that contain feedback loops, also known as recurrent connections [63]. These recurrent connections allow RNN

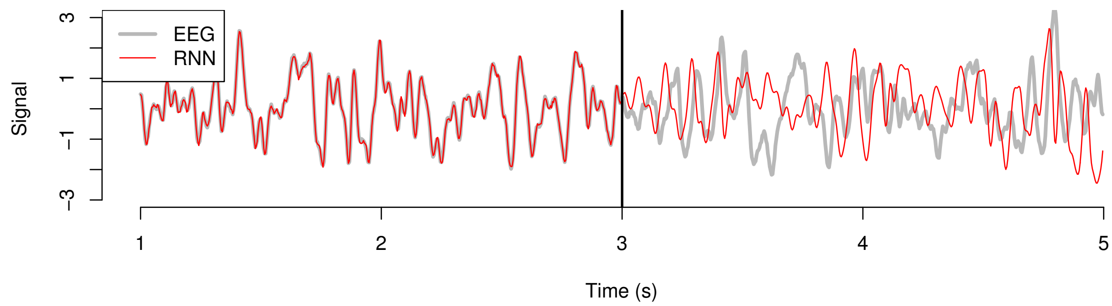

Figure 7: An Elman Recurrent Network transitioning from forecasting to an iterated model at the eight-second mark.

to have memory and maintain state without requiring TDE. Forney and Anderson have explored the use of two kinds RNN: Elman Recurrent Networks and Echo State Networks. In these works, it was demonstrated that RNN are able to forecast EEG signals with a relatively high degree of accuracy. Forney and Anderson have also explored the dynamics that RNN are able to capture in EEG signals by placing a feedback loop from the network outputs back into the network inputs, known as autonomous or iterated models. When these models were constructed with a large number of artificial neurons, they were able to generate complex dynamics that were not clearly periodic and had energy spectra similar to the underlying EEG signals. Figure 7 shows an Elman Network forecasting an EEG signal before the eight-second mark and transitioning to an iterated model after the eight-second mark.

Forney and Anderson have demonstrated that simply selecting the class label associated with the model that produced the lowest forecasting error outperformed LDA and QDA for assigning class labels [72]. These methods were tested offline for various four-task and two-task problems using data that was recorded from subjects with no motor impairments as well as from subjects with severe motor impairments in their home environments. The resulting classification accuracies varied between 15–68% for four mental tasks and 50–95% for two mental tasks. A statistically significant difference was found between these two groups of potential users.

Although non-linear time-series models appear to be able to capture complex patterns in EEG signals, there are a number of drawbacks to these approaches that may make them less than ideal for use in practical BCI. First, neural networks, and particularly recurrent networks, often require large amounts of computational resources. Although the Echo State Networks explored by Forney and Anderson can be trained very quickly, they have many parameters that must be tuned manually. Since BCI must be able to operate in real-time, this may be prohibitive. Second, these models often have difficulties learning long-term patterns. The necessity for learning long-term patterns may depend on the sampling rate of EEG signal. In other words, these models may have difficulties capturing patterns at both low and high frequencies simultaneously. In order to remedy this problem, methods should be sought for modeling EEG signals at multiple time-scales. Finally, interpreting and visualizing the patterns found by these models can be difficult.

4 Discussion

This research exam has looked at a number of methods for representing and classifying EEG signals within the framework of BCI that utilize imagined mental tasks. This analysis has lead to a number of conclusions about the capabilities and limitations of each of these methods while also raising a number of unanswered questions. Empirical evidence was also gathered regarding the performance of these approaches in various offline and online studies. In this section, the results that have been collected throughout these studies is summarized and a final overview of each of method is given. Finally, some potential directions for future research are suggested.

4.1 Summary and Conclusions

Table 1: Summary of Classification Accuracies Across Studies.

Author	Paper	Representation	Classifier	# Tasks	Delay	Accuracy
Keirn	[8]	PSD	QDA	2	2s	76–84%
Millán	[22]	PSD	Mixture of Gaussians	3	0.5s	>70%
Zhiwei	[40]	CWT	Linear SVM	2	1s	67–100%
Zhiwei	[40]	CWT	Linear SVM	4	1s	65–92%
Iáñez	[41]	CWT	LDA	2	1s	20–60%
Gysels	[44]	PLV+PSD	Linear SVM	3	1s	60–75%
Anderson	[48]	TDE	ANN	4	3s per correct	
Anderson	[52]	TDE+PCA	LDA	5	1-2s	68–78%
Zhang	[53]	PSD+ICA	LDA	2	1s	62–86%
Zhang	[53]	PSD+ICA	LDA	5	1s	38–69%
Friedrich	[23]	TDE+CSP	LDA	4	2s	Mean 50%
Keirn	[8]	AR	QDA	2	2s	78–88%
Anderson	[60]	AR	ANN	2	1s	88–96%
Curran	[61]	AR	Variational Learning	2	1s	63–74%
Coyle	[66]	Time-Series+ANN+LDA		2	5s	85–91%
Forney	[72]	Time-Series+ERN+Best Fit		2	1s	52–93%
Forney	[72]	Time-Series+ERN+Best Fit		4	1s	28–68%
Forney	[73]	Time-Series+ESN+Best Fit		2	2s	50–95%
Forney	[73]	Time-Series+ESN+Best Fit		4	2s	15–65%

The empirical results that were found in the various works reviewed in this research exam are summarized in Table 1. Unfortunately, drawing conclusions from across these studies is difficult for a number of reasons. First, different EEG data was used in the majority of the studies reviewed. This means that differences between individuals and between recording methods may be responsible for much of the variation seen in classification performance. Second, many of these studies use a different number of subjects, possibly contributing to the difference in variability seen among the methods. Third, many of these studies use different delays between successive class label assignments. Typically, there are trade-offs between the rate at which class labels are assigned, classification performance and user experience. Finally, the number and kind of mental tasks used is different between many of these studies. Although a number of methods that better represent BCI performance have been proposed, few of these studies report results using these metrics [74].

Nevertheless, some very general conclusions can be drawn from Table 1. It appears that some of the highest classification accuracies were achieved using CWT representations with linear SVM classifiers. Time-series models using either AR models with ANN or recurrent networks with best fit classifiers also appear to achieve high classification accuracies. It is also important to mention that TDE combined with ANN and LDA achieved good results using four and five mental tasks. Linear classifiers, namely LDA and linear SVM, were the most common choices. QDA and other Bayesian classifiers as well as ANN were also used in a number of these studies. Unfortunately, it does not seem reasonable to draw more detailed conclusions from across these studies. This means that the relative performance of these methods remains largely unknown.

Despite the difficulties encountered when attempting to draw conclusions across these studies, useful experiments were performed within these studies. This includes a few direct comparisons between signals representations and classification algorithms. Additionally, analyzing these methods has lead to a number of interesting observations and conjectures.

This analysis began by investigating methods for representing EEG signals in the frequency-domain and by looking at the classification algorithms that have been used with these representations. Power Spectral Densities (PSD) are one such representation that have been used in other disciplines for analyzing spontaneous EEG signals. There are also several well-known neurophysiological phenomenon that have been identified in PSD that have been used to construct BCI. However, PSD representations appear to suffer from at least three limitations that may prevent them from expressing some forms of patterns in EEG signals. First, PSD typically undergo a smoothing process. Although smoothing reduces the dimensionality of the representation, it is possible that this process removes important information. The trade-off between smoothing, dimensionality reduction and resolution has not been explored in the work reviewed. Second, PSD representations typically discard phase information. This prevents them from capturing some forms of spatial patterns that may occur across EEG sensors. Third, since PSD are typically derived by representing the signal as a sum of sinusoids, they may have poor resolution through time. This is particularly problematic because EEG signals are non-stationary. Among the studies that were reviewed, Bayesian classifiers, such as LDA, QDA and Mixture-of-Gaussian Models, appear to be popular choices. However, the assumptions of normality that they rely upon were not verified. Furthermore, regularization parameters and other non-linear classifiers were rarely reported.

Continuous Wavelet Transforms (CWT) are another option for representing EEG signals in the frequency-domain. Since CWT rely on wavelets, which have a localized response, and since EEG signals are non-stationary, they often have higher resolution than PSD. However, CWT representations also have a trade-off between smoothing and resolution that has yet to be thoroughly explored. Furthermore, the CWT that were examined did not include phase information, leaving them susceptible to some of the same problems found in PSD. In the research reviewed, LDA and SVM with several kernels were explored. Although linear SVM outperformed non-linear kernels in one study, a detailed examination of regularization parameters and a comparison with other non-linear classifiers was not performed.

Phase Locking Value (PLV) was the final frequency-domain representation examined. Although PLV may allow some phase-related patterns to be captured, it only measures a moving average of changes in instantaneous phase. This means that a phase differences that are stable over a relatively short period of time cannot be captured. Again, only LDA and linear SVM were explored for assigning class labels and a detailed analysis of regularization parameters was not presented. Although the results from these experiments were somewhat conflicting, it appears that PLV was useful for discriminating between some mental tasks.

This analysis of frequency-domain representations suggests that these methods are able to capture oscillatory components of EEG signals. These representations also align well with many EEG analysis techniques used in other disciplines. However, each of these methods may be limited by a lack of ability to capture some forms of spatial or temporal patterns. Furthermore, there is a lack of work thoroughly exploring the various parameters that should be tuned in the representations and classification algorithms. Clearly, these questions must be answered before it can be determined if the limitations of these representations are significant.

A number of approaches were then examined for representing EEG signals in the time-domain. Time-Delay Embedding (TDE) is capable of capturing both spatial and temporal patterns. However, TDE can be difficult to interpret, yields high dimensionality and is not invariant to shifts in time. Nevertheless, good performance was achieved using ANN and evidence accumulation in an approach that relied more heavily on the classifier rather than the signal representation for identifying patterns.

Linear Signal Transformations can be used to remove artifacts and reduce the dimensionality of TDE representations. However, it is difficult to know without more empirical evidence which transformations to use and which signal components to remove. Furthermore, these approaches may not be suitable for use in practical BCI if they cannot be fully automated. Among the studies reviewed, LDA was used predominantly; although one study examined the use of Decision Trees. Presumably, these approaches relied more heavily on the representation than on the classifier. It is difficult to know, however, if other classifiers would achieve better results when combined with these representations.

Time-series models are another option for representing and classifying EEG signals. These approaches have the advantage of acting as filters by removing aspects of the signals that are not predictable. Additionally, they may yield more temporal invariance and lower dimensionality than TDE. First, linear Autoregressive (AR) models were examined for this role. AR models can be used to represent EEG signals using only the model coefficients and are well-understood. It remains unclear, however, if the assumptions of linearity required by AR models are valid. QDA, ANN, Logistic Regression and Variational Learning were all examined for classification. Although ANN achieved the best results, only a few hidden units were required, suggesting that the class boundaries may be only slightly non-linear.

Time-series models using ANN were also explored in order to investigate the possibility that non-linear models may be better at modeling EEG signals than linear AR models. Feedforward networks were explored for this role and were demonstrated favorably when compared to linear AR models in one study. Recurrent networks were also explored and it has been demonstrated that they are able to capture complex patterns in EEG signals without requiring TDE. However, it is difficult to asses the advantages of using recurrent networks because there is a lack of direct comparison with AR models and feedforward networks. Additionally, the ability to train and tune recurrent networks in a real-time BCI is questionable. Both linear and non-linear time-series models may have difficulties modeling EEG signals across a range of time scales.

This analysis of time-domain representations suggests these methods may be well-suited for capturing spatial and non-periodic patterns. However, this often comes at the expense higher dimensionality and difficult interpretation. Unfortunately, there appears to be a general lack of direct comparison and real-time experiments among all of the methods explored in this research exam. Although there are a number of potential advantages and disadvantages to each of these approaches, it remains unknown what types of patterns are required for distinguishing between imagined mental tasks. It is therefore difficult to determine the types of signal representations and classification algorithms that are suitable for use in BCI.

4.2 Potential Research Directions

A number of open questions regarding the topics discussed in this research exam remain unanswered and should be addressed in future work. Clearly, more comparisons of EEG signal representations and classification algorithms should be performed. Direct comparisons of each of the signal representations discussed should be performed on a single dataset consisting of EEG recordings from a number of subjects each performing multiple imagined mental tasks. Such comparisons would help to determine which types of patterns found in EEG signals are important for discriminating mental tasks. For example, if PSD and CWT perform as well as TDE and multivariate time-series models, then it would appear that differences in phase across channels may not be important in this setting.

Another interesting possibility that should be explored is the inclusion of phase information in PSD and CWT representations. Since it is possible to extract phase information using Discrete Fourier Transforms and Wavelet Transforms, the use of this information should be considered. Such approaches may offer advantages over PLV since they might be able to express constant differences in phase and because they may not need to generate a separate metric for each pair of channels in the EEG montage.

A more thorough comparison of classification algorithms should also be performed. Currently, it seems unclear whether or not linear classifiers yield better results than non-linear classifiers; although it certainly depends on the signal representation used. This experimentation would help to determine if non-linear patterns in EEG signals are important for discriminating mental tasks. The assumptions that these classifiers rely upon should also be checked. For example, the assumptions of normality that are required by LDA and QDA should be verified before concluding whether these classifiers are appropriate or not.

Unfortunately, adequate exploration and tuning of the various parameters in the methods explored in this research exam were not often presented. For example, the trade-off between resolution and dimensionality was not explored for PSD and CWT representations. Additionally, the regularization parameters for the various classifiers used in these studies were often not explored. Without exploring these parameters using a validation procedure, it is difficult to know if one method outperforms another overall or if it simply yields better results in the absence of careful regularization.

Among the works investigated in this research exam, AR models and time-series models using feedforward and recurrent networks have achieved promising results. However, a number of questions remain to be answered about these approaches. First, it is not yet clear if non-linear time-series models are necessary. Although Coyle, et al., have found that ANN outperform AR models in a direct comparison, these results varied largely across subjects and only investigated two-class problems. Multi-class problems should be investigated using a larger number of subjects in order to determine if, overall, these methods are advantageous over linear AR models [66]. Forney and Anderson also assert that recurrent networks have advantages over feedforward networks and linear AR models but have not yet performed a direct comparison [72]. Furthermore, the ability and necessity for time-series models to capture long-term patterns in EEG signals is unknown. In order to explore this possibility, time-series models capable of modeling signals at multiple time-scales should be investigated as well as models that are able to achieve longer-term memory.

Finally, real-time experiments with users that have motor impairments are extremely important. Ultimately, the performance of methods for representing and classifying EEG signals should be measured in the context of actual BCI. Although offline evaluations and comparisons may be informative, the final word in the performance of these approaches depends on how well they work in practice, in real-time BCI and with potential users.

Acknowledgments

I would like to thank Kate Ericson and Maggie VanDenBerg for proofreading this research exam. I would also like to thank Maggie for helping watch our son Parker while I worked on this exam. I would also like to thank Charles Anderson for providing his CWT and MSF code. Last, but certainly not least, I would like to thank all of the people that participated in our BCI studies.

References

[1] Fred Plum and Jerome B Posner. *The Diagnosis Of Stupor & Coma*, volume 19. Oxford University Press, USA, 1982.

[2] J.D. Bauby. *The diving bell and the butterfly: A memoir of life in death*. Vintage, 1998.

[3] Jonathan Wolpaw and Elizabeth Winter Wolpaw. *Brain-computer interfaces: principles and practice*. Oxfort University Press, 2012.

[4] Adrian M Owen and Martin R Coleman. Detecting awareness in the vegetative state. *Annals of the New York Academy of Sciences*, 1129(1):130–138, 2008.

[5] Fiachra Matthews, Barak A Pearlmutter, Tomas E Ward, Christopher Soraghan, and Charles Markham. Hemodynamics for brain-computer interfaces. *IEEE Signal Processing Magazine*, 25(1):87–94, 2008.

[6] Leigh R Hochberg, Daniel Bacher, Beata Jarosiewicz, Nicolas Y Masse, John D Simeral, Joern Vogel, Sami Haddadin, Jie Liu, Sydney S Cash, Patrick van der Smagt, et al. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 485(7398):372–375, 2012.

[7] Eric C Leuthardt, Kai J Miller, Gerwin Schalk, Rajesh PN Rao, and Jeffrey G Ojemann. Electrocorticography-based brain computer interface-the seattle experience. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 14(2):194–198, 2006.

[8] Zachary A Keirn and Jorge I Aunon. A new mode of communication between man and his surroundings. *IEEE Transactions on Biomedical Engineering*, 37(12):1209–1214, 1990.

[9] P.L. Nunez and R. Srinivasan. *Electric fields of the brain: the neurophysics of EEG*. Oxford University Press, USA, 2006.

[10] AC Metting Van Rijn, A Peper, and CA Grimbergen. High-quality recording of bioelectric events. *Medical and Biological Engineering and Computing*, 28(5):389–397, 1990.

[11] Alexander C MettingVanRijn, Anthony P Kuiper, Taco E Dankers, and Cees A Grimbergen. Low-cost active electrode improves the resolution in biopotential recordings. In *Proceedings of the 18th Annual International Conference of The IEEE Engineering in Medicine and Biology Society*, volume 1, pages 101–102. IEEE, 1996.

[12] Lawrence Ashley Farwell and Emanuel Donchin. Talking off the top of your head: toward a mental prosthesis utilizing event-related brain potentials. *Electroencephalography and clinical Neurophysiology*, 70(6):510–523, 1988.

[13] Benjamin Blankertz, Steven Lemm, Matthias Treder, Stefan Haufe, and Klaus-Robert Müller. Single-trial analysis and classification of erp components–a tutorial. *Neuroimage*, 56(2):814–825, 2011.

[14] Eric W Sellers, Theresa M Vaughan, and Jonathan R Wolpaw. A brain-computer interface for long-term independent home use. *Amyotrophic lateral sclerosis*, 11(5):449–455, 2010.

[15] Joseph N Mak, Dennis J McFarland, Theresa M Vaughan, Lynn M McCane, Phillippa Z Tsui, Debra J Zeitlin, Eric W Sellers, and Jonathan R Wolpaw. Eeg correlates of p300-based brain–computer interface (bci) performance in people with amyotrophic lateral sclerosis. *Journal of neural engineering*, 9(2):026014, 2012.

[16] F Nijboer, EW Sellers, J Mellinger, MA Jordan, T Matuz, A Furdea, S Halder, U Mochty, DJ Krusienski, TM Vaughan, et al. A p300-based brain–computer interface for people with amyotrophic lateral sclerosis. *Clinical neurophysiology*, 119(8):1909–1916, 2008.

[17] P Brunner, S Joshi, S Briskin, JR Wolpaw, H Bischof, and G Schalk. Does the "p300" speller depend on eye gaze? *Journal of neural engineering*, 7(5):056013, 2010.

[18] J Kalcher, Doris Flotzinger, Ch Neuper, S Gölly, and G Pfurtscheller. Graz brain-computer interface ii: towards communication between humans and computers based on online classification of three different eeg patterns. *Medical and Biological Engineering and Computing*, 34(5):382–388, 1996.

[19] G Pfurtscheller, Ch Neuper, D Flotzinger, and M Pregenzer. Eeg-based discrimination between imagination of right and left hand movement. *Electroencephalography and clinical Neurophysiology*, 103(6):642–651, 1997.

[20] G. Pfurtscheller and C. Neuper. Motor imagery and direct brain-computer communication. *Proceedings of the IEEE*, 89(7):1123–1134, 2001.

[21] E Friedrich, R Scherer, J Faller, and C Neuper. Do user-related factors of motor impaired and able-bodied participants correlate with classification accuracy? In *Procceedings of the 5th International Brain-Computer Interface Conference*, pages 156–159. Graz University of Technology, 2011.

[22] J del R Millán, Josep Mouriño, Marco Franzé, Febo Cincotti, Markus Varsta, Jukka Heikkonen, and Fabio Babiloni. A local neural classifier for the recognition of EEG patterns associated to mental tasks. *IEEE Transactions on Neural Networks*, 13(3):678–686, 2002.

[23] Elisabeth VC Friedrich, Reinhold Scherer, and Christa Neuper. Long-term evaluation of a 4-class imagery-based brain–computer interface. *Clinical Neurophysiology*, 2013.

[24] Mark H Libenson. *Practical approach to electroencephalography*. Saunders, 2012.

[25] Joseph C Doyle, Robert Ornstein, and David Galin. Lateral specialization of cognitive mode: Ii. eeg frequency analysis. *Psychophysiology*, 11(5):567–578, 1974.

[26] Howard Ehrlichman and Marjorie S Wiener. Eeg asymmetry during covert mental activity. *Psychophysiology*, 17(3):228–235, 1980.

[27] AS Gevins, GM Zeitlin, JC Doyle, CD Yingling, RE Schaffer, E Callaway, and CL Yeager. Electroencephalogram correlates of higher cortical functions. *Science*, 203(4381):665–668, 1979.

[28] Mariko Osaka. Peak alpha frequency of eeg during a mental task: Task difficulty and hemispheric differences. *Psychophysiology*, 21(1):101–105, 1984.

[29] Petrie Rappelsberger and Hellmuth Petsche. Probability mapping: power and coherence analyses of cognitive processes. *Brain Topography*, 1(1):46–54, 1988.

[30] Gert Pfurtscheller. Event-related synchronization (ers): an electrophysiological correlate of cortical areas at rest. *Electroencephalography and clinical neurophysiology*, 83(1):62–69, 1992.

[31] C Neuper and G Pfurtscheller. Event-related dynamics of cortical rhythms: frequency-specific features and functional correlates. *International journal of psychophysiology*, 2001.

[32] Elisabeth VC Friedrich, Reinhold Scherer, and Christa Neuper. Stability of event-related (de-) synchronization during brain–computer interface-relevant mental tasks. *Clinical Neurophysiology*, 2012.

[33] G. Heinzel, A. Rüdiger, R. Schilling, and T. Hannover. Spectrum and spectral density estimation by the discrete fourier transform (dft), including a comprehensive list of window functions and some new flat-top windows. 2002.

[34] Trevor. Hastie, Robert. Tibshirani, and J Jerome H Friedman. *The elements of statistical learning*, volume 1. Springer New York, 2001.

[35] José del R Millán, Frédéric Renkens, Josep Mouriño, and Wulfram Gerstner. Brain-actuated interaction. *Artificial Intelligence*, 159(1):241–259, 2004.

[36] José Del R. Millán, Pierre W Ferrez, Ferran Galán, Eileen Lew, and Ricardo Chavarriaga. Non-invasive brain-machine interaction. *International Journal of Pattern Recognition and Artificial Intelligence*, 22(05):959–972, 2008.

[37] Ferran Galán, Marnix Nuttin, Eileen Lew, Pierre W Ferrez, Gerolf Vanacker, Johan Philips, and J del R Millán. A brain-actuated wheelchair: asynchronous and non-invasive brain–computer interfaces for continuous control of robots. *Clinical Neurophysiology*, 119(9):2159–2169, 2008.

[38] Jerome H Friedman. Regularized discriminant analysis. *Journal of the American statistical association*, 84(405):165–175, 1989.

[39] Paul S Addison. *The illustrated wavelet transform handbook: introductory theory and applications in science, engineering, medicine and finance*. Taylor & Francis, 2010.

[40] Li Zhiwei and Shen Minfen. Classification of mental task eeg signals using wavelet packet entropy and svm. In *8th International Conference on Electronic Measurement and Instruments*, pages 3–906. IEEE, 2007.

[41] Eduardo Iáñez, José María Azorín, Andrés Úbeda, José Manuel Ferrández, and Eduardo Fernández. Mental tasks-based brain–robot interface. *Robotics and Autonomous Systems*, 58(12):1238–1245, 2010.

[42] Jean-Philippe Lachaux, Eugenio Rodriguez, Jacques Martinerie, Francisco J Varela, et al. Measuring phase synchrony in brain signals. *Human brain mapping*, 8(4):194–208, 1999.

[43] Florian Mormann, Klaus Lehnertz, Peter David, and Christian E Elger. Mean phase coherence as a measure for phase synchronization and its application to the eeg of epilepsy patients. *Physica D: Nonlinear Phenomena*, 144(3):358–369, 2000.

[44] Elly Gysels and Patrick Celka. Phase synchronization for the recognition of mental tasks in a brain-computer interface. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 12(4):406–415, 2004.

[45] Dean J Krusienski, Dennis J McFarland, and Jonathan R Wolpaw. Value of amplitude, phase, and coherence features for a sensorimotor rhythm-based brain–computer interface. *Brain research bulletin*, 87(1):130–134, 2012.

[46] Norman Richard Draper, Harry Smith, and Elizabeth Pownell. *Applied regression analysis*, volume 3. Wiley New York, 1966.

[47] Benjamin Hamner, Robert Leeb, Michele Tavella, and José del R Millán. Phase-based features for motor imagery brain-computer interfaces. In *2011 Annual International Conference of The IEEE Engineering in Medicine and Biology Society*, pages 2578–2581. IEEE, 2011.

[48] Charles Anderson, Elliott Forney, Douglas Hains, and Annamalai Natarajan. Reliable identification of mental tasks using time-embedded eeg and sequential evidence accumulation. *Journal of Neural Engineering*, 8(2):025023, 2011.

[49] Carl W. Baum and Venugopal V. Veeravalli. A sequential procedure for multihypothesis testing. *IEEE Transactions on Information Theory*, 40(6), 1994.

[50] Charles W Anderson, James N Knight, Tim O'Connor, Michael J Kirby, and Artem Sokolov. Geometric subspace methods and time-delay embedding for eeg artifact removal and classification. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 14(2):142–146, 2006.

[51] Charles W Anderson, James N Knight, Michael J Kirby, and Douglas R Hundley. Classification of time-embedded eeg using short-time principal component analysis. *Toward Brain-Computer Interfacing*, pages 261–278, 2007.

[52] Charles W Anderson and Jeshua A Bratman. Translating thoughts into actions by finding patterns in brainwaves. In *Proceedings of the Fourteenth Yale Workshop on Adaptive and Learning Systems*, pages 1–6, 2008.

[53] Li Zhang, Wei He, Chuanhong He, and Ping Wang. Improving mental task classification by adding high frequency band information. *Journal of medical systems*, 34(1):51–60, 2010.

[54] Elisabeth VC Friedrich, Reinhold Scherer, and Christa Neuper. The effect of distinct mental strategies on classification performance for brain–computer interfaces. *International Journal of Psychophysiology*, 84(1):86–94, 2012.

[55] Reinhold Scherer, Josef Faller, David Balderas, Elisabeth VC Friedrich, Markus Pröll, Brendan Allison, and Gernot Müller-Putz. Brain–computer interfacing: more than the sum of its parts. *Soft Computing*, 17(2):317–331, 2013.

[56] Johannes Müller-Gerking, Gert Pfurtscheller, and Henrik Flyvbjerg. Designing optimal spatial filters for single-trial eeg classification in a movement task. *Clinical neurophysiology*, 110(5):787–798, 1999.

[57] Herbert Ramoser, Johannes Muller-Gerking, and Gert Pfurtscheller. Optimal spatial filtering of single trial eeg during imagined hand movement. *IEEE Transactions on Rehabilitation Engineering*, 8(4):441–446, 2000.

[58] Peter J Brockwell and Richard A Davis. *Introduction to time series and forecasting*. Springer Verlag, 2002.

[59] Charles W Anderson, Erik A Stolz, and Sanyogita Shamsunder. Discriminating mental tasks using eeg represented by ar models. In *17th Annual Conference of The IEEE Engineering in Medicine and Biology Society*, volume 2, pages 875–876. IEEE, 1995.

[60] Charles W Anderson, Erik A Stolz, and Sanyogita Shamsunder. Multivariate autoregressive models for classification of spontaneous electroencephalographic signals during mental tasks. *IEEE Transactions on Biomedical Engineering*, 45(3):277–286, 1998.

[61] Eleanor Curran, Peter Sykacek, Maria Stokes, Stephen J Roberts, Will Penny, Ingrid Johnsrude, and Adrian M Owen. Cognitive tasks for driving a brain-computer interfacing system: a pilot study. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 12(1):48–54, 2004.

[62] K-R Muller, Charles W Anderson, and Gary E Birch. Linear and nonlinear methods for brain-computer interfaces. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 11(2):165–169, 2003.

[63] Simon S Haykin. *Neural networks and learning machines*, volume 3. Prentice Hall New York, 2009.

[64] Lalit Gupta, Mark McAvoy, and James Phegley. Classification of temporal sequences via prediction using the simple recurrent neural network. *Pattern Recognition*, 33(10):1759–1770, 2000.

[65] Ikusaburo Kurimoto Shinichi Oeda and Takumi Ichimura. Time series data classification using recurrent neural network with ensemble learning. *Lecture Notes in Computer Science*, 4253:742–748, 2006.

[66] Damien Coyle, Girijesh Prasad, and Thomas Martin McGinnity. A time-series prediction approach for feature extraction in a brain-computer interface. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 13(4):461–467, 2005.

[67] Elliott M Forney and Charles W Anderson. Classification of eeg during imagined mental tasks by forecasting with elman recurrent neural networks. In *Proceedings of The 2011 International Joint Conference on Neural Networks*, pages 2749–2755. IEEE, 2011.

[68] Damien Coyle, Thomas M McGinnity, and Girijesh Prasad. Creating a nonparametric brain-computer interface with neural time-series prediction preprocessing. In *28th Annual International Conference of The IEEE Engineering in Medicine and Biology Society*, pages 2183–2186. IEEE, 2006.

[69] Damien Coyle, Girijesh Prasad, and Thomas M McGinnity. Extracting features for a brain-computer interface by self-organising fuzzy neural network-based time series prediction. In *26th Annual International Conference of The IEEE Engineering in Medicine and Biology Society*, volume 2, pages 4371–4374. IEEE, 2004.

[70] Damien Coyle. Neural network based auto association and time-series prediction for biosignal processing in brain-computer interfaces. *Computational Intelligence Magazine*, 4(4):47–59, 2009.

[71] Damien Coyle, Thomas M McGinnity, and Girijesh Prasad. A multi-class brain-computer interface with sofnn-based prediction preprocessing. In *Proceedings of The 2008 IEEE International Joint Conference on Neural Networks*, pages 3696–3703. IEEE, 2008.

[72] Elliott M Forney. Electroencephalogram classification by forecasting with recurrent neural networks. Master's thesis, Colorado State University, 2011.

[73] Elliott Forney, Charles Anderson, William Gavin, and Patricia Davies. A stimulus-free brain-computer interface using mental tasks and echo state networks. In *Proceedings of the Fifth International Brain-Computer Interfaces Meeting (To appear.)*. Graz University of Technology Publishing House, 2013.

[74] Julien Kronegg, Svyatoslav Voloshynovskiy, and Thierry Pun. Analysis of bit-rate definitions for brain-computer interfaces. In *International Conference on Human-computer Interaction*, volume 12, page 18, 2005.