Chapter 1

Introduction

The continuous improvement of acquisition and computer simulation techniques results in increasing amounts of numerical data. Yet, the aim is not simply to acquire data, but to gain insight in the measured or computed data. Visualization aids the interpretation by transforming large quantities of (complex) data into visual representations. Medical imaging belongs to the better-known application fields of visualization, with neuroimaging as one of its more prominent disciplines. Neuroimaging comprises structural imaging (brain anatomy) and functional imaging (brain activity).

Developed in the 1920s, *electroencephalography (EEG)* is the oldest noninvasive functional neuroimaging technique, which records electrical brain activity. Other noninvasive functional

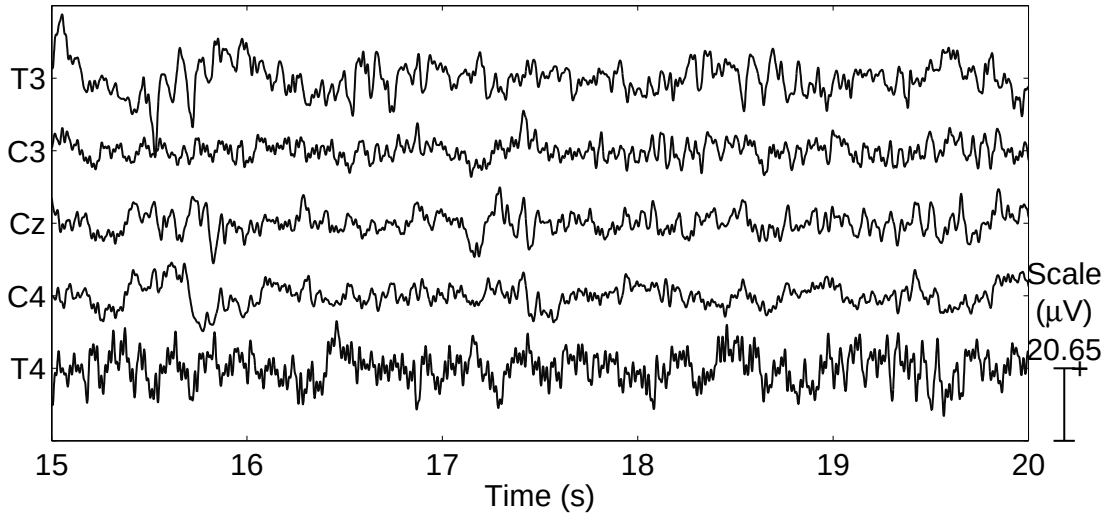

Figure 1.1. *Conventional EEG representation. Measured voltages (in μV) are shown against time (in seconds), for five electrodes (indicated by the labels T3, C3, Cz, C4, and T4).*

neuroimaging techniques are magnetoencephalography (MEG), which measures magnetic fields induced by electrical brain activity (Cohen 1968, Hämäläinen *et al.* 1993), and functional magnetic resonance imaging (fMRI), which measures changes in cerebral blood oxygenation levels in the brain (Ogawa *et al.* 1990, Logothetis *et al.* 2001). When comparing the three noninvasive techniques, the temporal resolution is superior for EEG and MEG, whereas the spatial resolution

is superior for fMRI. However, EEG has much lower costs and has a higher equipment transportability than MEG and fMRI. Moreover, EEG allows participants more freedom to move than MEG and fMRI.

Nowadays, EEG has various clinical as well as scientific applications in different fields, including medicine, pharmacy, psychology, linguistics, and biology. EEG is one of those fields dealing with increasing amounts of data as a result of technological advances. One improvement is the increased number of electrodes that can be attached to the scalp, leading to *multichannel* (or *high-density*, or *high-resolution*) EEG. For example, a typical setting for multichannel EEG involves 128 electrodes, while simultaneously measuring 1000 potentials per second, yielding 7.68 million data elements per minute, or 460.8 million data elements per hour. Further, automated signal processing by computer may provide additional data.

However, data is not yet information. Therefore, several different visualization methods are applied to assist the interpretation of the EEG. In a conventional EEG visualization, the time-varying EEG is represented by one time series per electrode, displaying the measured potential as a function of time (Fig. 1.1).

Synchronous activity between brain regions is associated with a functional relationship between those regions. EEG coherence, calculated between pairs of electrode signals as a function of frequency, is a measure for this synchrony. A common visualization of EEG coherence is a graph layout. In the case of EEG, graph vertices (drawn as dots) represent electrodes and graph edges (drawn as lines between dots) represent similarities between pairs of electrode signals (Fig. 1.2). Traditional visual representations are, however, not tailored for the current number of electrodes.

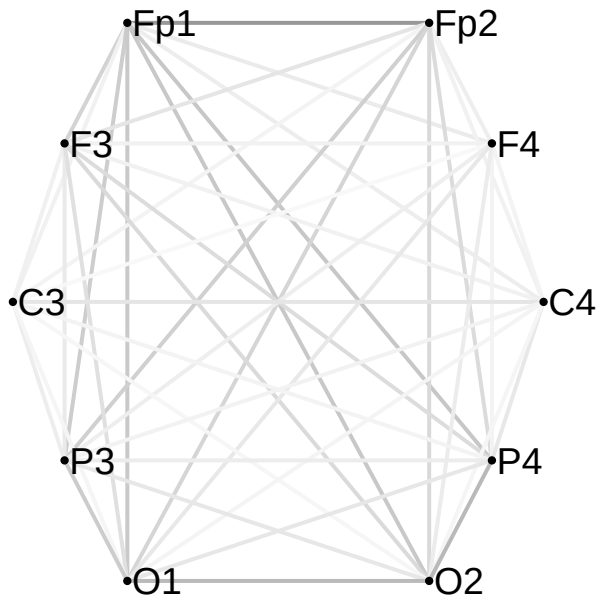

Figure 1.2. *Conventional graph visualization for EEG coherence. Electrodes (labeled Fp1, Fp2, F3, F4, C3, C4, P3, P4, O1, O2) are represented by dots, similarities between pairs of electrode signals (above a certain threshold) by lines. Darker lines indicate higher similarities.*

This thesis introduces visualizations for time domain analysis of multichannel EEG data and frequency domain analysis restricted to EEG coherence, which are specifically designed to overcome shortcomings of existing techniques. Evaluations of these new methods include a qualitative study and a user study for the time domain visualization, and two case studies for the frequency domain visualization.

The remainder of this chapter is organized as follows. Section 1.1 provides background information on EEG and distinguishes time and frequency domain analysis of EEG data. Basic visualization techniques for time-varying multivariate data in the time domain and relational (coherence) data in the frequency domain are included in Section 1.2. Section 1.3 describes the contributions of this thesis, and gives an outline of the remainder of this thesis.

1.1 Electroencephalography (EEG)

Electrical potentials generated within the brain can be measured with electrodes at the scalp during an EEG recording. The measured EEG signals reflect rhythmical activity varying with brain state. Specific brain responses can be elicited by the presentation of external stimuli. For EEG analysis, the time domain, frequency domain, and time-frequency domain are distinguished.

1.1.1 Brain Potentials

The outer part of the human brain which evolved most recently is called the *neocortex*. It is often briefly referred to as *cortex*. The neocortex has an estimated number of $2 \cdot 10^{10}$ neurons (nerve cells), a number decreasing with age and stabilizing in adults (Pakkenberg *et al.* 2003). These neurons transfer signals between them. The point of contact between two neurons is a so-called synapse. The number of synapses in the brain is about $1.5 \cdot 10^{14}$ (Pakkenberg *et al.* 2003).

Neurons exchange ions (electrically charged particles) with their environment, thus changing their potential. The *resting potential* of a neuron (approximately $-65\,\mathrm{mV}$) is the consequence of an uneven distribution of ions on both sides of the membrane (the outer layer of a neuron). A reduction of the resting potential (depolarization) by at least $10\,\mathrm{mV}$ causes an *action potential*, with an amplitude of about $110\,\mathrm{mV}$ and a duration of about $1\,\mathrm{ms}$. A cluster of action potentials results in so-called *burst firing* and tends to be synchronous among multiple neurons in the same population. Action potentials in presynaptic cells (sending neurons) can cause a *postsynaptic potential* change in postsynaptic cells (receiving neurons) with an amplitude of a few microvolts and a duration of over $100\,\mathrm{ms}$. An inhibitory (excitatory) postsynaptic potential caused at an inhibitory (excitatory) synapse increases (decreases) the membrane potential.

A pyramidal neuron (or pyramidal cell) is a so-called multipolar neuron, having over 100,000 synapses on the receiving end. Pyramidal neurons are the most important cells for potential changes in the cortex. Functional vertical columns in the cerebral cortex consist of hundreds of pyramidal cells located closely to each other and oriented in parallel. A group of several thousands of (postsynaptic) pyramidal cells can be activated simultaneously if they are connected to the same (presynaptic) neuron.

As a result of *volume conduction*, an electrical current flows from the generator in the brain through different tissues (e.g., brain, skull, skin) to a recording electrode on the scalp. The measured EEG is mainly generated by neuronal (inhibitory and excitatory) postsynaptic potentials and burst firing in the cerebral cortex. Isolated action potentials have a much smaller or even negligible effect on the EEG. Variations of the (summed) postsynaptic potentials in time and location cause the temporal and spatial variations of the EEG. Measured potentials depend on the source intensity, its distance from the electrodes, and on the conductive properties of the tissues between the source and the recording electrode (Spehlmann and Fisch 1991).

1.1.2 Anatomical References

Here, some background information is provided on anatomical references for a human head (Beatty 1995, Bear *et al.* 1996). The most prominent features of the human brain are its two (cerebral) *hemispheres*, with the *midline* running through the middle. Structures towards the midline are *medial*, structures away from the midline are *lateral* (Fig. 1.3). *Contralateral* refers to something on the other side, *ipsilateral* to something on the same side. *Homologous* refers to a similarity in structure or position.

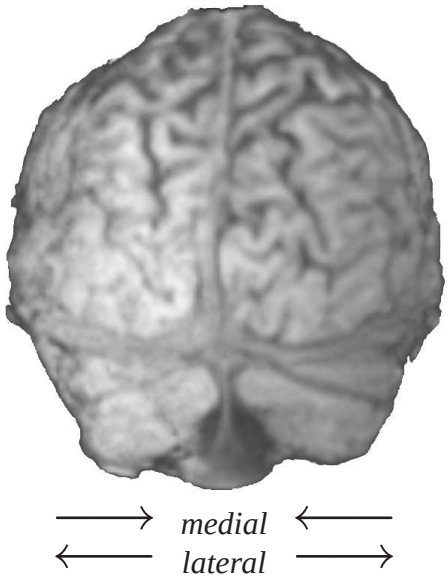

Figure 1.3. *Neocortex, backside of the head. Medial (towards the middle) and lateral (towards the sides). Image created with MRIcro (Rorden and Brett 2000).*

When standing or sitting, directing towards the top of the head is called *superior* (or dorsal). The opposite is *inferior* (or ventral). Towards our nose is referred to as *anterior* (or rostral) and towards the back of the head as *posterior* (or caudal).

The convolutions of the surface of the neocortex consist of series of hills and valleys, referred to as *gyri* and *sulci*, respectively. Some deeper sulci are referred to as *fissures*. The *interhemispheric* (or (medial) longitudinal) fissure separates the left and right hemispheres. Other fissures

divide both hemispheres into four regions referred to as *lobes*. The *frontal* lobe is the most anterior lobe. The frontal and parietal lobes are separated by the *central* (or Rolandic) fissure. The *temporal* lobe is inferior to the frontal and parietal lobes, and is separated from those two lobes by the *lateral* (or Silvian) fissure. The most posterior lobe is the *occipital* lobe. The parietal and occipital lobes are separated by the *parieto-occipital* fissure. The prefixes *pre* and *post* mean directly anterior and posterior, respectively. Anatomical references and lobes are illustrated in Fig. 1.4. Directly posterior to this is the *parietal* lobe.

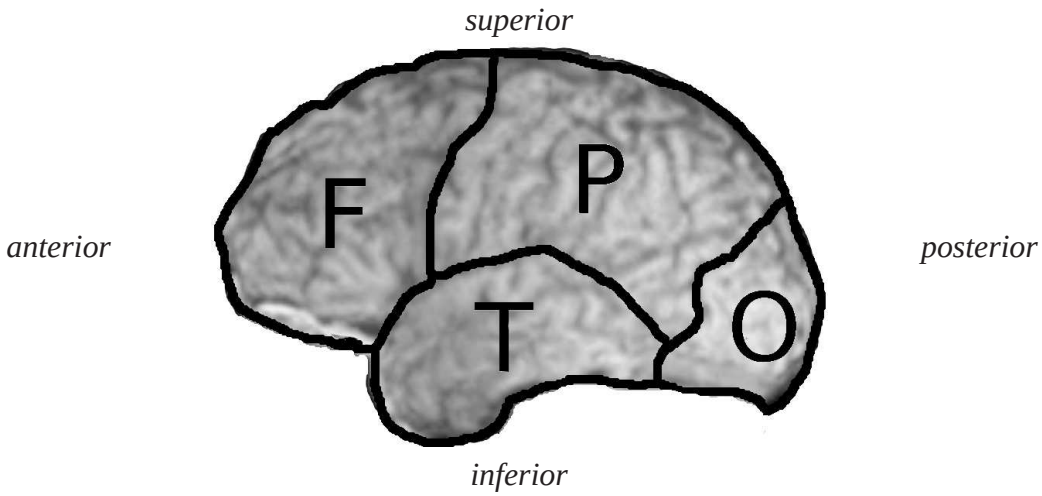

Figure 1.4. *Frontal (F), parietal (P), temporal (T), and occipital (O) lobe (neocortex, lateral view of the left hemisphere, with the nose on the left). Additionally, anatomical references (anterior, posterior, inferior, superior) are indicated. Image created with MRIcro (Rorden and Brett 2000).*

1.1.3 EEG Recording

During an EEG recording, up to 512 electrodes can be attached to the scalp at different positions. A conductive gel is applied between skin and electrodes to reduce impedance. An electrical potential is measured from all electrodes simultaneously, typically between 125 and 2000 times per second. Clinical recordings may take 15 to 30 minutes, scientific (e.g., sleep) recordings may take hours. Electrodes are often held in fixed positions by an elastic cap. Each electrode carries a unique labeling, usually by a combination of letters and numbers. The letters refer to a nearby lobe or fissure. Odd (even) numbers refer to the left (right) hemisphere. A label with a 'z' instead of a number refers to a midline electrode. A label 'Fp' refers to a prefrontal electrode. Letters may refer to the central fissure (C), or to the frontal (F), parietal (P), temporal (T), or occipital (O) lobe. Labels of other electrodes use a letter combination of nearby electrodes. Electrodes located between the earlier labeled electrodes are labeled frontocentral (FC), centroparietal (CP), parieto-occipital (PO), frontotemporal (FT), and temporoparietal (TP). Electrodes anterior to the frontal electrodes are labeled anterofrontal (AF). Additional electrodes between FC and C are referred to as FCC, et cetera. For an example with 118 labeled electrodes, see Fig. 1.5.

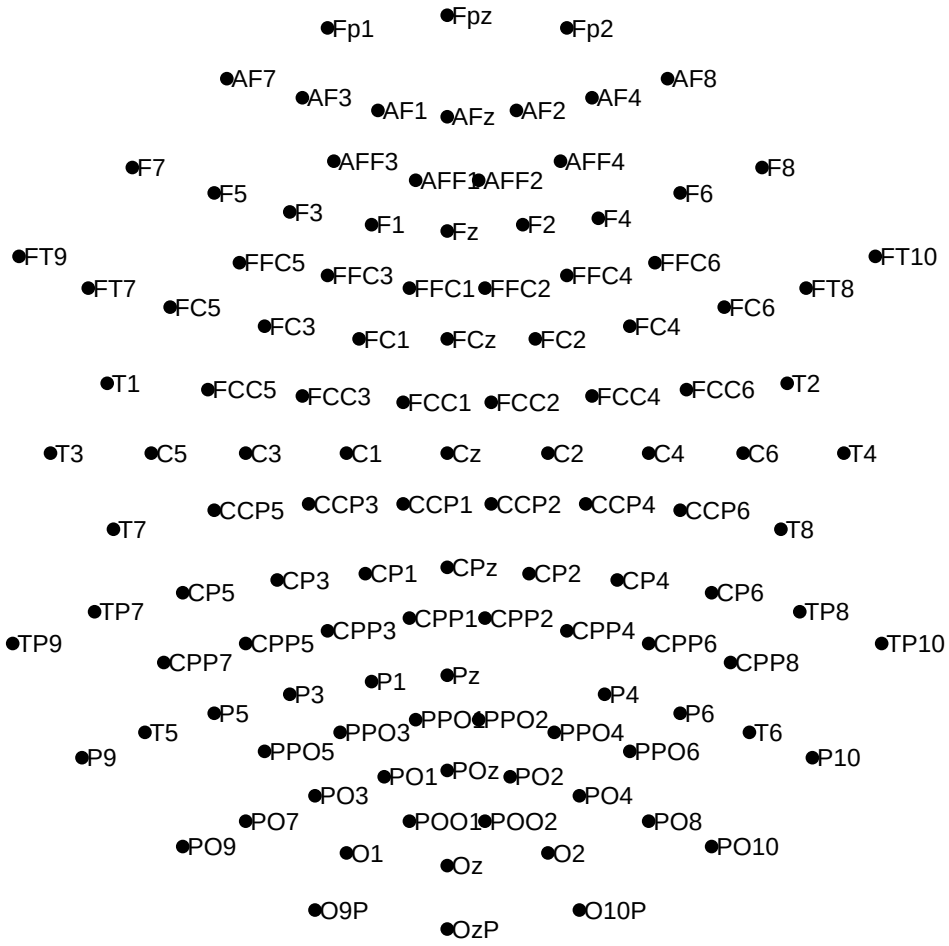

Figure 1.5. *Positions (dots) and labels of 118 EEG electrodes, top view of the head (nose on top).*

The measured signal from each electrode is amplified, resulting in one recording channel for every electrode. For every electrode, the potential is expressed as the difference between the measured signal and a reference signal (computed per time step). For multichannel EEG with a large number of widely spread electrodes, an *average reference* is often applied which is the average of all measured scalp potentials. Alternatively, a *common reference* uses the signal of one selected electrode as a reference for all other electrodes, or it may combine two signals, e.g., recorded from electrodes attached to both ears. Potential differences between scalp electrodes are usually of the order of 10 μV to 100 μV.

Usually, a *high-pass* (*low-pass*) filter is applied to reduce frequencies below (above) a certain threshold outside the range of frequencies of interest. A typical value for a high-pass filter is 0.16 Hz, for a low-pass filter 70 Hz. An additional *notch* filter can be used to reduce effects of the electric mains (50 Hz in the Netherlands).

1.1.4 Rhythms

The complex summation of neuronal activity measured by EEG results in rhythmical activity. Different frequency bands are generally recognized (delta, theta, alpha, beta, gamma), although there is no clear consensus on the boundaries between these bands. Delta may refer to frequencies up to 4 Hz (Nunez and Silberstein 2000), theta between 4 and 8 Hz, alpha between 8 and 12 Hz, beta between 12 and 24 Hz, and gamma above 24 Hz (Stein *et al.* 1999). The main frequency in the EEG usually depends on the state of the subject. For example, (medium amplitude) theta activity in the human EEG indicates drowsiness, whereas (high-amplitude) delta activity is observed during non-dreaming sleep (Beatty 1995).

1.1.5 Event-Related and Evoked Potentials

Often, brain potentials are recorded following (e.g., visual, auditory, or somatosensory) stimulus presentation. Generally, if two stimuli are identical, then two nearly identical brain responses are expected. However, from trial to trial, responses will differ due to the varying background EEG (i.e., ongoing EEG which is not the result of the stimulus). Because the background EEG is expected to be on average equal to zero, consecutive responses are usually averaged (ensemble-averaging). When low-level brain responses are the objective of the study, the response is called evoked potential (EP). When higher-level cognitive responses are the objective, the response is called event-related potential (ERP). In comparison with EPs, ERPs generally exhibit longer latencies, higher amplitudes, and lower frequencies. EPs and ERPs can be evaluated in the time domain or the frequency domain. The identification of peaks can be supported by a plot showing the measured potential against time, with one time axis per electrode or with one time axis for all electrodes (butterfly plot). Positive and negative peaks, referred to as *components*, are identified with their amplitudes and latencies. For illustrative purposes, we explain one EP and one ERP in more detail. Both are applied in clinical practice and are used later in this thesis to demonstrate new visualizations.

A *somatosensory evoked potential* (*SEP*) is the result of external stimulation of the peripheral nervous system (the nervous system other than the brain and the spinal cord). This stimulation is usually electrical, sometimes mechanical (Aminoff and Eisen 1998). We focus here on electrical stimulation of the median nerve (near the wrist). A SEP has a characteristic contralateral negative peak about 20 ms after the stimulus in the EEG signal recorded near the parietal cortex, referred to as N20 component, or N20 for short (Fig. 1.6). The N20 is thought to reflect the arrival at the cortex of a volley of signals caused by the electrical stimulus (Tomberg and Desmedt 1999). A butterfly plot may additionally show changes in the overall variation of the EEG (Fig. 1.7).

An often used ERP in clinical practice is the *P300* (Fonaryova Key *et al.* 2005). It is usually elicited by a so-called *oddball* paradigm, presenting two types of stimuli with different probabilities in a random order, during an attention demanding task. The participant is required to respond to the infrequent stimulus, usually by mentally counting or a button press. For example, during an auditory oddball task, stimuli may consist of a majority of low tones (85 %), randomly alternated with high tones (15 %) as the infrequent stimulus (Maurits *et al.* 2006). The name 'P300' refers to the characteristic *p*ositive peak measured about 300 ms after the stimulus. The

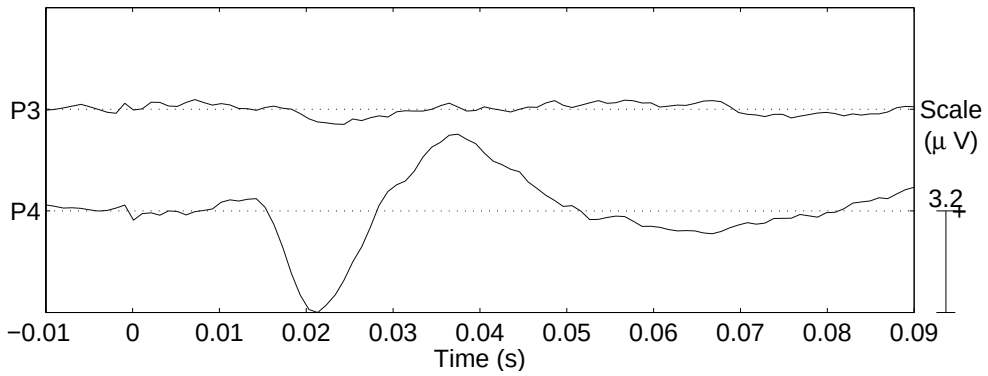

Figure 1.6. *Somatosensory evoked potential (SEP) , left-hand stimulation of the median nerve (near the wrist). The average EEG over approximately 500 electrical stimuli is displayed for the electrodes labeled P3 and P4 (attached near the left and the right parietal cortex, respectively). The dotted line indicates the zero-level. The first negative peak for electrode P4, just after 20 ms, is referred to as the N20.*

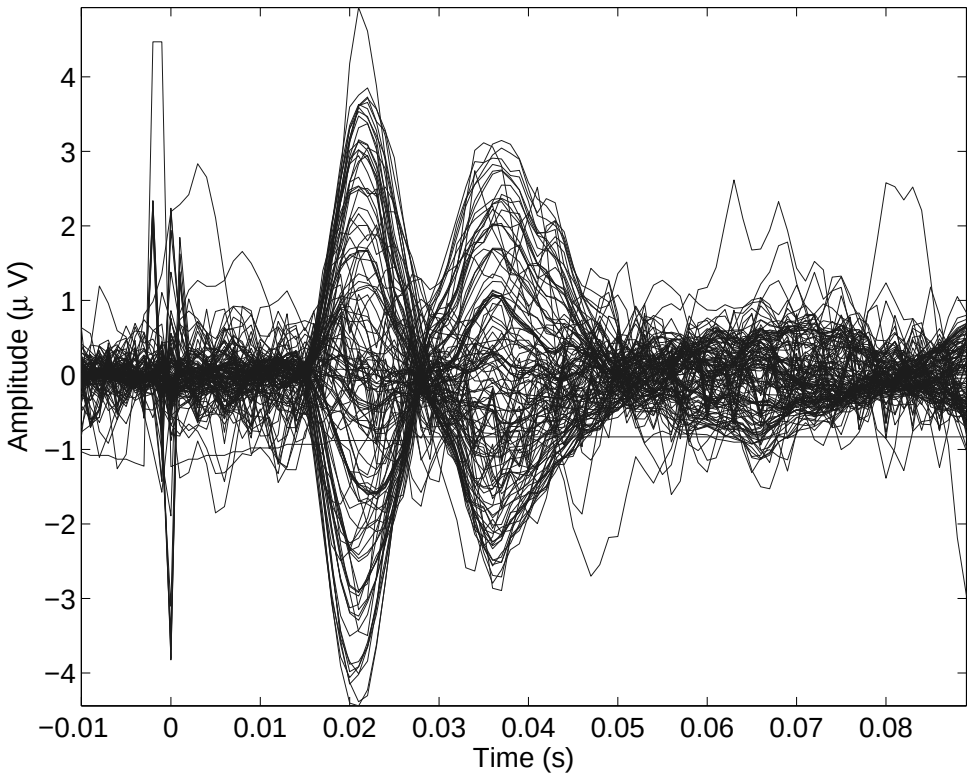

Figure 1.7. *Butterfly plot, with superimposed SEPs for 128 electrodes. The EEG shows local maxima in the variation about 0.021 s and 0.035 s after the stimulus. (The pattern between approximately 0.015 and 0.060 s resembles a butterfly shape.)*

P300 potential is also referred to as P3, i.e., the third positive peak after the stimulus. Nowadays, it is known that the P3 peak consists of an anterior (P3a) component and a posterior (P3b)

component (Comerchero and Polich 1999).

1.1.6 EEG Analysis

For the analysis of EEG signals, we first briefly consider the time domain. Next, we indicate how frequency domain and time-frequency domain analysis are employed to study rhythmical activity.

Time Domain

Time domain analysis of EEG mainly involves the identification of positive and negative peaks of interest and their amplitudes and latencies. Human experts (visual inspection) or automatic methods can detect, e.g., specific (sequences of) peaks associated with epileptiform activity, referred to as spikes and spike-and-wave complexes (Wilson and Emerson 2002, Jerger *et al.* 2001).

Alternatively, the relationship between such simultaneously recorded signals can be analyzed by computer. Cross-correlation is a classical measure for the interdependence between two time series as a function of their delay time. The linear *cross-correlation* for two time series $x(t)$ and $y(t)$ as a function of time lag τ is (Pereda *et al.* 2005)

$$C_{xy}(\tau) = \frac{\sum_{k=1}^{N-\tau}(x_{k+\tau}-\mu_x)(y_k-\mu_y)}{\sqrt{\sum_{k=1}^{N-\tau}(x_{k+\tau}-\mu_x)^2}\sqrt{\sum_{k=1}^{N-\tau}(y_k-\mu_y)^2}}, \tag{1.1}$$

where x_t and y_t denote discrete sample values at time t with N the total number of discrete samples, and μ_x and μ_y denote the average of both respective signals. Values of C_{xy} are in the range $[-1,1]$. A value $C_{xy}(\tau) = -1$ suggests a complete linear inverse correlation between signals $x(t)$ and $y(t)$ for time lag τ, $C_{xy}(\tau) = 1$ a complete linear direct correlation, and $C_{xy}(\tau) = 0$ a lack of linear interdependence. A negative sign of C_{xy} indicates that both signals tend to have similar absolute values but opposite signs. A positive sign of C_{xy} indicates that both signals tend to have similar absolute values and the same sign. Although the value of τ which maximizes C_{xy} is often used as an estimation of the delay between two signals x and y, it cannot be directly interpreted as the propagation time of, e.g., an electrical signal in the cortex (Pereda *et al.* 2005).

Frequency Domain and Time-Frequency Domain

The *minimum frequency* that can be distinguished in a discrete time-varying signal depends on the segment length T (in seconds) and is equal to $\frac{1}{T}$, which is also the *frequency resolution* (the minimum distance between two frequencies that can be distinguished). The *maximum frequency* that can be distinguished in the same signal depends on the sampling rate φ (in Hz) and is equal to $\frac{\varphi}{2}$, according to the Nyquist-Shannon sampling theorem. The *sampling interval* is $\frac{1}{\varphi}$ s.

Frequency components of signals can be quantified in different ways. Spectral analysis can decompose a *single signal* into its frequency components, e.g., using the Fourier transform. Let $x(t)$ be a zero-mean time series and x_t again its discrete sample value at time t. A discrete *Fourier*

transform of the l^{th} segment ($l = 1 \ldots L$) of length T from time series x at frequency λ_j is defined as

$$\hat{d}_x^T(\lambda_j, l) = \sum_{t=(l-1)T}^{lT-1} e^{-i\lambda_j t} x_t, \tag{1.2}$$

where λ_j are the Fourier frequencies ($\lambda_j = 2\pi j/T$ for $j = 0 \ldots (T-1)/2$).

Alternatively, a similarity measure can be formulated for *signal pairs* as a function of frequency. A common measure for this synchrony is coherence. Let $y(t)$ be another zero-mean time series, with y_t denoting its discrete sample value at time t. The discrete *cross-spectrum* $\hat{f}_{xy}$ between discrete time series x and y at frequency λ_j is defined, suppressing the dependency on l, as (Halliday *et al.* 1995)

$$\hat{f}_{xy}(\lambda_j) = \frac{1}{2\pi LT} \sum_{l=1}^{L} \hat{d}_x^T(\lambda_j, l) \overline{\hat{d}_y^T(\lambda_j, l)}, \tag{1.3}$$

where the overbar "$^-$" indicates a complex conjugate. The discrete *auto-spectrum* is defined similarly and is denoted by $\hat{f}_{xx}(\lambda_j)$. The *coherence* $\hat{c}_{\lambda_j}$ as a function of frequency λ for two discrete time signals x and y is defined as the absolute square of the discrete cross-spectrum $\hat{f}_{xy}$ normalized by the discrete autospectra $\hat{f}_{xx}$ and $\hat{f}_{yy}$ (Halliday *et al.* 1995):

$$\hat{c}_{\lambda_j}(x,y) = \frac{\mid \hat{f}_{xy}(\lambda_j) \mid^2}{\hat{f}_{xx}(\lambda_j)\, \hat{f}_{yy}(\lambda)}. \tag{1.4}$$

Coherence has values in the interval $[0,1]$, with $\hat{c}_{\lambda_j}(x,y) = 0$ indicating that both time series are independent at frequency λ_j, whereas $\hat{c}_{\lambda_j}(x,y) = 1$ indicates that both time series are maximally linearly correlated at this frequency.

The cross-spectrum for large T (and $\lambda_j \neq 0$) has a form similar to the complex covariance parameter ($\text{cov}\{A,B\} = E\{(A - E\{A\})\overline{(B - E\{B\})}\}$ which reduces to $\text{cov}\{A,B\} = E\{A\overline{B}\}$ for zero-mean processes A and B) and can be interpreted as the covariance at frequency λ_j. Because correlation is defined as $\text{corr}\{A,B\} = \frac{\text{cov}\{A,B\}}{\text{var}\{A\}\text{var}\{B\}}$, the coherence $\hat{c}_{\lambda_j}$ can be regarded as a correlation of two signals at a frequency λ_j.

Synchronous electrical activity in different brain regions is generally assumed to imply functional connectivity between these regions. Therefore, EEG coherence, calculated between signal pairs recorded from electrodes attached at different positions, is employed for the study of functional brain connectivity. EEG coherence can be regarded as the correlation of two electrode signals at a certain frequency.

The frequency content of an EEG signal changes over time. To follow these changes, one segment can be divided into multiple *windows* which commonly have the same width and overlap. The center of a window associated with a time step is systematically adapted in a *moving window* approach. When the frequency content of consecutive (moving) windows is quantified, this approach is referred to as time-frequency analysis. Examples of time-frequency analysis include the short-time Fourier transform (Zhan *et al.* 2006) and the wavelet transform (Adeli *et*

al. 2003, Bartnik *et al.* 1992, Glassman 2005, Zhan *et al.* 2006). There is always a trade-off between a high time and a high frequency resolution, due to the effects of segment length on the frequency resolution and the minimum frequency.

1.2 Visualization

Within the field of visualization, three (partially overlapping) areas are distinguished. *Scientific visualization* involves data representing physical phenomena. *Information visualization* represents more abstract data. *Visual analytics* uses interactive visual interfaces to support analytical reasoning. Alternatively, visualizations can be classified by their application (e.g., medicine, geography and cartography, fluid flow, astronomy), data type (e.g., scalar, vector, tensor, multivariate, time-varying, relational), or interaction technique (e.g., zooming, linked views).

Here, we discuss visualizations related to the data types for time domain and frequency domain EEG. Time domain multichannel EEG concerns one time series per electrode and can be considered as *time-varying multivariate data* (Section 1.2.1). Correlation in the time domain is not considered in this thesis. Frequency domain analysis is here restricted to the study of functional brain connectivity using EEG coherence. EEG coherence analysis results in *relational data* for pairs of electrode signals, typically represented by *graphs* (Section 1.2.2).

The discussion is here restricted to some basic representations for the data types under concern, not intending to give a complete overview of the visualization of time-varying, multivariate, or relational data. More specific visualizations are discussed when relevant in the following chapters.

1.2.1 Multivariate and Time-Varying Data

The visualization of time-varying and multivariate data is challenging. Sometimes, time is considered to be just another dimension.

Time series

The oldest known visualization of time-varying data is a time series plot dating from the tenth or eleventh century, showing multiple inclinations of planetary orbits as a function of time (Tufte 1983). This visualization implies a natural ordering for the time scale, because if a time step t is visualized between two other time steps t_a and t_b (with $t_a < t_b$), it means that $t_a < t < t_b$. Nowadays, the time series plot is the most popular of all graphic representations. Examples are the conventional EEG and the averaged EEG, each with one horizontal time axis per electrode (Figs. 1.1 and 1.6), and the butterfly plot with multiple electrode signals in combination with one time axis (Fig. 1.7).

Other options for time-varying data are animations instead of static visualizations (Ware 2004). Further, cyclic time has other characteristics than linear time and offers other possibilities, such as a circular layout (Müller and Schumann 2003).

Pixel-Oriented

One approach to represent large quantities of any type of data is to map each data value to the color of the smallest element of a screen, a pixel. This approach is referred to as pixel-oriented (or screen-filling, or dense pixel) technique. Different pixel orderings (row by row, column by column, or recursive patterns) can be applied (Keim 2000). One drawback is that the number of pixels cannot easily be extended. Moreover, relations between data elements might be lost by their ordering on the screen. For example, on a screen with 1500×2000 pixels, there is no suitable ordering for elements representing 3 million consecutive time steps, such that two elements are displayed close to each other *if and only if* the two elements are close to each other in time.

Scatter Plot, Projection Pursuit, and Grand Tour

A *scatter plot* (Fig. 1.9) shows relations between two dimensions of a dataset, showing for each data vector one symbol (e.g., dot, cross) at the position corresponding to the values in the two selected dimensions. Whereas scatter plots are restricted to two-dimensional projections, projection pursuit and grand tour similarly show series of static lower-dimensional projections and animations of lower-dimensional projections, respectively. *Projection pursuit* selects those static projections which are interesting according to a criterion based on the choice of the user (Huber 1985), while ignoring other projections. *Grand tour* (Asimov 1985) is supposed to represent all possible projections, however, it may miss the most interesting projections. Moreover, for a large number of dimensions the number of possibly interesting projections may be large.

Tabular Layout

A contingency table sets out features along rows and columns and shows relations between row and column features. Visualizations of such tables often use a tabular (or matrix) layout. Each table (or matrix) entry visualizes the relation between the corresponding row and column features. For example, an EP (or ERP) image for EEG sets out EPs (or ERPs) along rows and time steps along columns (Jung *et al.* 2001). A procedure to plot EP images is available in EEGLAB (Delorme and Makeig 2004). Each table entry is a square with an electrical potential value mapped to a gray value (Fig. 1.8). This technique can be applied to any table with scalar values, provided that the number of rows and columns does not exceed the vertical or horizontal screen resolution. A similar example from bioinformatics is referred to as Eisen plot, with genes along rows and time steps along columns, and gene expression mapped to a color (Eisen *et al.* 1998).

Other tabular layouts visualize more than one data value per table entry. In a so-called *scatter plot matrix*, the dimensions for an N-dimensional dataset are set out along both rows and columns, with one scatter plot for each pair of dimensions (Fig. 1.9). It is an appropriate representation, given the human ability to observe patterns for the identification of interesting relationships between pairs of dimensions. However, there may be a higher-dimensional relationship which is not directly visible.

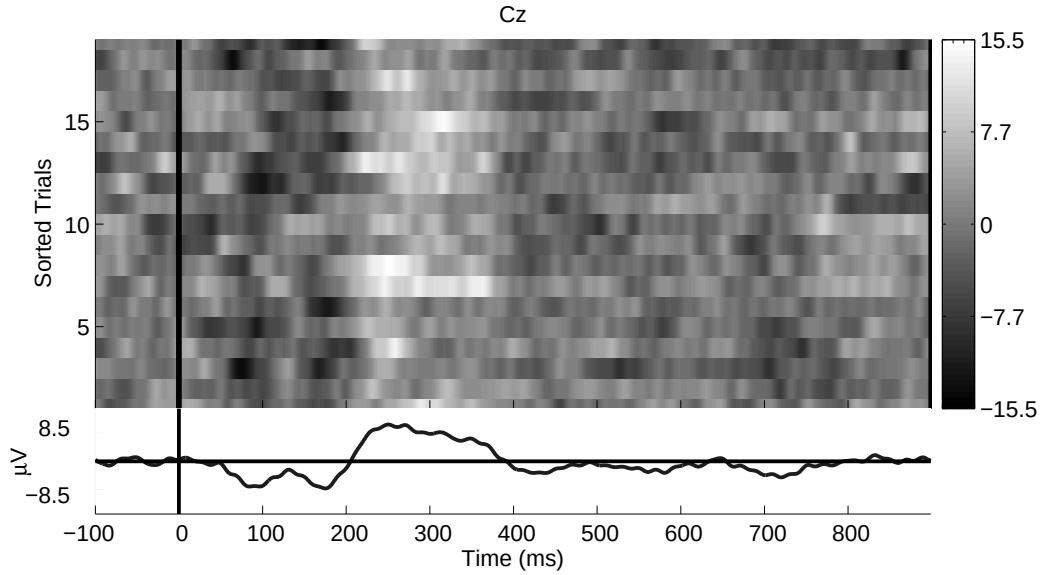

Figure 1.8. *ERP image for the electrode labeled Cz. Sixteen P300 responses (rows) are color-coded separately. At the bottom, a time series of the (average) ERP is shown.*

Whereas the previously mentioned tabular layouts use regularly spaced table entries, others vary row and column width. For example, table lens allows users to increase the area of table entries of interest while preserving the context of all other (smaller) table entries (Rao and Card 1994). A mosaic display visualizes table entries as tiles whose visual area depends on the part of the data represented (Friendly 2002).

However, there are $\frac{N(N-1)}{2}$ table entries to inspect for N dimensions. Moreover, if there is no optimal one-dimensional ordering of the dimensions, then there is no optimal ordering of the dimensions along the rows and columns of a tabular layout.

Glyphs

A graphical object called *glyph* (or icon) conveys multiple data values at a single position (not a single pixel) (Ward 2002, Schroeder and Martin 2004, Ware 2004). A simple illustration is a glyph called whisker (Fig. 1.10, left). Each data value is represented by a line segment connected to a central point (with a longer line segment for a higher value). Consecutive data values are separated by a certain angle. An extension of the whisker is called star plot (or star glyph) (Ware 2004), showing a line connecting the ends of the line segments (Fig. 1.10, right).

Glyph positioning and the use of large numbers of glyphs, resulting in a visual texture, are important issues to consider (Ward 2002, Ware 2004). The number of data values to convey with one glyph is not easily extended without losing an overview.

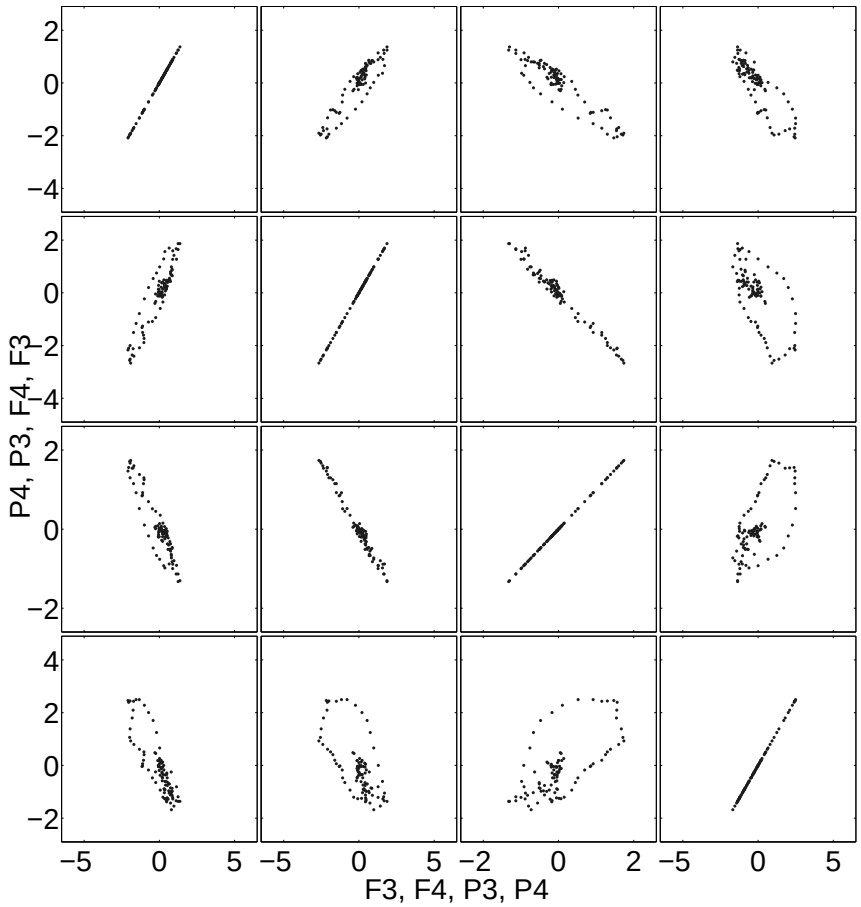

Figure 1.9. *Scatter plot matrix for four dimensions corresponding to four electrodes (labeled F3, F4, P3, P4, from top to bottom and from left to right). Each matrix element is a scatter plot. Data represent 100 time steps of an (averaged) SEP, in this example following electrical stimulation of the median nerve near the left wrist. High values for F3 and F4 and low values for P3 and P4 co-occur, and vice versa.*

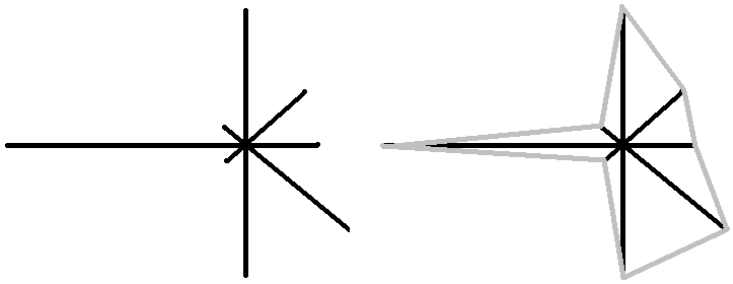

Figure 1.10. *A whisker representing an eight-dimensional data vector (left) and its corresponding star plot (right).*

Parallel Coordinates, Circular Coordinates, and Extruded Parallel Coordinates

The parallel coordinate method (Inselberg 1985) represents each data dimension as a (usually) vertical axis. For an N-dimensional vector $(x_1, x_2, ..., x_N)$, N uniformly spaced parallel axes are

used. To display a single vector, each vector element is indicated by a dot at the corresponding vertical axis, and all dots for a single vector are connected by a single polyline (Fig. 1.11).

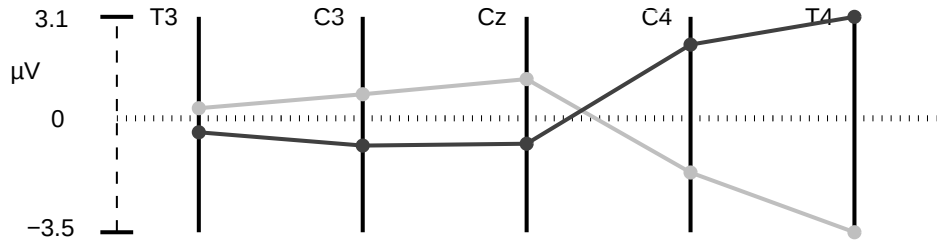

Figure 1.11. *Parallel coordinate representation for two five-dimensional vectors, each of which represents one time step. The data have been recorded from five EEG electrodes simultaneously (labeled T3, C3, Cz, C4, and T4). For each vector, one polyline is drawn. The voltage (µV) is set out vertically (see scale on the left).*

Circular coordinates organize multiple axes as spokes in a wheel (Siegel *et al.* 1972), resulting in an organization similar to whiskers and star plots. In fact, a star plot can be considered as a circular coordinate plot for one polyline, with axis length depending on the value in the corresponding dimension. Extruded parallel coordinates use a two-dimensional organization of parallel axes, which are orthogonal to the plane that intersects the middle of all axes (Wegenkittl *et al.* 1997). The same axis organization is referred to as three-dimensional parallel coordinates in a more recent publication (Rübel *et al.* 2006), whereas Ware (2001) would probably refer to it as 2.5D parallel coordinates. Additional features may represent the variation of data vectors along an axis (Hauser *et al.* 2002, Yang *et al.* 2003, Stump *et al.* 2003) or between axes (Johansson *et al.* 2007).

Although the axes can theoretically be put in any desired order, in practice the ordering might affect the data analysis. Further, the use of a large number of polylines causes visual clutter. Solutions include the visualization of clusters of polylines (Fua *et al.* 1999). An extension of parallel coordinates to three-dimensional space potentially adds occlusions.

1.2.2 Relational Data: Graphs

A graph (or network) consists of vertices (or nodes) and edges (connecting pairs of vertices). A graph can be represented by an adjacency matrix, with vertices organized along both the rows and the columns. A matrix entry then indicates an edge value for the corresponding vertex pair.

Straightforward graph visualizations display vertices as dots and edges as lines (Fig. 1.2). In graph layouts, vertex and edge densities can be so high that individual vertices and edges cannot be distinguished. Moreover, edges and vertices can obscure each other and other visual information. Multiple solutions exist to reduce such cluttering, e.g., changing the layout of the vertices (Fruchterman and Reingold 1991) or the edges (Wong *et al.* 2003). Other solutions vary visual attributes of vertices and edges (Srinivasan *et al.* 1999, Herman *et al.* 2000, Chen *et al.* 2003, Wong *et al.* 2003, Salvador *et al.* 2005*a*).

Alternatively, mathematical methods are employed to recognize groups of similar vertices, i.e., graph clustering. However, there is no single algorithm that produces the best clustering for any application. Therefore, a wide variety of graph clustering methods exists with quality depending on the application (van Dongen 2000).

1.3 Thesis Contribution & Organization

This thesis describes new visualizations tailored for the use with multichannel EEG. One is for time domain analysis, the other for frequency domain analysis of EEG.

For the time domain (time-varying multivariate data), Chapter 2 introduces an EEG visualization based on a tiled organization of parallel coordinates. The so-called *tiled parallel coordinate (TPC) map* uses one tile per electrode and schematically preserves electrode positions by using a two-dimensional tabular layout of these tiles. For time steps of interest, potentials measured at all electrodes are visualized by parallel coordinates. Meanwhile, potentials measured for other time steps are preserved as a context. Time information is preserved by linking the TPC map with conventional time series visualizing the overall variation of the EEG. The TPC map is assessed by both a qualitative evaluation and a user evaluation. Contrary to other methods, the TPC map is able to visualize the combination of a large number of electrodes with many time steps. In comparison with an existing clinical EEG visualization method, the TPC map is found to be faster without a loss of information for a typical EEG assessment task.

Frequency domain analysis on the basis of EEG coherence, employed for the study of functional brain connectivity, is discussed in Chapter 3. A method is presented for data-driven coherence analysis which is designed for the use with multichannel EEG. In a graph representing EEG coherence, the method detects data-driven regions of interest (ROIs) referred to as functional units (FUs). For individual dataset analysis, a so-called *FU map* visualizes the coherence between FUs. In comparison with other data-driven visualizations of multichannel EEG coherence, the FU map strongly reduces visual clutter. Additionally, two group maps are employed for group analysis of multichannel EEG coherence. First, the *group mean coherence map* preserves dominant features from a collection of individual FU maps. Second, the *group FU size map* visualizes the average FU size per electrode across a collection of individual FU maps. The FU map and both group maps all preserve electrode locations.

This new method for multichannel EEG coherence analysis based on the FU map for individual analysis combined with both group maps is applied in Chapter 4 to a mental fatigue case study without hypotheses. Contrary to a hypothesis-driven approach which depends on previous experiments, our method leads to a selection of data-driven ROIs and coherences of interest between those ROIs which are specific for the actual experiment. It summarizes extensive experimental results which otherwise would be very difficult and time-consuming to assess.

Chapter 5 provides a summary of this thesis and conveys conclusions. Finally, directions are indicated for future research.