Erkka Heinilä

Exploring EEG recordings of Focused Attention Meditation with Fourier-ICA

Master's Thesis in Information Technology

May 6, 2017

University of Jyväskylä

Department of Mathematical Information Technology

Author: Erkka Heinilä

Contact information: erkka.heinila@jyu.fi

Supervisors: Tapani Ristaniemi, and Tiina Parviainen

Title: Exploring EEG recordings of Focused Attention Meditation with Fourier-ICA

Työn nimi: Keskittymismeditaation tunnistettavuus EEG-aineistossa Fourier-ICA-menetelmällä

Project: Master's Thesis

Study line: Computational science

Page count: 50+9

Abstract: Practicing meditation is often associated with many health benefits. With Electroencephalogram (EEG) it is possible to keep track of temporal changes in the nervous system during a meditation session. Findings in this area can give information about the neuronal responses of meditation, but also in general about the brain in all naturalistic conditions. This study uses a variant of Independent Component Analysis (ICA) called Fourier-ICA to analyze EEG data from meditation sessions and shows that ICA can be used to do much more than just remove artifacts, which is the normal use case in EEG studies. In particular, the answer is sought for the question of whether it is possible by these means to differentiate between neural correlates of focused and wandering mind.

Keywords: meditation, mindfulness, alpha, eeg, ica, tfr, fourier-ica

Suomenkielinen tiivistelmä: Meditaation harjoittamiseen liitetään usein useita terveyshyötyjä. Elektroenkefalografia (EEG)-menetelmällä on mahdollista seurata hermostossa meditaatioistunnon aikana tapahtuvia ajallisia muutoksia. Löydöt tällä alueella antaisivat lisää tietoa erityisesti meditaation aivovasteista, mutta myös yleisempää tietoa aivojen toiminnasta luonnollisissa olosuhteissa. Tässä tutkielmassa käytetään Fourier-ICA:ksi kutsuttua itsenäisten komponenttien analyysin muunnelmaa meditaatioistunnoista saatujen EEG-aineistojen analyysiin ja näytetään että itsenäisten komponenttien analyysiä voi käyttää muuhunkin kuin vain häiriöiden poistamiseen, mikä on sen tavallisin käyttötapaus EEG-tutkimuksissa.

Erityisesti tutkielmassa etsitään vastausta kysymykseen ovatko keskittymisen ja ajatusten harhailun aivovasteet tällä menetelmällä erotettavissa toisistaan.

Avainsanat: meditaatio, mindfulness, alfa, eeg, ica, tfr, fourier-ica

Glossary

CLT	Central Limit Theorem
EEG	Electroencephalogram
FFT	Fast Fourier Transform
fMRI	functional Magnetic Resonance Imaging
FT	Fourier Transform
ICA	Independent Component Analysis
MEG	Magnetoencephalogram
PCA	Principal Component Analysis
PSD	Power Spectral Density
PP	Projection Pursuit
STFT	Short-Time Fourier Transform
TFR	Time-Frequency Representation
TSE	Temporal Spectral Evolution
TM	Transcendental Meditation

List of Figures

Contents

1 Introduction

Learning and practicing meditation is often associated with health benefits such as emotion regulation, cognitive control, stress reduction and feeling of fuller life. It is not however clear what the underlying brain mechanism behind the health benefits of meditation is. In this study we are interested in yet another open question: what happens in the brain during meditation practice. Reliable information on this area should help to understand the meditation itself, but also its claimed benefits.

As Tang, Holzel, and Posner (2015) states, meditation research is still at its infancy and many studies have had low methodological quality. Issues include small sample sizes, relatively high bias for publishing positive results, and lack of proper control for variables that might be the real reasons behind the results but not being integral aspects of meditation. The problems are not necessarily a sign of research conducted poorly, as there are also specific complicating factors in studying meditation, for example in variable controlling. And as it is often difficult to collect data in great quantity, some novel methods are needed to overcome this problem.

For this study Electroencephalogram (EEG) data from 30 meditation practitioners of different expertise was used. They each participated in a 10 minutes session where they were asked to focus their undivided attention to their breathing. If they had any thoughts unrelated to the task, they were instructed to press a button and then return their focus to the task.

Independent component analysis (ICA) of Time-Frequency Representations (TFR) was then used to extract relevant information out of the noisy EEG recordings. This method allowed such improvement on the signal-to-noise ratio, that the frequency distributions calculated from the ICA components revealed a difference in the amplitude of 10 hertz neural oscillations at the back of the head in the state of mind wandering and the state of focused attention. This is in line with a hypothesis that 10 hertz oscillations have an inhibiting role, and the vision-related brain areas at the back of the head are inhibited more when the task is to keep focus in breathing sensations.

The same feature extraction method was also used in combination with Temporal Spectral

Evolution (TSE) to allow more concise look at the behavior of 10 hertz oscillations near the button presses. After averaging all of the button-related TSE data segments together, some interesting patterns emerged, but because the time when mind wandering actually started was unknown, it was hard to get any conclusive information about the relationship of meditation and mind wandering.

The report will proceed as follows. First the necessary theory and underlying concepts are explained; we will start from meditation and basics of EEG and continue up to theory behind the Independent Component Analysis used on TFR's. Then the analysis section explains the paradigm used in the study and shows and analyzes the results of using the methods and theory developed before to the collected data.

2 Underlying concepts

2.1 Focused Attention Meditation

Why study meditation? Doesn't the task of focusing on breathing and keeping other conscious thoughts at minimum, make brains less active or less interesting? To really appreciate the effort to study calm mind without no specific problem solving, it should be understood that for a brain the involvement with the environment by sensing and acting is only a temporary distraction in a complex self-sustained brain activity (Buzsaki 2011). Traditional approach in EEG studies has been to show a research subject some visual or auditory stimuli and look for the directly following responses. But in that case the information about neural responses might only reflect transient phenomena leaving out the activity that is present all the time. To reveal the world behind only sensing and acting, other kind of view point could be taken. Meditation gives a good window to look what it is that actually happens in the brain when it is not distracted. In addition to this, the claimed health benefits of meditation should justify the quest for more comprehensive understanding of meditation.

Meditation comes in many varieties that include but are not limited to Yogic meditation, Vipassana meditation, Transcendental Meditation (TM), and more secular Mindfulness meditation. They all have their own nuances and purposes, but many of them have concentration as one of their key aspects or methods (Saggar 2011). As Mindfulness meditation has received a lot of attention in neuroscience research in the past few decades (Tang, Holzel, and Posner 2015), this study also focuses on a practice contained in that tradition, namely Focused Attention Meditation, where the meditator is focusing her attention (usually) to sensations of breathing. Here the interest is in finding neurophysiological correlates of participant being concentrated to the task in comparison to her having wandering thoughts.

It is known that task of attending to breathing has two key aspects. One is staying focused on the breathing and the other is being aware of how well it is going (Saggar 2011). When mind wandering is noticed, attention is then directed back towards the breathing. Functional Magnetic Resonance Imaging (fMRI) studies confirm that the moment of mind wandering in the middle of a breathing task can be divided to four different phases: mind wandering,

awareness of mind wandering, shifting of attention back to the task, and being focused to the task. In the study by Hasenkamp et al. (2012) all the phases were given a block of 3 seconds and had distinguishable brain responses.

Even though the EEG studies have had a lot of problems including mixed results (Tang, Holzel, and Posner 2015), there are some results that can be carefully stated. It would seem that neural oscillations of around 10 times a second on many areas of the brain are stronger for people with meditation experience compared to people with no experience, and on average are stronger during meditation session compared to not meditating for all people. It is only correlation though, and it is not easy to attribute these changes specifically to meditation (Cahn and Polich 2006).

2.2 Electric activity of the brain

It was stated before that neural oscillations around 10 times a second somehow reflect the meditative state. What are these oscillations?

To do information processing brains need to transfer information and this information is transferred by electric signals. Electric signals are ions, which are certain kind of atoms, moving from place to another. Electric signals are mediated by neurons, which are cells that specialize in electric signaling and locate in brain. Information processing then is activating and deactivating collections of neurons in specific patterns and locations. Signaling can happen in nonrecurring way in which case the information is transmitted onward from one neuron to other and no further actions are needed. Sometimes that is not the case, and then there either is a source that makes neuron fire in repetitive fashion or there is a loop of activity which results in the neuron firing repetitively. When the firing is recurring, the activity is said to be oscillatory.

Electric activity, or the behavior of ions, can be measured with electrodes. Electrodes can be inserted inside the brain to track directly firing of even single neurons, but measurements can also be done non-invasively from outside the brain, exchanging accuracy for convenience. The data in this study comes from an Electroencephalogram (EEG) device. An EEG device consists of a grid of electrodes attached to scalp and it measures electric potential differences

between these electrode sites and some reference site. In simple terms electric potential means the ability of a electric current to stream from one place to another, and in this case the measurable electric currents come from the inhibitory (i.e making neuron further away from firing) and excitatory (i.e making neuron closer to firing) ions flowing in and out from neuron cells giving different ion concentrations that attract each other (Luck 2005). This information about the ion concentration differences then also gives indirect information about the firing of neurons.

Potentials from single neurons are hardly measurable because of the resistivity of the scalp, the skull and the cerebrospinal fluid, and because of the long distance between electrodes and source of activity. For this reason the potentials that can be measured by electrodes will have contributions from large populations of neurons that are well synchronized, i.e their potentials sum up instead of canceling each other out. This synchronous activity that mainly comes from cerebral cortex then gets recorded by the electrode-pairs in a form of a time series where voltage, a measure of difference between potentials, is represented as a function of time (Teplan 2002). These representations can then be converted into digital form, i.e to sample at some rate and store to a hard drive with desired accuracy, to allow further offline processing. Figure 1 shows example of data recorded and stored from a EEG measurement.

EEG devices come in many forms with different topographical layouts and with different numbers of electrodes. In this study the used EEG device had 128 site electrodes spread across the scalp. Because electricity spreads out in every direction in a conductor and because of the varying resistivity properties of the brain matter between electrodes and sources, different electrodes will also share contributions from same sources of activity and the contributions do not necessarily come uniformly. This means that even if there is seemingly good accuracy on the scalp, it is not trivial to deduce the source locations of the activity (Luck 2005).

When doing a EEG measurement in continuous fashion, often many oscillatory patterns emerge. As these patterns emerge from large populations of synchronous neurons with small delays, this sum of activity visible in EEG can actually look like a curve of harmonic oscillator, i.e the curve is sinusoidal (Figure 2). These oscillative patterns emerge at many different

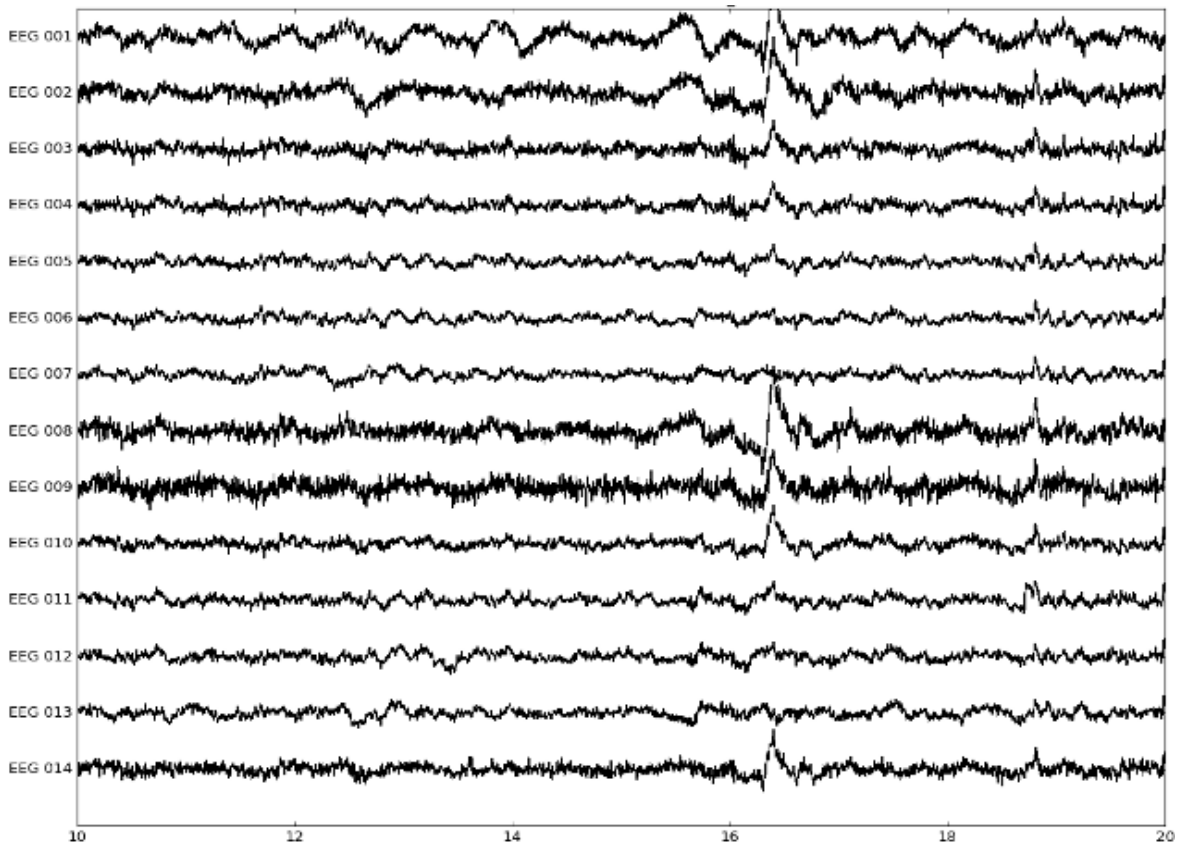

Figure 1. EEG data. 10 seconds of EEG data shown with MNE-python (Gramfort et al. 2013).

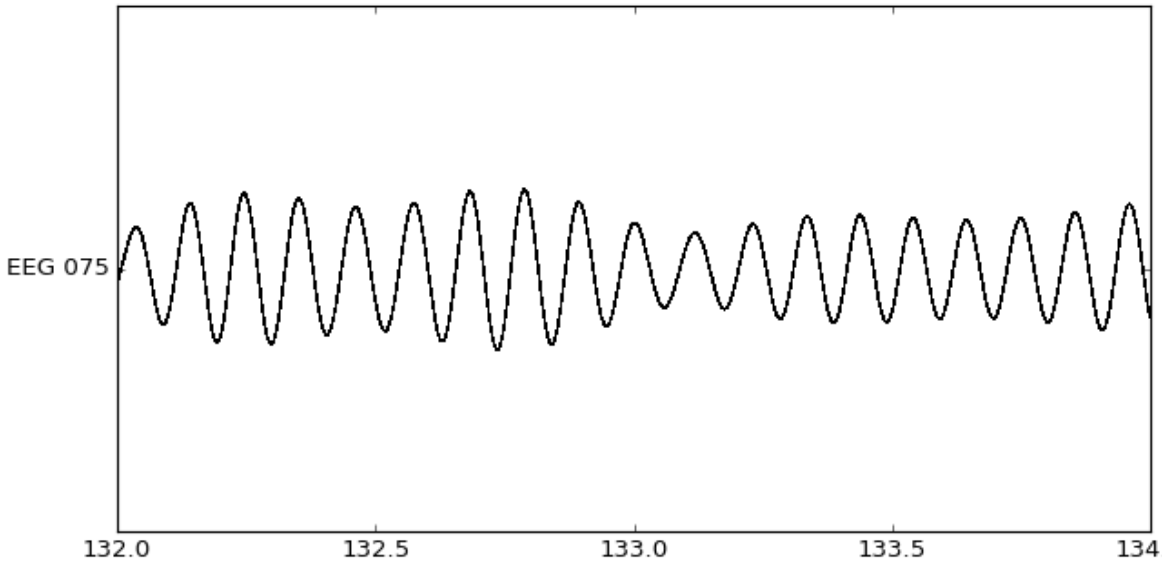

Figure 2. Alpha oscillations. Example of neural oscillations occurring at 10 times a second. Data in this channel was first band-pass filtered to 9-11 hertz and then plotted with MNE-python (Gramfort et al. 2013).

frequencies (i.e how many times a event happens in a second) and some of the most usual ones have been named. Slow delta waves, that are often associated with different stages of sleep, are firing close to 0.5-4 times a second. Little faster theta waves, believed to partake in learning and memory retrieval, oscillate at around 4-8 times a second. There are beta waves, that relate to movement control, at 12-30 times a second, and gamma waves, believed to solve a binding problem of different neuronal populations, at $>$ 30 times a second. (Buzsaki 2011) For this study some preliminary exploration of the data and previous studies guided to look specifically for oscillations around 10 times a second, that have been named as alpha waves.

Alpha waves are neuronal oscillations having a frequency around 10 times a second. The most prominent alpha waves can usually be found at the posterior (i.e back of the head) sites. The generators and the role of alpha are still a matter of debate, but it is often thought to have a inhibiting function, meaning that the brain area that is exhibiting a powerful alpha rhythm is not participating in any task-related processing so that some other brain area can function more efficiently. Basic example of this can be seen by closing one's eyes. When eyes are open, the occipital lobe (located at the back of the head) has a lot of visual information

processing to do and the alpha power at that area is weak. When eyes are then closed, there is much less visual processing to do and the alpha rhythm becomes more powerful. (Klimesch, Sauseng, and Hanslmayr 2007)

As a simple note from this it could be hypothesized that in the breathing task there is a decrease of alpha in somatosensory areas, as one might feel the breathing when focusing to it, and increase of alpha in other sensory areas. Indeed, in this study the posterior alpha is looked for to find that kind of effects.

Next sections will cover the methods that will be used to analyze the time series data collected with EEG.

2.3 Fourier Analysis

A very often used method for analyzing time series data is to look at its spectral content. The word spectrum originally referred to colors of the rainbow, but has a very similar meaning in time series analysis. Seeing different colors of the light means sensing electromagnetic waves of different wavelengths. White light contains all of the colors at the same time. Similarly a time series can consist of waves of different wavelengths, and sometimes some wavelengths dominate the others. And as spectral differences of light can reveal information about the events of atmosphere, so can spectral differences in time series reveal information about the underlying phenomena of interest. But where are the waves in time series?

According to Fourier Analysis, any time series can be decomposed to simple waveforms that when summed together will form the original time series. This is illustrated in Figure 3, where the bolded curve represents a complicated time series which actually is just a pointwise sum of the few selected curves. These curves are sine waves, and they are completely defined by their amplitude (i.e the height), phase (i.e the horizontal location) and frequency, which is just the inverse of wavelength and means how quickly the waveform repeats. Spectrum of a time series then tells how largely each of the sine waves (with unique frequency) contributes to the time series, not unlike how the spectrum of light gives the relative strengths of different colors. One might think that there is a need for multiple sine waves with the same frequency but with different phases, but this is not the case as a sum of two sines with same

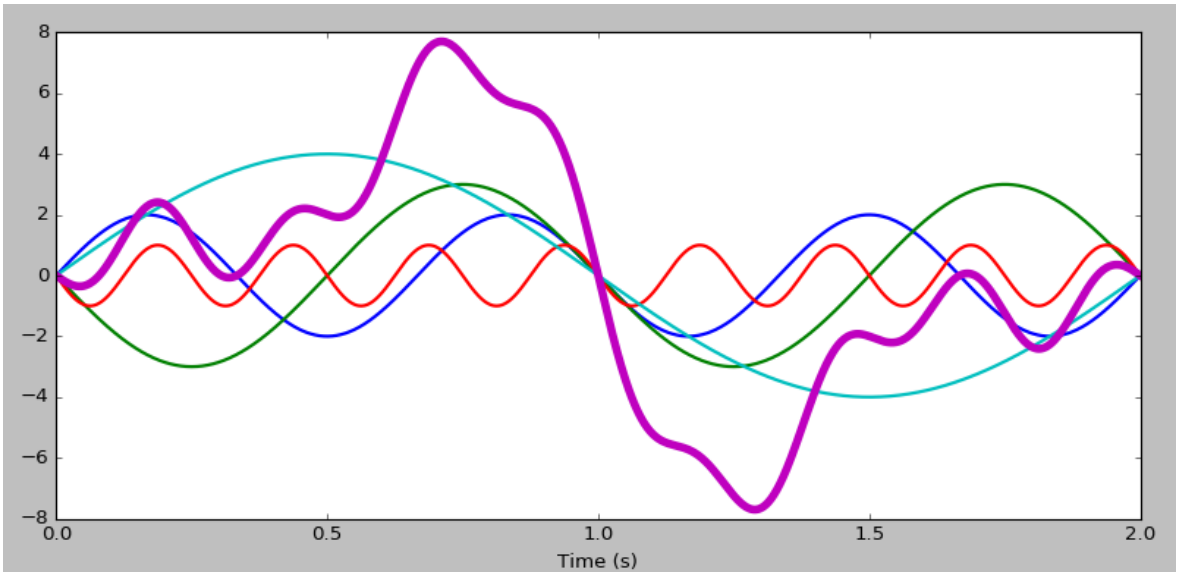

Figure 3. Time series decomposed to sinusoids. Picture shows one complicated time series (bolded) and the sine waves that can be used to construct it.

frequency and different phases is actually also a sine wave with the same frequency and can be used in the representation instead of the two (III 2007).

Representation of time series as a sum of its sinusoidal components is called its Fourier Transform (FT) and computationally it can be found with Fast Fourier Transform (FFT)-algorithm, for which the details can be found for example in the original paper Cooley and Tukey (1965). FFT calculates a discrete version of FT called the Discrete Fourier Transform (DFT). DFT is frequency-domain sampled version of Discrete-Time Fourier Transform (DTFT), which relates to FT by the following equation:

$$X_{dt}(f) = (1/dt) * X(f), \tag{2.1}$$

where $X_{dt}(f)$ is the DTFT, dt equals to the time interval between adjacent samples and $X(f)$ is the FT. If only the strengths of underlying sine waves in a time series are of interest, i.e phase information is not needed, a Power Spectral Density (PSD) is often calculated. The PSD is calculated as:

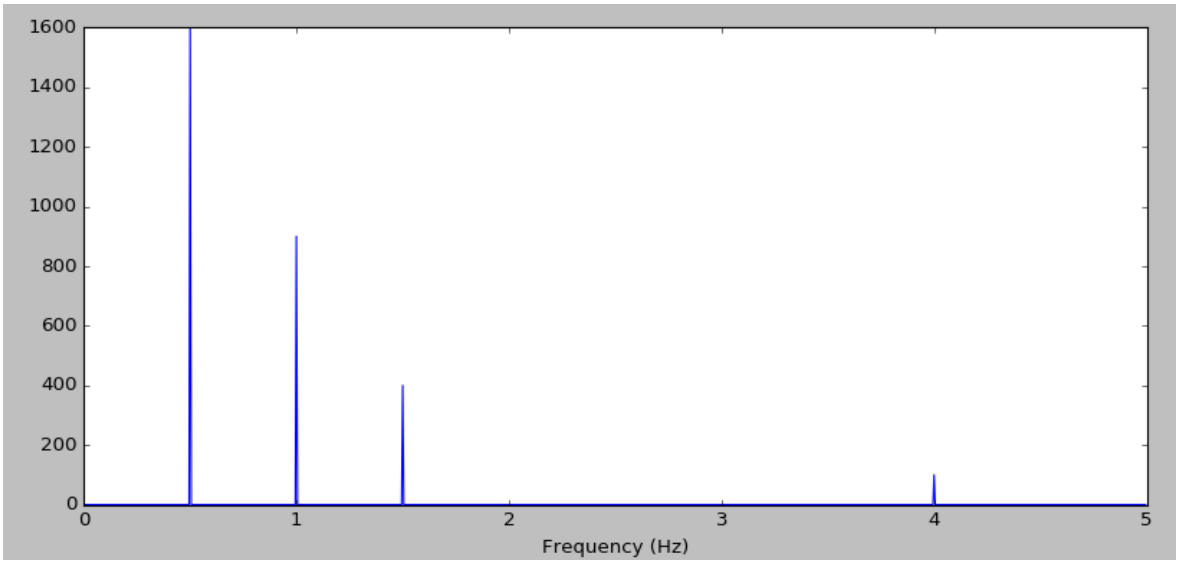

Figure 4. Spectrum of time series. Figure shows PSD of the time series in figure 3. Every peak corresponds to a sine wave with some amplitude and frequency.

$$PSD(f) = |X(f)|^2/T, \tag{2.2}$$

where T is the total duration in time for the time window (Stoica and Moses 2005). Combining these two one thus gets the discrete PSD as

$$PSD(f_k) = |X_{dt}[f_k] * dt|^2/T \tag{2.3}$$

Basically this gives squared amplitudes (in the context of EEG, the unit is V^2/Hz) as a function of frequency. Figure 4 shows the PSD of the curve in figure 3, where one can clearly see the peaks of four different frequencies that represent the strengths of the underlying sine waves. The oscillatory activity present in brains can also often be seen as peaks in the PSD's as will be seen later on.

One important assumption that needs to be obeyed when calculating Fourier Transform so that the results will be sensible is that the time series should be stationary. This means that the spectral content of time series should not change in time or equivalently that the underlying sine components should not modulate in time, i.e their amplitude should stay constant. This

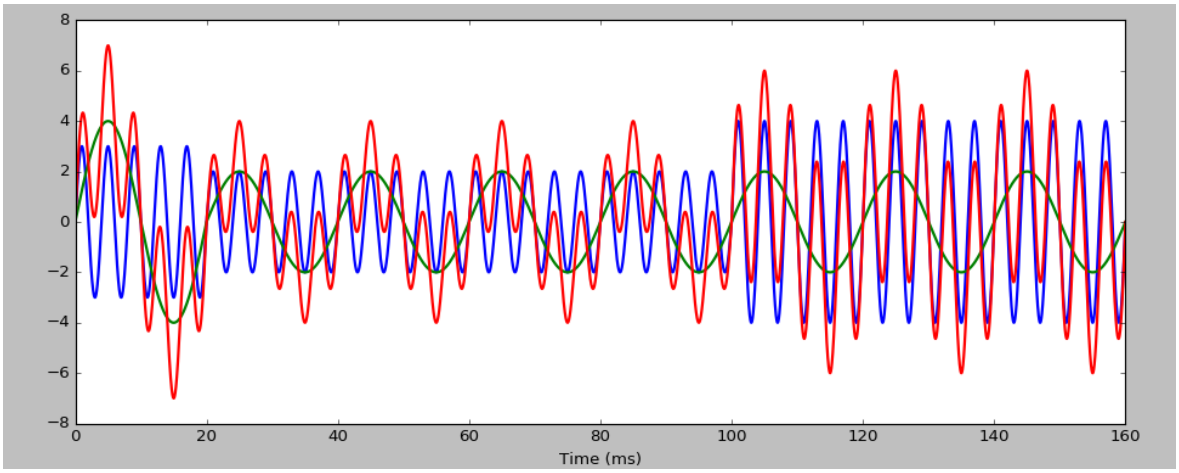

Figure 5. Time series decomposed to sinusoids 2. Red series is a sum of blue and green series that are amplitude-modulated (i.e their amplitude changes as a function of time).

is not normally the case in brain data where the underlying components do modulate even when the subject is just resting. What is usually done is that the Fourier Transform is computed in short time windows, where the stationarity holds better, and resulting transforms are averaged together. This is called the Welch's method and it gives the average spectral content of a time series.

Sometimes, however, the temporal changes of spectral content can also be a matter of interest. This is the case in this study, where it is assumed that there are different states that the subject is going through in her mind, i.e she could be focused or not focused, and it could be that there are amplitude changes in the oscillatory activity reflecting these states of mind. Then the method of calculating spectrums at different time windows and averaging them together will lose just that information. On the other hand, if the averaging part is left out, one will have a set of spectrums each representing the spectral content of the corresponding time windows. This is exactly what was looked for as now the spectral changes can be tracked by following the sequence of spectrums. This method, called the Short-Time Fourier Transform (STFT), leads to Time-Frequency Representation (TFR), which is a colored grid of points where each element represents the power of the sine wave corresponding to some frequency and time, and going along one axis changes the frequency while going along the other axis changes the time. This is illustrated in figures 5 and 6.

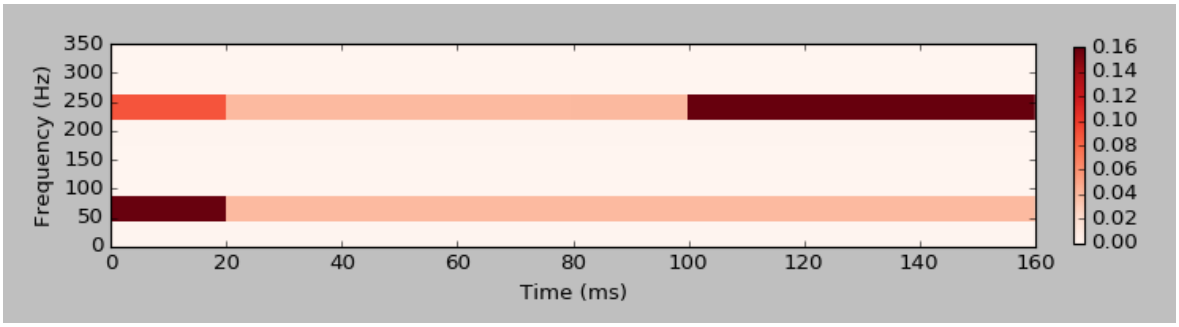

Figure 6. TFR of time series seen in 5. In the figure one can see two lines representing the two sine waves and travelling along the x-axis shows how their amplitudes change (indicated by the color).

It should be noted that this procedure will for information theoretic reasons necessarily lead to a compromise between accuracy of time and frequency. This relationship is captured by the formula:

$$Fr = Sr/Ws, \tag{2.4}$$

where Fr is the frequency resolution (i.e the width of frequency bins), Sr is the sampling rate and Ws is the window size. When the time series is divided into smaller segments, which means more accuracy in time, the spectrum calculations will have less samples to use and this results in less accuracy for frequencies. Other thing to note is that frequency resolution can be improved by downsampling the signal, but this will also narrow the frequency bandwidth.

Related to Time-Frequency Representations, Temporal Spectral Evolution (TSE) also targets to track changes in oscillatory activity, but instead of focusing on the whole spectrum, a small band of frequencies is selected for analysis. TSE can be seen as a special case of TFR, but is usually calculated and visualized in a different way. As a first step time series is band-pass filtered, which means retaining frequency content inside some band of frequencies and leaving everything else out. In terms of sine waves this means that sine waves with frequencies out of this band are subtracted away from the time series. Then the resulting series is rectified, i.e absolute values are taken, and smoothened, to allow plotting of amplitude as a function of time. With this the possibly transient nature of the frequency band gets revealed

(Hari et al. 1997).

Now that the basics of time series analysis are known, we are ready to go into the realm of Independent Component Analysis (ICA).

2.4 Independent Component Analysis

EEG data might get recorded with hundreds of sensors and a sampling rate of over a 1000 samples a second for long periods of time for multiple subjects, which results in hundreds of millions of biomedical values stored in a hard drive. The answers that are looked for, on the other hand, are usually kind of simple, like does some statistic differ significantly in two different conditions, and thus there is a need for methods that step by step can extract the relevant information out from the data, finally answering the question asked. Using a 128-sensor EEG cap, like in this study, allows for more accuracy on the scalp than for example a 19-channel EEG cap. But as there might only be a limited amount of concurrent neural processes distributed to a limited amount of brain areas, then if all the sensors are capturing real brain activity, this means that same brain signals end up recorded in multiple sensors, thus on principle giving every sensor some kind of mixture of all the brain activity in different areas but with different strengths (though some being negligible). This makes the analysis harder as it is not possible to label for example some group of sensors to only relate to some specific neural process.

Traditional EEG research is based on having a subject repeat a given task so many times, that it is possible to average out the task-independent brain activity from the recording when doing the analysis. This works because it is assumed that the brain process of interest is having similar kind of activity every time the task is done, but other processes are not locked to the task in the same way. Thus their activity will cancel out on average, but the activity from the process of interest will not (Luck 2005). This way the mixture problem vanishes and the researcher can then compare the simple waveforms in different conditions to get the results being looked for. This approach cannot be used with data that does not have this kind of repeated task structure, however, which is the case in this study, where there is just one continuous task of resting or meditating. For this kind of continuous data, other

kind of approach is needed. One way is to select one sensor, or a group of nearby sensors, and use Fourier Transform to analyze only a small band of frequencies and assume that the mixture problem becomes small enough. This can be done, but in this study it did not lead far enough, perhaps because there was too much noise and more data would be needed, or maybe there still was some cancellation even inside such small band of frequencies. Anyway more advanced signal processing was needed, and as was hinted in the beginning of this section, this is a data mining problem: a lot of data, and the answers lie there somewhere. Could the general data mining methods be of use?

For clearness, a simple-to-use representation for data is introduced here. If data consisted of two simultaneous time series recordings, one can imagine that there is a point in *xy*-plane corresponding to each time point. If at the first time point the first recording had value 1 and the second recording value 2, then this would correspond to a point $(1,2)$ on the *xy*-plane. As the time goes on, the point then travels in this two-dimensional space. In a similar way, if there was 128 or n simultaneous recordings, then this would correspond to a point travelling in 128- or n-dimensional space, which might be hard to imagine, but works analogously. With this idea, let $x^i = (x_1^i, x_2^i, ..., x_n^i)$ be a vector representing ith sample from a n-dimensional time series.

Most of the following content in this section is based on Hyvärinen, Karhunen, and Oja (2001). With Principal Component Analysis (PCA) it is possible to transform the data to a lower dimensional space in a way that the resulting data will retain most of the information possible given the dimensionality reduction, and is for this reason often used in data compression. Formally the transformation to k-dimensional space can be done with a matrix multiplication:

$$y^i = Mx^i, \tag{2.5}$$

where $y^i = (y_1^i, y_2^i, ..., y_k^i)$ is the transformed sample and M is a $k \times n$ projection matrix. The nice property of this transformation is that by selecting a good value for k, a lot of redundancy (and noise) can be removed from the data without affecting the bigger picture much. However, this could introduce even more difficulty to interpret the resulting data, as then the

time series will not directly correspond to either brain processes or the sensors, but instead the largest directions of spread, which might not be very intuitive. Thus even if the PCA helps with the redundancy problem, it is not a tool that makes the data necessarily more understandable. However as will be seen, it can be used as a helpful preprocessing step for another more suited method.

Projection Pursuit (PP) is a similar method to PCA in the sense that the data is transformed to another space, perhaps lower-dimensional, but instead of trying to find most varying directions, it puts emphasis on the "interestingness" of the directions. It turns out that the heuristics for maximizing "interestingness" of the projections that got developed in the field of Projection Pursuit are very similar to the principles of maximizing statistical independence of the projections.

Independent component analysis (ICA) is a generic computational method to divide a multidimensional time series to it's underlying statistically independent components. It specifically tries to solve the mixture problem that was encountered earlier, given that the original sources that were mixed together were independent of each other, by unmixing the observed time series and finding out the original underlying time series, that in the case of EEG could be the different neural processes working independently. Formally ICA assumes the following model:

$$x = As, \tag{2.6}$$

where $x = (x_1, x_2, ..., x_n)$ is the observed n-dimensional random vector, A is a $n \times n$ unknown mixing matrix, and $s = (s_1, s_2, ..., s_n)$ is the hidden original unknown n-dimensional random vector, the source vector. Our data comes in form of a time series, but it can be thought as a collection of realizations of a random vector if the time structure is forgotten and it is assumed that there is a probability distribution that when sampled from might give the values of the time series. After the unmixing matrix has been found using the random vector interpretation, it is then possible to transform the time-structured observed series with unmixing matrix to the time-structured source series. Matrix A is not constrained in any way so it is possible for all the observed variables to be weighted sums of all the source variables. Di-

mensions are assumed to be equal, meaning that there are as many source variables as there are observed variables. This might not be realistic assumption, but there are ways to take care of that which will be discussed later on. The question ICA then asks is that: is there a way to uncover the mixing matrix and the source distributions when only the observed mixtures are known? As it turns out, this is possible if s_i are mutually statistically independent and their statistical distributions are nongaussian.

The first condition, the mutual independence of random variables, means that knowing the realization of any variable does not give information about the realizations of the others. Formally with conditional probabilities this can be stated as follows:

$$p_{X_i}(x_i|x_j) = p_{X_i}(x_i), \forall i, j \in I, i \neq j, \tag{2.7}$$

where $p_{X_i}(x_i)$ is the probability density function of the random variable X_i and I contains indices of the random variables. This condition for the method is intuitive as the task is to find these existing independent components.

The second condition, necessity of nongaussianity of source distributions, comes from the fact that joint distribution of independent gaussians with equal variances is spherically symmetric and is invariant to rotations. Thus the correct rotation of the source distribution cannot be found, and it is not possible to define the mixing matrix as uniquely as it is possible with nongaussian distributions. The best one can do with a mixture of gaussians is to uncorrelate them, which is basically the same that the PCA does. Often the distributions in applications, if they are not sums of independent distributions (as they tend towards gaussians), will not be gaussians, and thus this will not be a problem. It also turns out that the key to find the independent components is actually to find as maximally nongaussian distributions as possible. This intuitively follows from Central Limit Theorem (CLT), as linear mixtures of independent random variables tend to be more gaussian than the original ones, but this can also be proven rigorously. This pursuit of nongaussianity can intuitively be accomplished by finding probability densities from the mixtures that resemble a gaussian as little as possible, usually meaning that the density is sparse, i.e it has a pointy peak and heavy tails.

It can be shown that under some natural assumptions the independent components found by ICA are unique up to a scale. Scale ambiguity comes directly from the model $x = As$, where multiplying any s_i with factor α can be canceled by multiplying corresponding row in A with $1/\alpha$. As the scale is arbitrary, one might as well define it to be such that every s_i has unit variance: $E\{s_i^2\} = 1$. This still leaves the ambiguity on sign, but that is usually not a significant problem.

So how then can the mixing matrix be estimated? Optimization landscape for finding one of the independent components is actually simple. Formally the objective can be stated as follows. Given the observed random vector $x = (x_1, x_2, ..., x_n)$, find $w = (w_1, w_2, ..., w_n)$ such that

$$y = w^T x, \tag{2.8}$$

and y is as nongaussian as possible and $E\{y^2\} = 1$. To get the whole unmixing matrix $W = A^{-1}$, this optimization problem is then solved n times with constraints that prevent finding the same optima twice. In practice the new solutions are constrained to be uncorrelated to previous ones, as it is known that independent components are always uncorrelated. Alternatively the whole unmixing matrix can be estimated at once, by estimating components in parallel and decorrelating all components at each iterative step, but the main principle is the same. The FastICA-algorithm, which does this estimation quickly and robustly, and is used in this study, is explained next.

First the data is centered and whitened, as this simplifies the calculations done later. Centering means subtracting the mean from the data and whitening means decorrelating and scaling in such a way that the covariance matrix becomes the identity matrix. Let $x = (x_1, x_2, ..., x_m)$ be the observed random vector as before and $z = (z_1, z_2, ..., z_n)$ be the whitened random vector. The task is to find $n \times m$ matrix V such that:

$$z = Vx \tag{2.9}$$

This can be done with PCA. Without going to too much detail, PCA tries to find the most

varying directions in the data, and they prove out to be the eigenvectors of the covariance matrix. When the data is then transformed to align with these eigenvectors, the covariance matrix becomes diagonal, and after scaling the data with the corresponding eigenvalues of the covariance matrix, the data becomes white. Formally V is just:

$$V = D^{-1/2}E^T, \tag{2.10}$$

where $D = diag(d_1, ..., d_n)$ with n first eigenvalues of the covariance matrix, and $E = (e_1, ..., e_n)$ is a matrix whose columns are the n first unit-norm eigenvectors of the covariance matrix.

In this phase it is also possible to reduce the dimension of data by letting $n < m$. By doing that one can reduce noise, and fix the situation where the number of independent components is smaller than the number of mixtures, by reducing the dimension to be the same as the number of independent components.

Turns out that after the data has been whitened to get z, only orthogonal transformation (i.e rotation) is needed to get the independent components. To maximize nongaussianity, which corresponds to finding maximally independent components, a measure for it is needed. Classical measures include negentropy and absolute value of kurtosis, but FastICA uses an approximation of negentropy to get better robustness than kurtosis but lower computational demands than negentropy. The measure can be defined as follows:

$$J(y) \approx k[E\{G(y)\} - E\{G(v)\}]^2 \tag{2.11}$$

Here v is a gaussian variable of zero mean and unit variance, k is a positive constant and G is a preferred nonquadratic function. By choosing G to be something that grows slowly, for example $G(y) = log(0.1 + y^2)$, one can make the algorithm not sensitive to outliers.

Finding independent components thus corresponds to finding maxima of $J(y)$, where $y = w^T z$ and $E\{y^2\} = 1$. This could be solved with gradient methods, but FastICA uses a fixed point algorithm that approximates the Newton's method. Steps for the algorithm are outlined below (weight vectors are estimated in parallel):

1. Select randomly initial weight matrix $W = (w_1, w_2, ..., w_n)$ with each w_i having unit norm and decorrelate W.
2. Let $w_i = E\{zg(w_i^T z)\} - E\{g'(w_i^T z)\}w_i$ for all i, where g is the derivative of non-quadratic function G.
3. Decorrelate W: $W = (WW^T)^{-1/2}W$
4. If not converged, go back to 2.

The decorrelation scheme above is in some sense optimal, as it projects W to space of orthogonal matrices in such way that the distance between old and new matrix is minimized. Expected values present in the algorithm are in practice replaced with sample averages. Using this simple, robust and fast procedure of FastICA one can thus obtain the unmixing matrix that can then be used to unmix the hidden source time series from the observed time series:

$$y^i = WVx^i, \tag{2.12}$$

where $y^i = (y_1^i, y_2^i, ..., y_n^i)$ is the ith sample of the n-dimensional source time series and W is the $n \times n$ unmixing matrix, V is the $n \times m$ whitening matrix, and $x^i = (x_1^i, x_2^i, ..., x_m^i)$ is the ith sample of observed m-dimensional time series.

2.5 Fourier-ICA

EEG data is two-dimensional; there is the dimension of sensors, and the dimension of time. ICA can be applied in both dimensions. If one looks for independent spatial maps (i.e values at every sensor), all the time points can be thought as random variables. Then each sensor gives a sample and ICA tries to find the spatial maps that have the most independent distributions. With continuous EEG data it is usually hard to do ICA this way, as the size of the covariance matrix is proportional to the square of number of variables (James and Davies 2008). In Sockeel and Schwartz (2016) this spatial variant of ICA was done to EEG data, but the data was transformed in such a way first that the number of time points went down. Normally in EEG the ICA is applied in its temporal variant. In that case the independent time courses are of interest and sensors are thought as random variables.

This is not all, however. As was described earlier, ICA works by maximizing nongaussianity. It is all good when the "interestingness" positively correlates with nongaussianity. When directly applying temporal variant of ICA to multi-channel EEG data, artifacts such as eye blinks or heartbeats are easily found, as will be seen in the analysis section when preprocessing the data, but it is not as easy to find oscillatory activity. This might happen because sine waves, even when being nongaussian when they are stationary, can become more gaussian-like when their amplitude modulates. This makes nonstationary sine waves hard to find for the ICA algorithm despite being really interesting neuroscientifically. For this reason another variant of ICA was proposed in Hyvärinen et al. (2010), idea being that the data is first be transformed with STFT to time-frequency domain, and ICA is then used on the transformed data. As it turns out, in this domain the previously gaussian-like amplitude-modulating oscillatory activity becomes nongaussian, and can be found with this spectro-temporal variant of ICA. Figure 7 illustrates the changes between distributions when transforming the data. Only the positive side of the distributions is plotted as in the Fourier-ICA case the objective function changes in such a way, that only positive values of distributions are being considered.

Some minor changes are needed in the ICA algorithm so that the TFR's given by the STFT become digestible to it. Two main things to consider are that the data will now be complex valued (as Fourier Transform results in complex values) and that the data is now in three-dimensional (time, frequency, channel) domain while the ICA that was explained in the previous section used data in two-dimensional domain. The second problem can be solved by flattening the time-frequency plane into time-frequency line by concatenating the FT's next to each other. After this the data will be in two-dimensional domain: other dimension is the channels and the other dimension consists of all the frequencies lined up next to each other as many times as there are time points in the TFR.

First problem is a bit more involved. We could use only the length of complex vectors for everything, but it turns out that the it is a good idea to utilize the phase information contained in the FT's, as then for example brain areas that are phase-locked, meaning that they are possibly working together, can be grouped into single independent component instead of having multiple components (Hyvärinen et al. 2010). To allow this, some small changes to

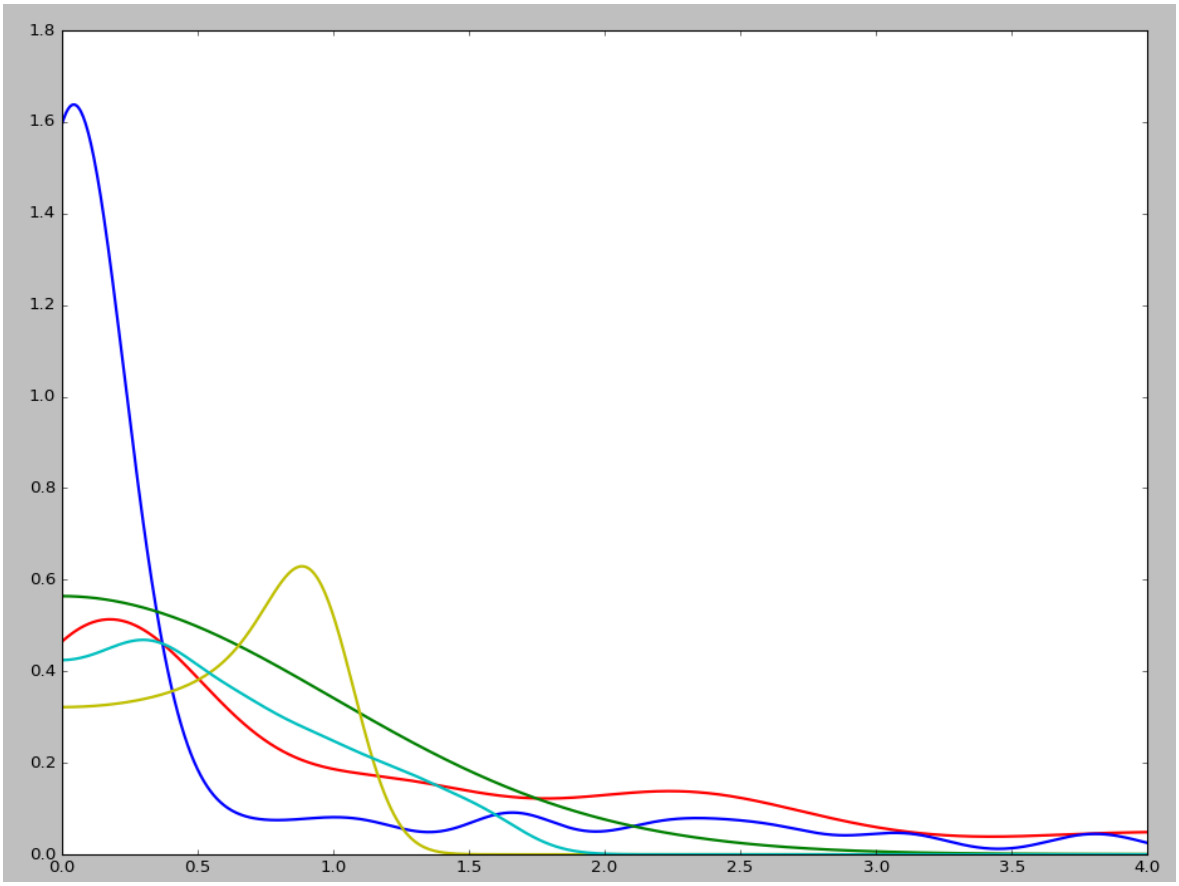

Figure 7. Illustration of gaussianities. The green curve shows a density of a folded gaussian variable. The yellow curve is a density of a sine wave in time domain and the cyan curve is a density of an amplitude-modulated sine wave in time domain. As can be seen that even if a sine wave is clearly nongaussian in time domain, it can become more close to a gaussian when amplitude-modulated. The red curve is a density of a mixture of two sines in concatenated time-frequency-domain and the blue curve shows a density of an amplitude-modulated sine wave in concatenated time-frequency domain (the bumps on the tail result from the amplitude peaks). Densities were estimated with gaussian kernels. This illustrates how a gaussian-like curve in time domain can become more nongaussian in time-frequency domain and gives incentive for such transformation step before using ICA.

the FastICA-algorithm are needed, but otherwise the procedure is the same.

Whitening transformation is done in the same way as before, but the transpose is replaced with conjugate transpose: $V = D^{-1/2}E^H$.

The goal in complex ICA is to maximize the independence of $y = W^H z$, where $y = (y_1, y_2, ...y_n)$ is the source random vector, W is the unknown unmixing matrix, and $z = (z_1, z_2, ..., z_n)$ is the whitened observed random vector. Solving this leads again (with small changes) to a fixed-point algorithm:

1. Select randomly initial weight matrix $W = (w_1, w_2, ..., w_n)$ with each w_i having unit norm and decorrelate W.
2. Let $w_i = E\{z(w_i^H z)^* g(|w^H z|^2)\} - E\{g(|w^H z|^2) + |w^H z|^2 g'(|w^H z|^2)\}w$ for all i, where $*$ denotes the complex conjugate and g is the derivative of a nonquadratic function.
3. Decorrelate W: $W = (WW^H)^{-1/2}W$
4. If not converged, go back to 2.

After the components are found, they can then be split back to the three-dimensional domain for further analysis. Now that we are fully equipped, the analysis section will follow.

3 Analysis

3.1 Task and subjects

The data for this analysis section comes from a study where 30 subjects with varying meditation experience had their electric brain activity recorded with a EEG device during a Focused Attention Meditation task. For the task the subjects were instructed to focus their attention to breathing for 10 minutes with their eyes open. They were also instructed to press a button if they noticed that their mind was off the task, and then to return their attention back to the task. To have some baseline data, also 3 minutes of resting was recorded, of which 90 seconds were eyes open.

Recordings were made with a Geodesic EEG system. The EEG Net had 128 channels spread over the scalp and they all recorded voltages at the frequency of 1000 hertz. This data was then saved to hard drive for later analysis. We will proceed to take a look at the data now, starting with preprocessing.

3.2 Preprocessing

EEG signals can be decomposed to two parts. In addition to the interesting part that could possibly tell us something about the underlying brain mechanisms that we are interested in, there is also a part that contains things that we are not interested in and make the analysis harder. Heartbeats and eye blinks are examples of activity that are usually cleaned out from the data as to remove the chance of them masking the real effects that are being looked for. Measurement devices or the environment can also cause artifacts; electrodes can sometimes be momentarily deattached from the scalp, or the electric line frequency might get recorded. It is always better to try to minimize the noise as early as possible (Luck 2005), but sometimes some disturbances are unavoidable, and EEG recordings could also be seen as inherently noisy: there is the part of the data that can be used to answer the research problem, and the rest of the data is noise even if it has its roots in the brain.

Next ICA in its temporal variant is used for the removal of artifacts, as the artifacts tend to

have very distinct time courses. ICA will not do the separation, though, as it will just give out the independent components, and then one must use knowledge of artifacts to identify which components are part of the interesting activity, and which components should be left out. After the components to be removed have been selected, their contribution to original recordings can be calculated (using the mixing matrix and then cross-correlating to find the scale) and subtracted away.

Essentially following code utilizing MNE-python (Gramfort et al. 2013) was run for every subject:

```
# open the data
raw = mne.io.Raw('data.fif')

# transform to source space
ica = mne.preprocessing.ICA(n_components=0.95, method='fastica')
sources = ica.get_sources(raw)

# plot for visual inspection
ica.plot_components(show=False)
sources.plot()

# select components to remove by inspecting visually
indices = ...

# project artifacts out
ica.apply(raw, exclude=indices)
```

Figure 8 shows independent components for resting data of one subject in the time domain. Figure 9 shows the same components projected to sensor topography, meaning that contributions of each independent component to the original sensor-level time series are calculated and these contributions are plotted over a head figure, thus giving some information about the locations of the components. Sometimes an component consisting of an artifact could be identified by just looking at the time course, like in the case of eye blinks, but sometimes it

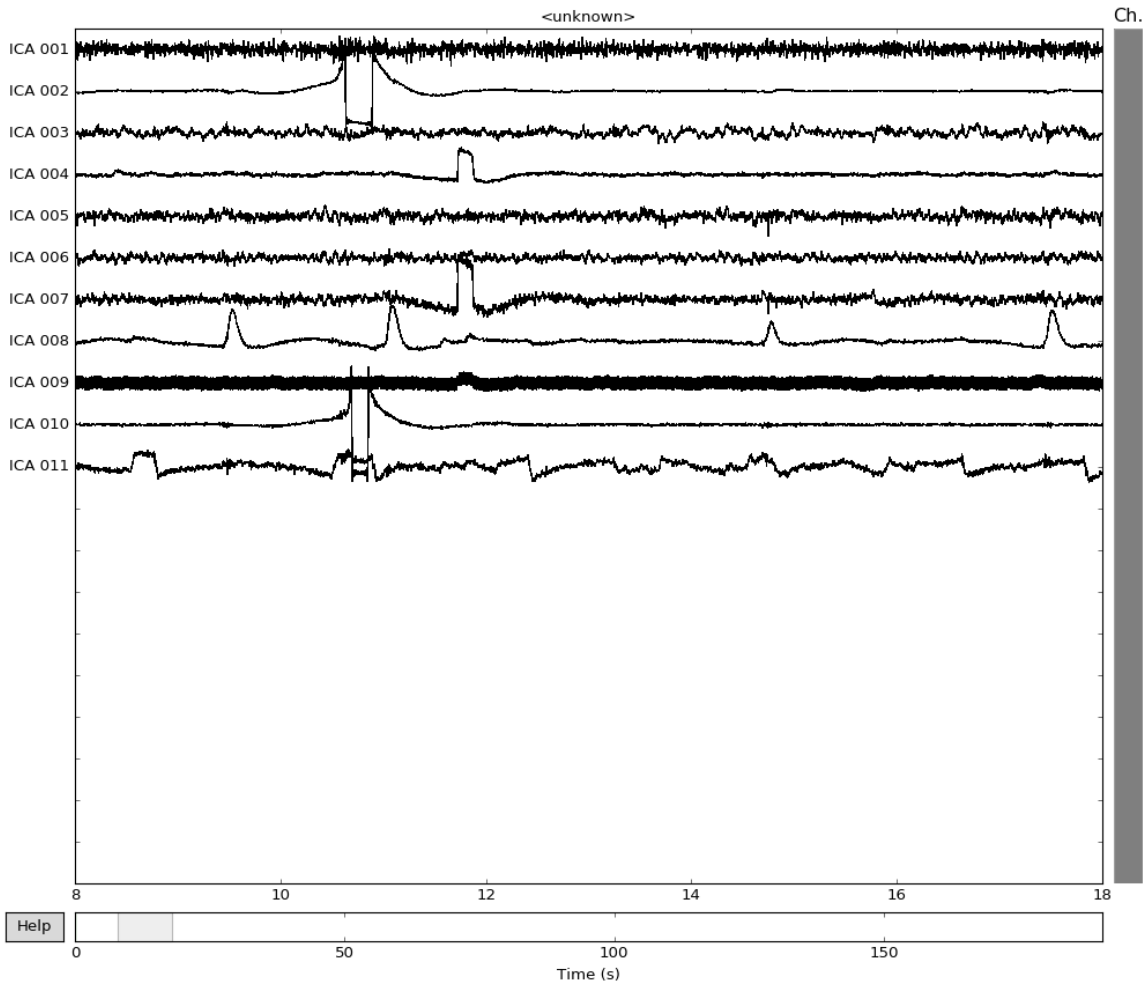

Figure 8. ICA components in time domain. FastICA algorithm implemented in MNE-python was able to find 11 independent time courses given that 95% of the variance was to be taken into account. As can be seen, many of these are artifacts.

helps to also look at the topographies. Activity stemming from the brain normally spreads to relatively large areas vanishing quietly, because electricity spreads to every direction in a conductor, and thus some components, like 1st, 2nd, 4th, 9th and 10th, are easily identified as artifacts by just looking at their isolated nature. Eye blinks can be seen in the eighth component.

This procedure was then used to clean data for all of the subjects. More noise could be removed by simply band-pass filtering out the uninteresting frequencies, but the analysis methods used later on will restrict themselves to interesting frequencies on their own. During

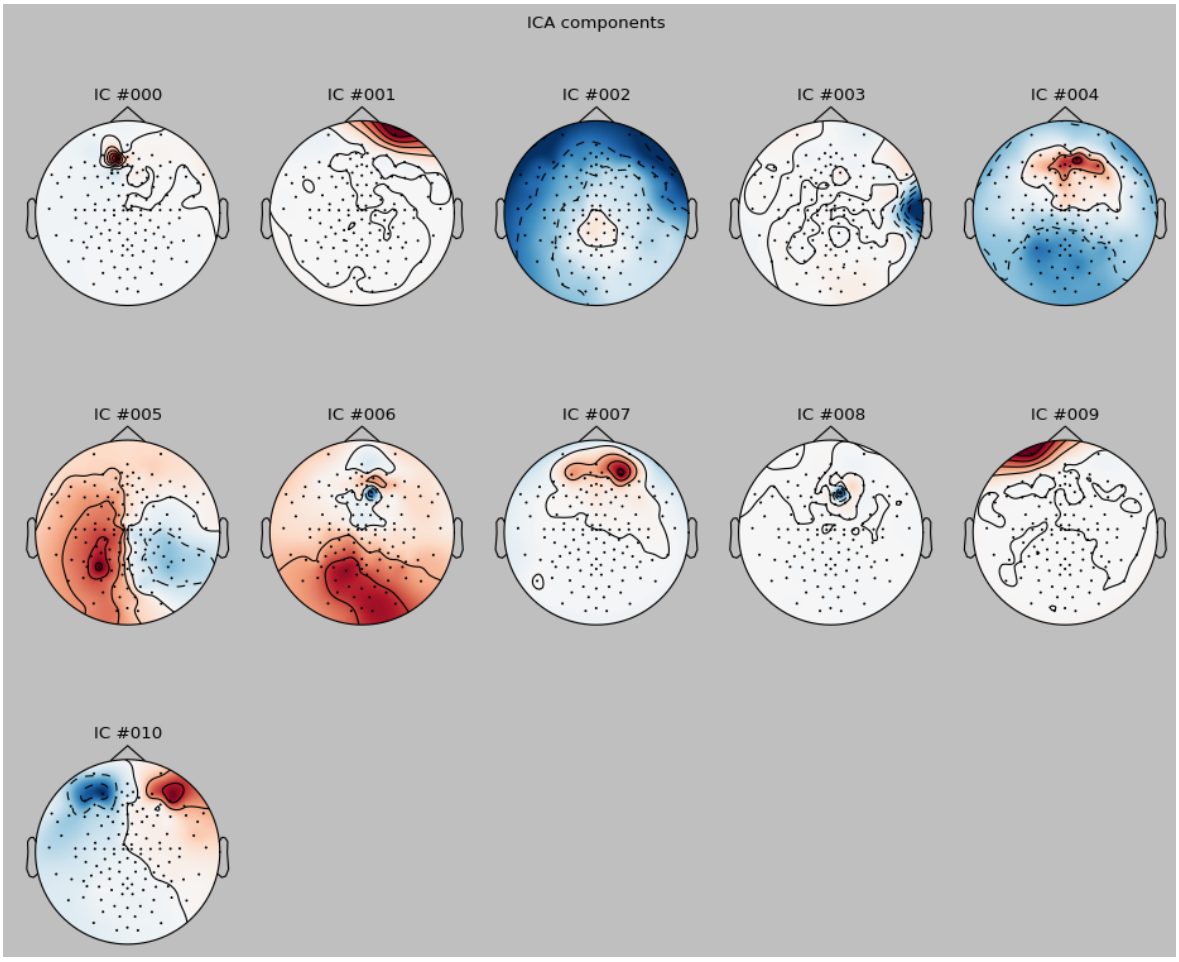

Figure 9. ICA components projected to head topography. Components that were present in 8 are shown from a different point of view: how much they contribute to different locations in the sensor topography.

the preprocessing step it was also found out that button press data was practically unusable for many of the subjects, either having too many buttons presses (which makes sense; it is hard to stay focused) so that any meaningful division to meditative state and mind wandering state was impossible, or too few, which also results in small amount of data. This will make it hard to do any analysis based on the button presses, but we will try it nevertheless.

3.3 Experiments

3.3.1 Comparing alpha power spectrums in resting and mindfulness state using sensor averages

First a comparison was done between alpha power (the peak around 10 hertz in the power spectrum) in two different states. Data for the first state comes from the meditation recording, where subject is focusing her attention to breathing. The data for the second state comes from the resting state recording, where the subject had her eyes open and was not instructed to do anything. Idea is that even though meditation can be also seen as a relaxing method, mind is still focused. This is in contrast with the resting state, where mind might be engaged in many thought patterns, and this difference in activity between the states might be visible in the spectrums.

To motivate the use of Fourier-ICA the comparison was first made without it. The Welch's method was used for each channel in both recordings to get power spectrums for the two conditions. In the method, a window of 2048 samples and a step size of 1024 samples was used. An effective band-pass filtering from 3 hertz to 17 hertz was achieved by dropping off the frequencies outside this region. The wide band was selected so that the inter-individual alpha frequency variability could be taken into consideration (Haegens et al. 2014). After getting the spectrums, the data was then spatially filtered to posterior site and averaged together thus giving one spectrum for each condition.

Two statistics, that were thought to reflect the strength of the alpha oscillations, were chosen for the statistical comparison: the value at the highest peak of the spectrum, and area under the curve in 7 hertz to 13 hertz region. When applying the paired t-test on this data, the means did not differ significantly in the resting and meditation conditions. P-value of paired

t-test for maximum peak statistic was 0.17, and p-value of paired t-test for the area statistic was 0.41. In both cases the calculation was done over the same 18 subjects that were used in the following sections, to allow better comparability. For all the 30 subjects alpha values were 0.18 for the peak value difference and 0.36 for the area difference.

3.3.2 Comparing alpha power spectrums in resting and mindfulness state using independent components

Even though the previous test was not able to show the difference, there is a reason to believe that the problem might be in the method and not in the nonexistence of the difference. Thus the theory about the Independent Component Analysis on Time-Frequency domain, that was presented in the theory sections of this thesis, is put to use next. The same data for the resting and mindfulness conditions is used as in the previous section, but instead of averaging over sensors, the features to be compared will be separated from the rest of the activity by ICA.

As was described in the theory sections, to have ICA prefer oscillatory activity the data is first transformed to time-frequency plane with STFT. For this a window of 2048 samples and a step size of 1024 samples were used. Then the result was band-pass filtered to band from 3 hertz to 17 hertz to have a wide but not unnecessarily wide band centered at 10 hertz. After that, time-frequency plane was flattened by concatenating spectrums of all time points next to each other. The data was whitened so that its covariance matrix equals to identity, and the FastICA-algorithm was applied to it, giving out the independent components. Number of components searched for was selected beforehand to 15. Components were then transformed back to the time-frequency plane by reversing the concatenation step. Using the calculated mixing matrix it is also possible to get the head topographies in the same way that it was possible in the preprocessing step. All of this was implemented in a python class, that was created for the purpose of using Fourier-ICA to EEG and MEG data, and this code can be found in the appendix. The code is based on the Fourier-ICA implementation found in the Spectrospatial decoding toolbox (Kauppi et al. 2013).

In the figures 10, 11 and 12 are the 5 highest scored Fourier-ICA components, their head topographies and spectrums (components averaged over time) for one of the subjects. Fifth

component looks like an artifact, but others seem to be some kind of oscillatory components peaking near the 10 hertz frequency. So it seems the method is able to isolate amplitude-modulating oscillatory components just as was assumed. It seems also that different components are spatially separated, and as the Fourier-ICA should be able to combine components with different but locked phase, it is possible that these alpha components are related to different brain mechanisms. If it is so, it is not hard to imagine that without unmixing the observed TFR's some canceling out might take place in the sensor level.

This procedure was then applied to all of the subjects resulting in a set of 15 components for each subject. Because of noise and maybe differences in the brain activity of the subjects, these sets of components can look very different. For some of them there might be a lot of oscillatory components from multiple alpha components to even beta components. For others there might be one good alpha component or even no alpha component at all. Luckily the posterior (i.e at the back of the head) alpha component was so prominent that it was included in most of the component sets. Next a template of posterior alpha topography and spectrum were created and these were used to select the most similar component from each subject. All such components (total of 18) were than collected together and are shown in the figures 13, 14, 15 and 16. It is possible that the components do not all represent the same brain function, as the selection was only rudimentary. Most often the spatial distance to other components was large, however, and it is known that there are strong alpha oscillations in that region, so there is reason to believe that most of the components in fact do represent the same function. Also the varying alpha peak location visible in the spectrums instead of being an artifact might actually tell something about the expertise of meditators (Tang, Holzel, and Posner 2015).

To allow robust comparison of alpha power between focused attention state and the resting state, the components were calculated for data where the resting state recording was concatenated to the focused attention recording. This made it possible to select components only once for each subject, thus reducing the variability coming from calculating and choosing components multiple times. Also it is good to note that by selecting the components based on only the spectrums and topographies, the selection was independent of the states of focused attention and resting, and thus was not biased.

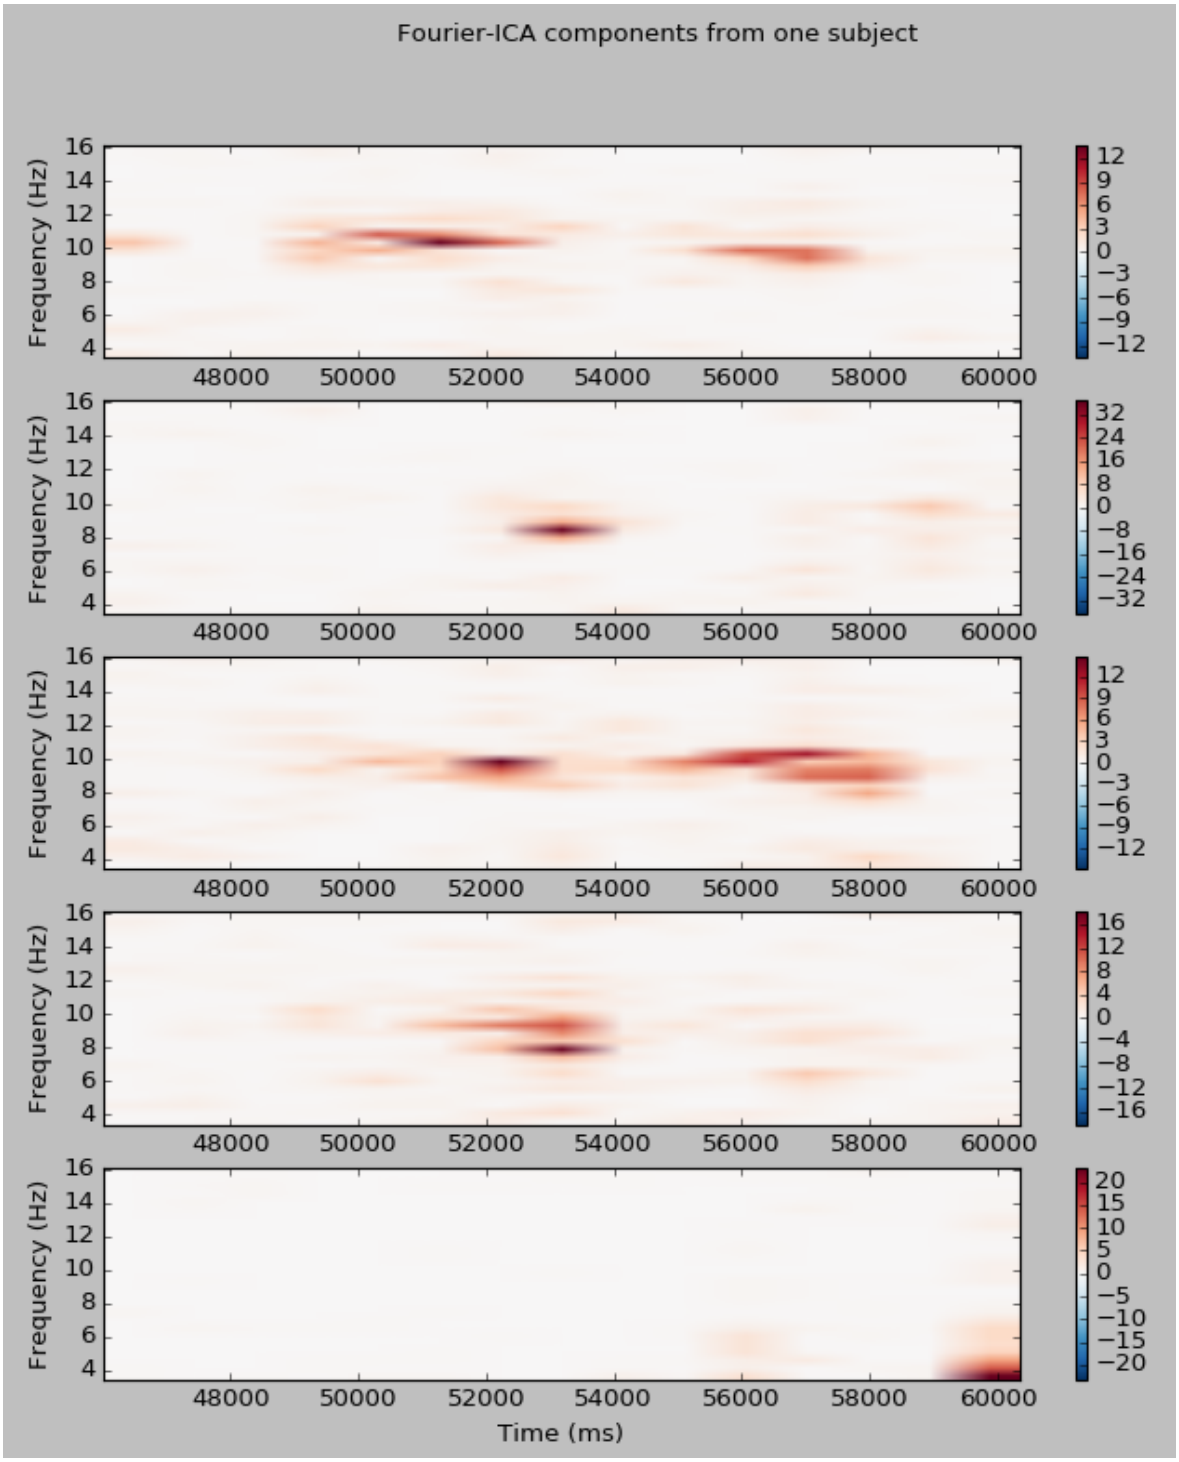

Figure 10. Fourier-ICA components from one subject. These TFR's represent the independent components that Fourier-ICA can unmix from the sensor-level TFR's. These particular components came from the first subject but are not unlike the components from other subjects.

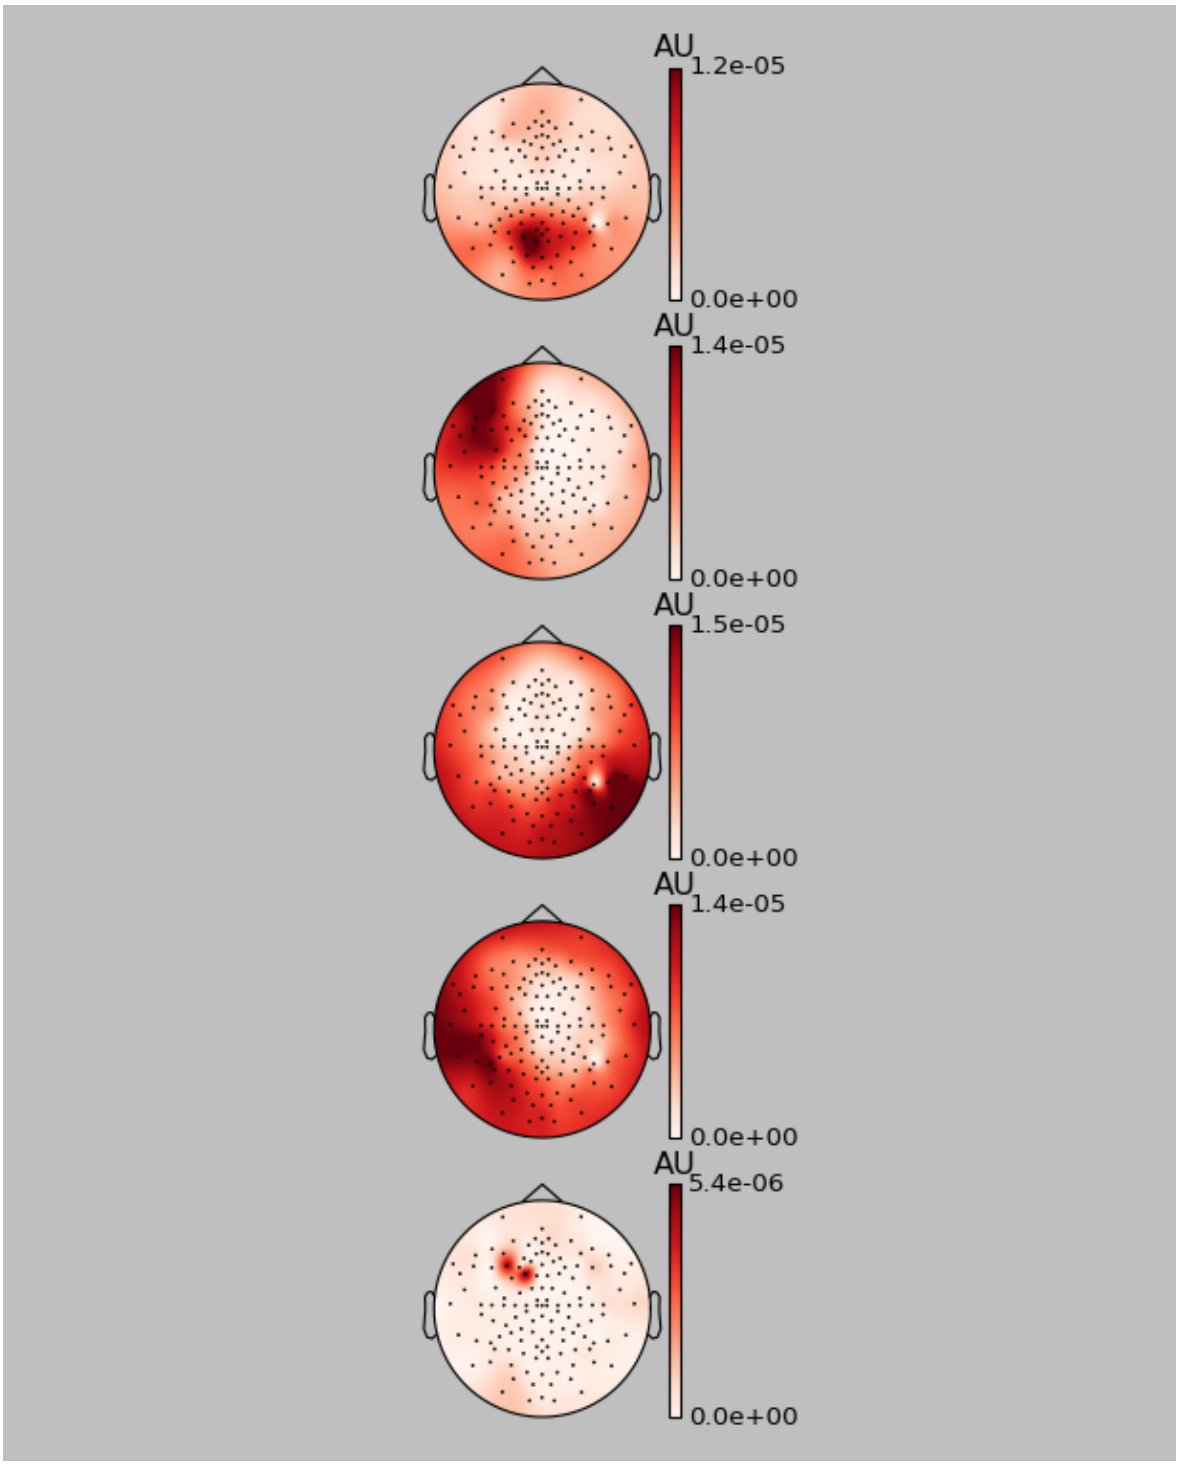

Figure 11. Topographies for Fourier-ICA components from one subject. These head topography plots represent the independent components unmixed by Fourier-ICA. These topographies are from the same subject as the TFR's in 10.

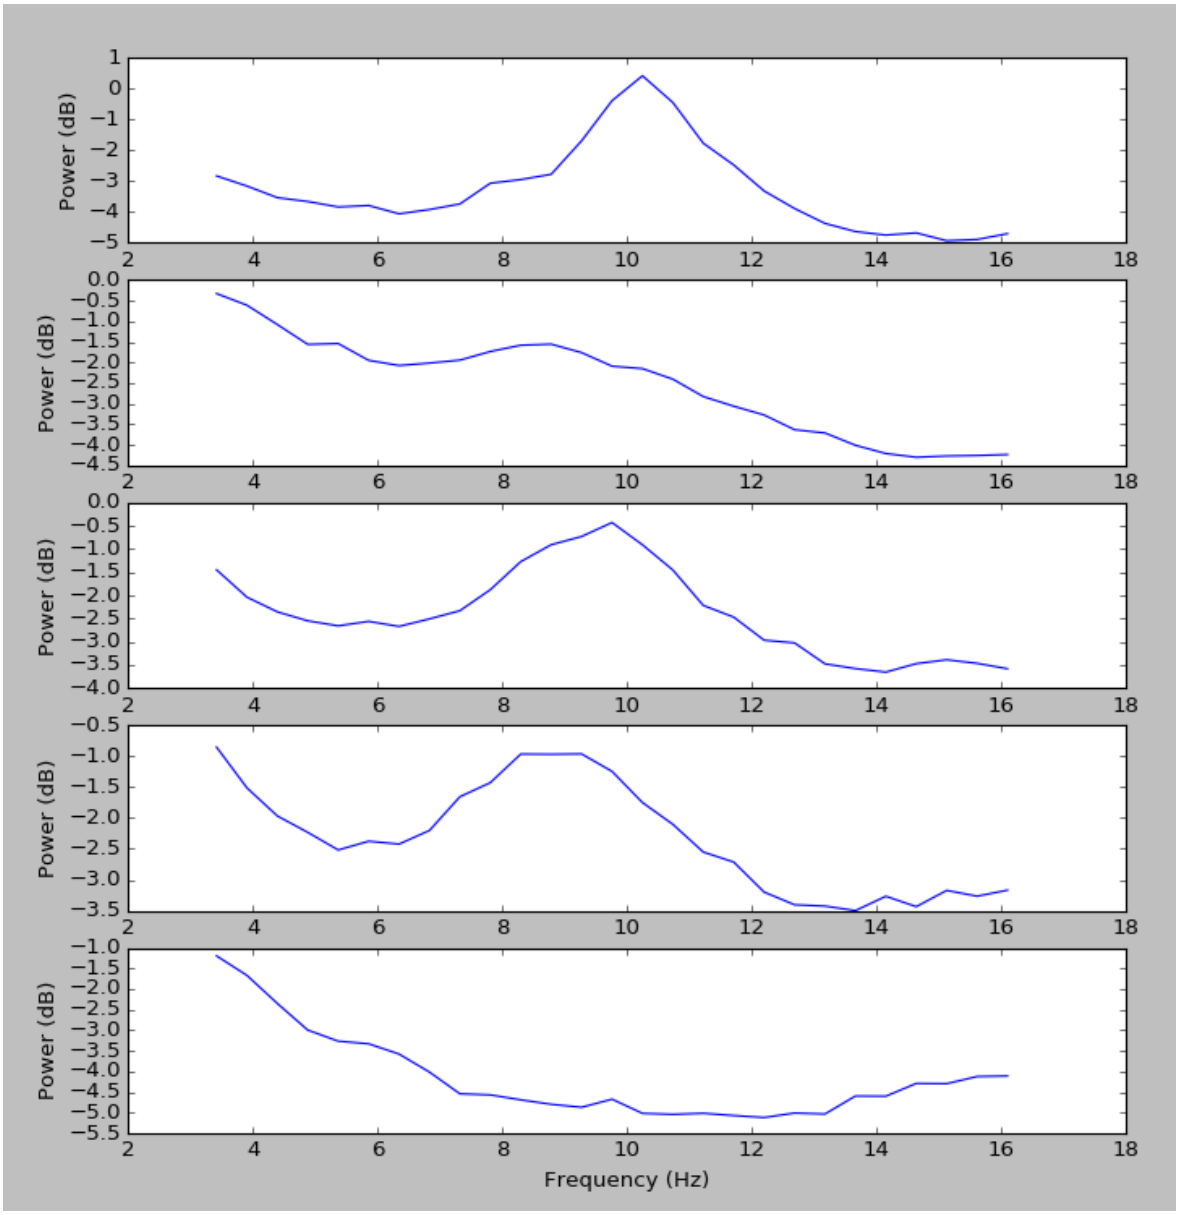

Figure 12. Spectrums for Fourier-ICA components from one subject. These spectrums were created by collapsing the time domain in the TFR's in 10. Looking at the spectrum gives a good overview, but also summarizes the component in a way that does not depend on the time structure.

Figure 13. Topographies for selected ICA components 1.

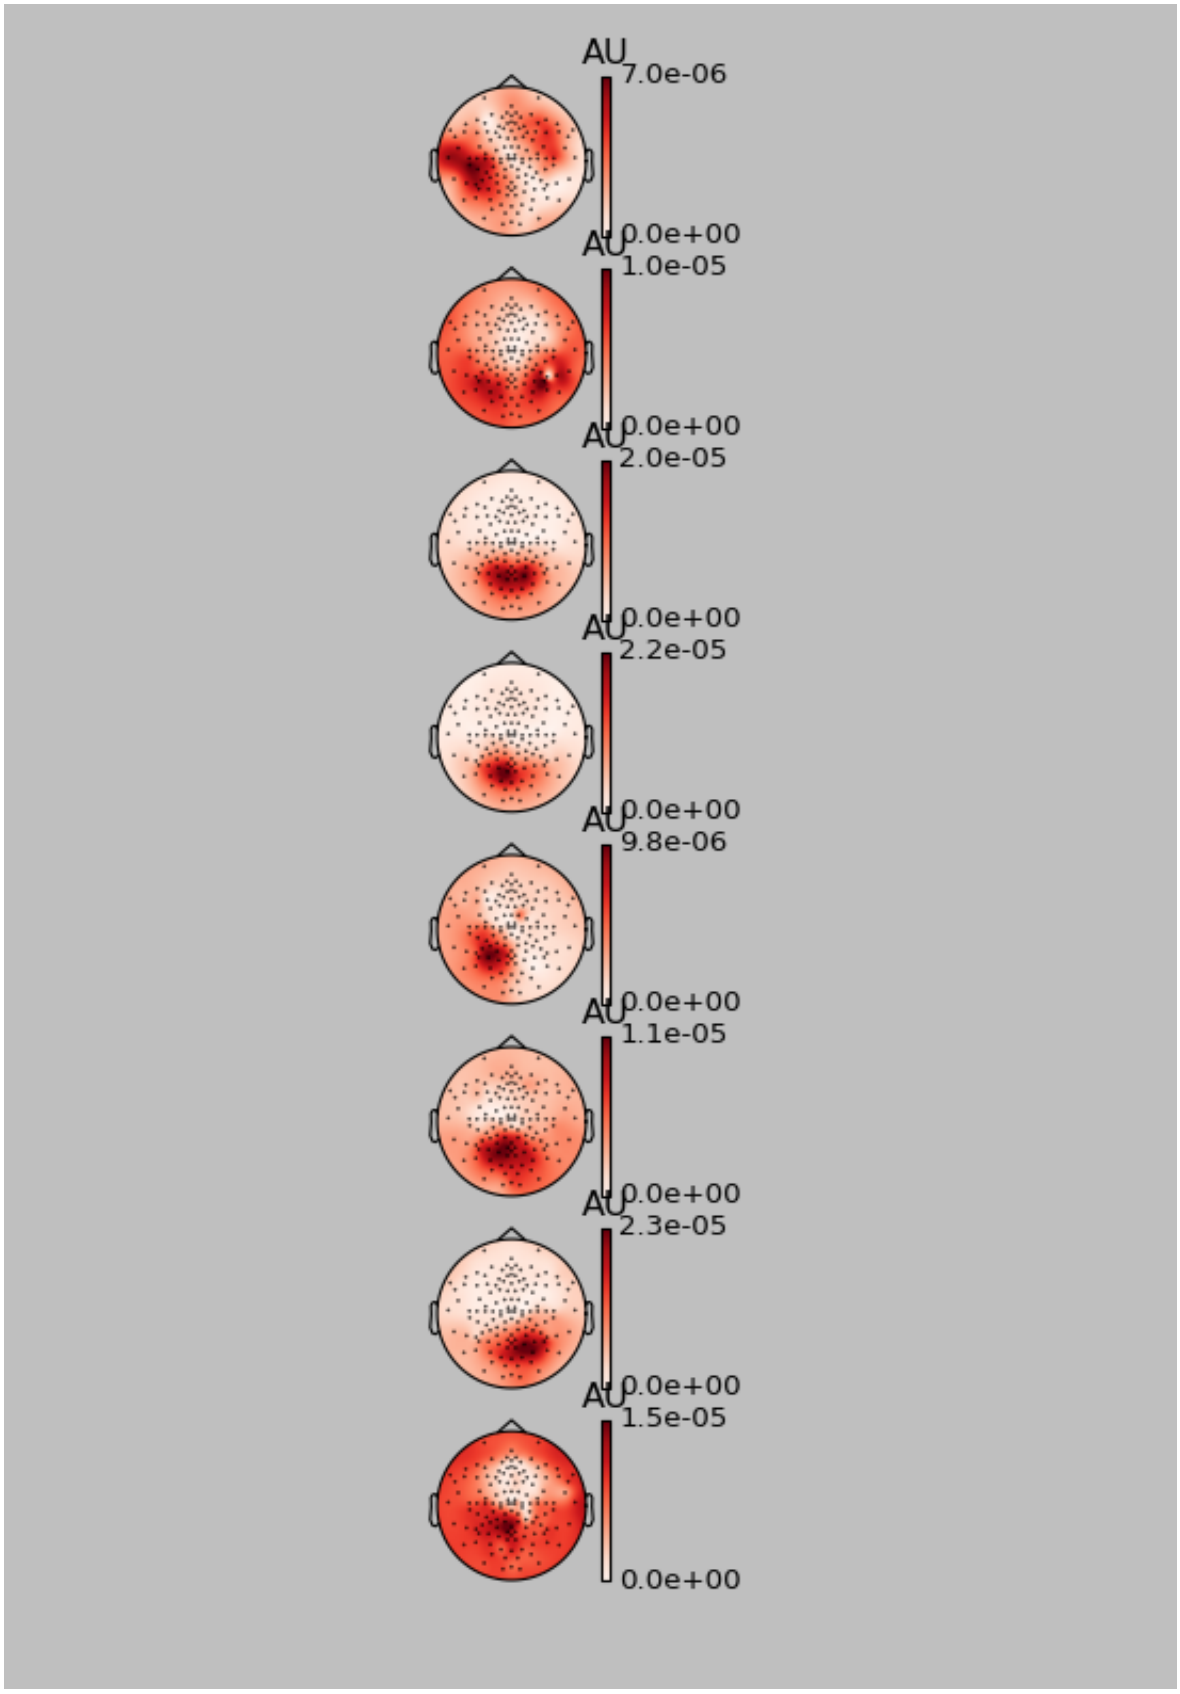

Figure 14. Topographies for selected ICA components 2.

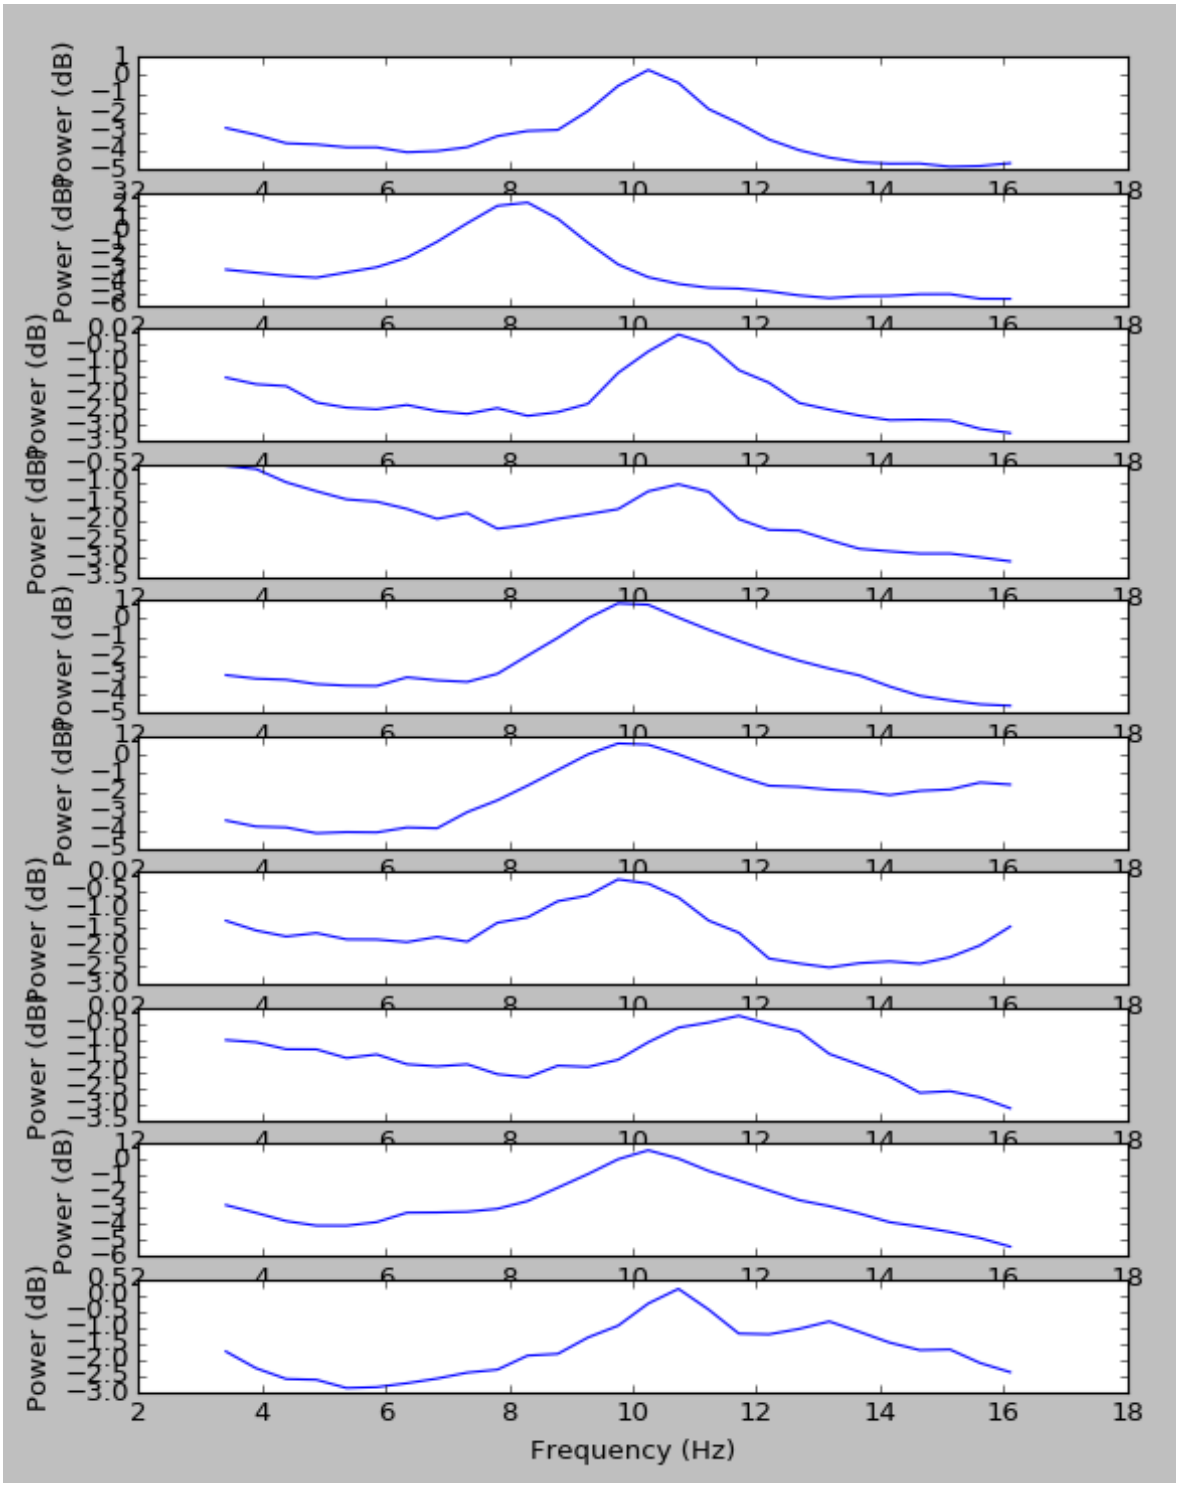

Figure 15. Spectrums for selected ICA components 1.

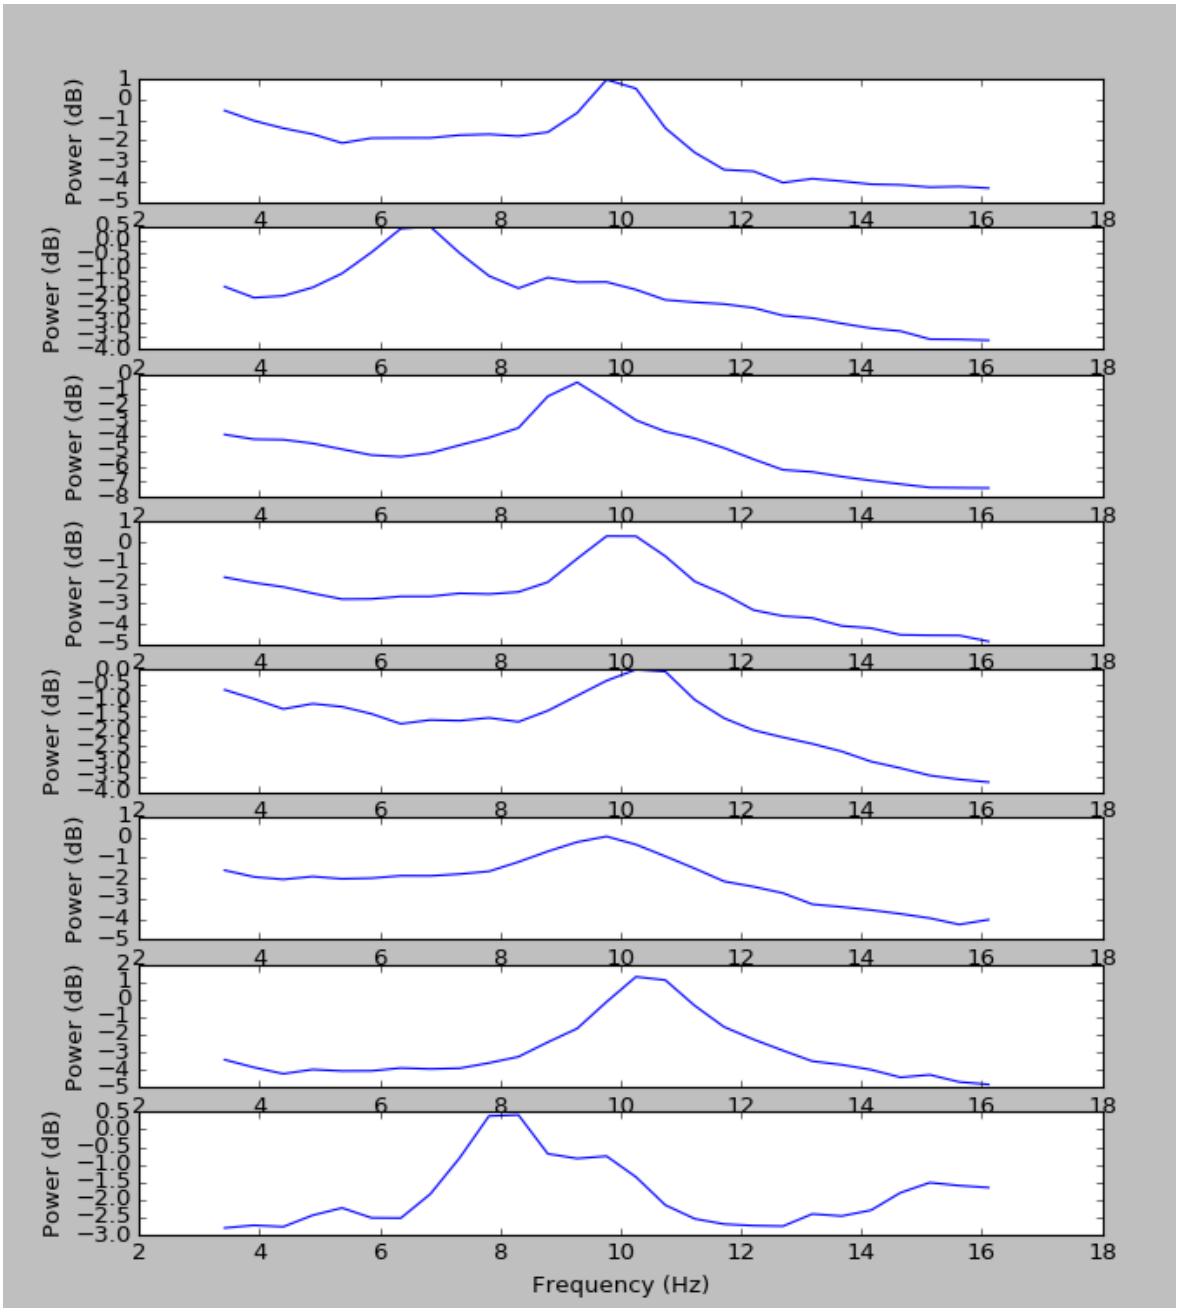

Figure 16. Spectrums for selected ICA components 2.

Spectrums for the components of resting state and focused attention averaged over all subjects can be seen in figure 17. As can be seen, it looks like the states actually do have differently behaving alpha oscillations. To be sure, statistical significance was tested. Two statistics were again chosen, peak value and area under the curve in 7-13 hertz range. Because the values are paired, i.e there are two state values for each subject, a paired t-test was selected. 18 subjects should be enough to compensate for the possible nongaussiness thus validating the use of the test. For the peak value difference a p-value of 0.0012 was got and a p-value of 0.0008 came from the area difference. There is thus reason to believe that posterior alpha power really tends to be higher when focus is concentrated on breathing, compared to the assumed mind wandering of resting state. This might be expected as by the inhibition hypothesis the brain area not responsible for doing the task should be inhibited when attention is constantly elsewhere.

3.3.3 Exploring patterns around button presses by averaging TSE's of independent components

An other way to approach mind wandering is to see if it could be detected inside a meditation task. The alpha component selection that was seen in the previous sections was used here too, but the components were calculated again because more accuracy in the time domain was needed to get a better look at the button presses. STFT was now calculated with a window of 1024 samples and a step size of 512 samples, giving 0.5s resolution in time domain.

These TFR-components were then transformed to TSE-representations. The result was similar to what was explained in the theory sections, but the starting point was a bit different, as now a TFR was transformed instead of a time series. This was done by selecting a band of frequencies in the TFR and averaging the amplitudes over those frequencies. As the alpha peak location differed in frequency from subject to subject, the peak was first located on every component, and then a small band of 3 hertz was selected centered to it. Averaging over that band gave the TSE, i.e. in this case alpha power as a function of time.

These TSE's were then split into segments and the segments around button presses were kept. A problem here was that even though there were 18 subjects meditating for 10 minutes, there

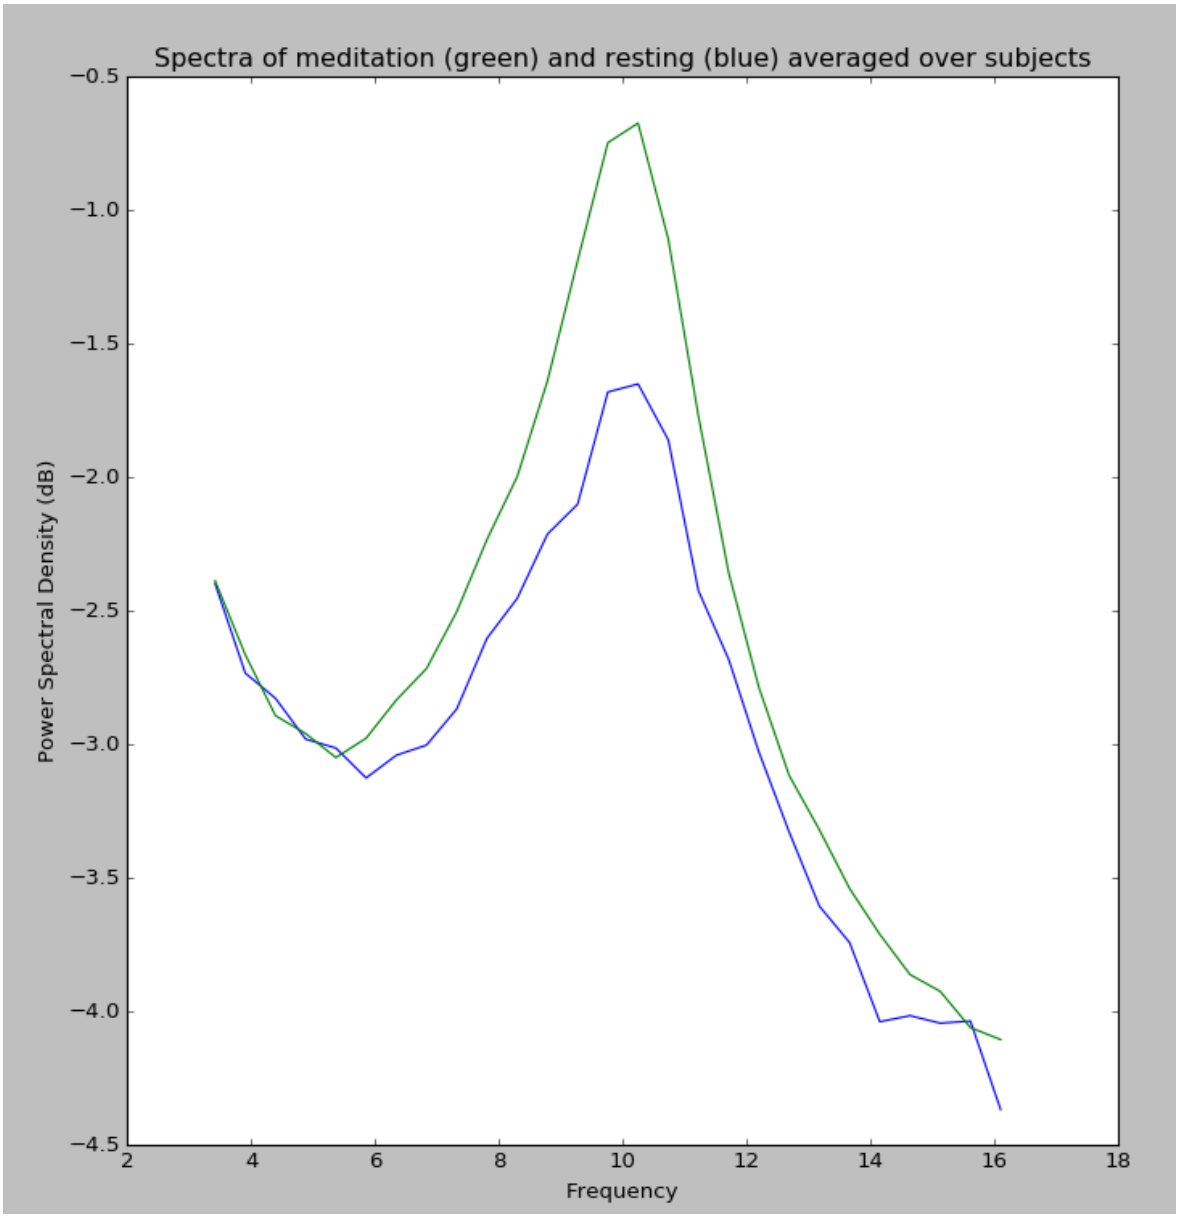

Figure 17. Spectra of meditation (green) and resting (blue) averaged over subjects.

were not that many proper button presses, meaning that many of the button presses were too close to each other, or too close to beginning or end of the recording, and certainly the button presses were not distributed evenly among subjects. To have as many as button presses as possible (total count was 66), the segment taken around a button press was only 15 seconds long: 10 seconds before the button press and 5 seconds after. This time window should be just about enough, though, if one trusts the fMRI study that was introduced in the section about Focused Attention Meditation. There it was said that the mind wandering event during meditation session could be divided into four 3 second phases of different metacognitive states. Thinking that the mind wandering might begin a while earlier than at the button press, the segment was selected to extend more to the past.

All these TSE segments from the button presses were then averaged together. The result can be seen in the figure 18. Even though the picture is not pretty, there are some things to note. It would seem that the alpha power dips at the moment of button press. If the inhibition hypothesis holds, this would mean that the area, where the posterior alpha component lies, gets activated during the button press. Assuming the component lies in the visual cortex, the subject might try to visualize the movement towards the button. Other thing to note is that right before the button press the alpha power is lower than what it is before that or after the button press. This possible difference might reflect the same property of alpha power as was seen in the previous section.

Given the uneven distribution and low amount of button presses, statistical analysis is omitted here. One main problem here also is the averaging process that can lose interesting information if the segments are not properly overlayed, and thus it would help if the moment of change from focused state to wandering mind was known beforehand. Nevertheless, even with these problems, it would seem that Fourier-ICA can enhance the analysis considerably for these kinds of tasks too.

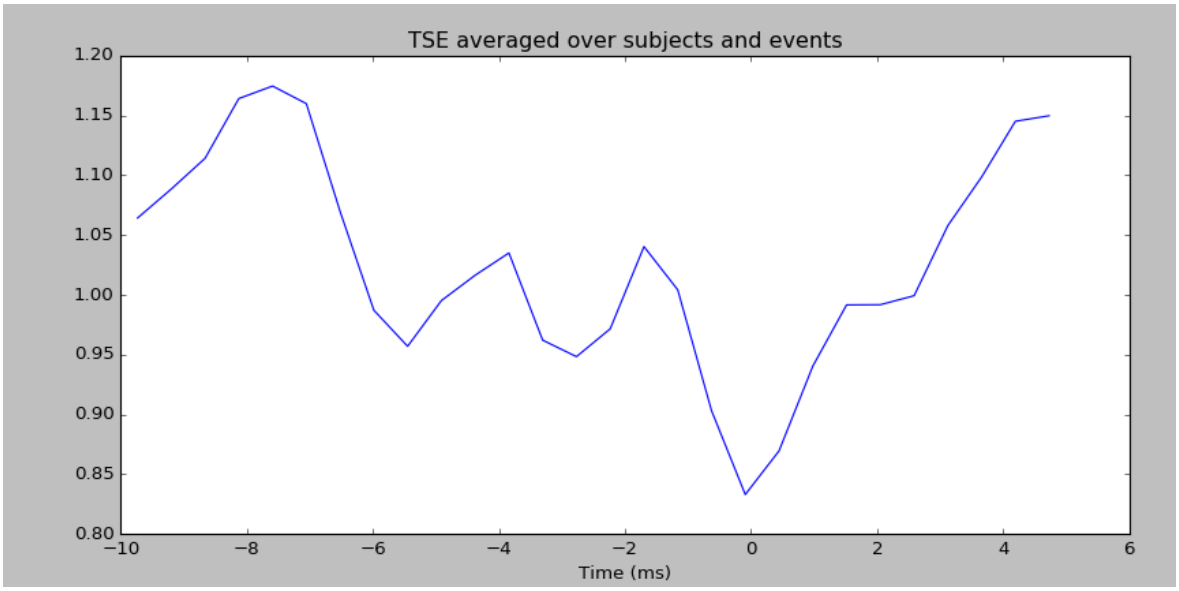

Figure 18. Grand average of TSE segments. Button press is located at the time of 0.

4 Discussion

This thesis has presented a way of analyzing continuous EEG data by combining Independent Component Analysis with Short-Time Fourier Transforms. Using this methodology, it was possible to demonstrate differences in oscillatory activity between different naturalistic states, which is something that the traditional EEG analysis methods including Event-Related Potentials cannot do, and in which the basic ICA alone does not excel either.

A significant difference was found in the behavior of posterior alpha activity between the states of meditation and resting, alpha power being higher in the state of meditation. This can be explained as inhibition of vision-related brain activity during the meditation task. Moreover it was shown that even if the task of extracting this difference with normal approach of averaging spectrums over sensors was difficult, the difference was easily extracted out with Fourier-ICA method. As to speculate on why did the effect not show with the averaging method, it could be that there are multiple differently behaving 10 hertz oscillations mixed in the posterior sensor site, and thus can be canceling each other out, or that statistical independence as the separating guideline was able to also separate out noise that was present in the sensors and was affecting the spectrum around 10 hertz too.

The voluntary button presses by the subjects during the meditation task, indicating the wandering of mind, proved out to be more difficult to analyze, problem mainly being in the voluntariness that results in uneven distributions of button press counts between the subjects and unusable data segments, and in the difficulty to determine the starting point of mind wandering. The difficulty to determine starting point of mind wandering results in losing relevant information when the mind wandering segments are averaged together. An interesting alternative to voluntariness would be to use experience sampling, i.e to ask the subject about how well she is doing during the task on semi-regular intervals. This would fix the problem of unevenness, although all of the samples could not be used as indicators of mind wandering and the starting points could still not be determined. Nevertheless the analysis with Fourier-ICA was able to hint about patterns emerging from the button presses, that could maybe be studied in the future. A main improvement would be to figure out a way to determine the points of change between focused state and the mind wandering.

In both Fourier-ICA tasks there was also a choice of selecting only one independent oscillatory component from each of the subjects. This was mainly for the simplicity but also for the noisiness of the data. Some of the subjects did show several oscillatory components but for some they remained hidden. It would be interesting to conduct this kind of study with data acquired by a Magnetoencephalogram (MEG) device, as then the improved signal-to-noise ratio could allow better discovery of the independent components. Combining information from multiple independent components could be a nice step forward. Automatic classifier utilizing Fourier-ICA that can combine information from multiple sources was actually implemented in Kauppi et al. (2013).

In this thesis only comparisons in the group-level were made, but given the amount of improvement in the feature extraction that Fourier-ICA can give in comparison to more simple methods like the sensor averaging, it could be possible to target even individual differences, and this should be looked into in the future studies.

The work done in this thesis tries to make its contribution for the analysis of oscillatory activity in the brain, bringing us a little bit closer to a more fundamental understanding of ourselves.

Bibliography

Buzsaki, Gyorgy. 2011. *Rhythms of the Brain.* Oxford University Press.

Cahn, BR., and J. Polich. 2006. "Meditation states and traits: EEG, ERP, and neuroimaging studies". *Psychologial Bulletin* 132.

Cooley, James W., and John W. Tukey. 1965. "An algorithm for machine calculation of complex Fourier series". *Mathematics of Computation* 19:297–301.

Gramfort, A., M. Luessi, E. Larson, D. Engemann, D. Strohmeier, C. Brodbeck, R. Goj, et al. 2013. "MEG and EEG data analysis with MNE-python". *Frontiers in Neuroscience* 7.

Haegens, Saskia, Helena Cousijn, George Wallis, Paul J. Harrison, and Anna C. Nobre. 2014. "Inter- and intra-individual variability in alpha peak frequency". *NeuroImage* 92:46–55.

Hari, R., R. Salmelin, J.P Mäkelä, S. Salenius, and M. Helle. 1997. "Magnetoencephalographic cortical rhythms". *International Journal of Psychophysiology* 26:51–62.

Hasenkamp, Wendy, Christine D. Wilson-Mendenhall, Erica Duncan, and Lawrence W. Barsalou. 2012. "Mind wandering and attention during focused meditation: A fine-grained temporal analysis of fluctuating cognitive states". *NeuroImage* 59:750–760.

Hyvärinen, Aapo, Juha Karhunen, and Erkki Oja. 2001. *Independent Component Analysis.* John Wiley & Sons, INC.

Hyvärinen, Aapo, Pavan Ramkumar, Lauri Parkkonen, and Riitta Hari. 2010. "Independent component analysis of short-time Fourier transforms for spontaneous EEG/MEG analysis". *NeuroImage* 49:257–271.

III, Julius O. Smith. 2007. *Introduction to Digital Filters.* W3K Publishing.

James, CJ., and ME. Davies. 2008. "Contrasting spatial, temporal and spatio-temporal ICA applied to ictal EEG recordings." *Proc. 30th IEEE EMBS Annual Int. Conf.*

Kauppi, J-P., L. Parkkonen, R. Hari, and A. Hyvärinen. 2013. "Decoding magnetoencephalographic rhythmic activity using spectrospatial information". *NeuroImage* 83.

Klimesch, Wolfgang, Paul Sauseng, and Simon Hanslmayr. 2007. “EEG alpha oscillations: The inhibition-timing hypothesis”. *Brain Research Reviews* 53:63–68.

Luck, Steven J. 2005. *An Introduction to the Event-Related Potential Technique.* The MIT Press.

Saggar, Manish. 2011. “Computational Analysis of Meditation”.

Sockeel, S., and D. Schwartz. 2016. “Large-Scale Functional Networks Identified from Resting-State EEG Using Spatial ICA”. *PLoS One.*

Stoica, Petre, and Randolph L. Moses. 2005. *Spectral Analysis of Signals.* Pearson Prentice Hall.

Tang, Yi-Yuan, Britta Holzel, and Michael Posner. 2015. “The neuroscience of mindfulness meditation”. *Nature Reviews Neuroscience* 16.

Teplan, M. 2002. “Fundamentals of EEG measurement”. *Measurement Science Review* 2.

Appendices

A Fourier-ICA implementation

```
# Authors: Erkka Heinila <erkka.heinila@jyu.fi>
#
# License: BSD (3-clause)

import sys

import numpy as np
import mne

from mne.time_frequency.stft import stft
from mne.time_frequency.stft import stftfreq

from scipy.linalg import sqrtm
from numpy.linalg import inv


class FourierICA(object):

    """M/EEG signal decomposition using STFT and Independent
       ↪ Component
    Analysis (ICA)

    This object can be used to explore and filter interesting
       ↪ components
    found by ica used to short-time fourier transformed time series
       ↪  data.

    """
```

```
def __init__(self, wsize, n_components, tstep=None,
             conveps=None, maxiter=None, zerotolerance=None,
             lpass=None, hpass=None, sfreq=None):
    self.wsize = wsize
    self.n_components = n_components
    self.tstep = tstep
    self.conveps = conveps
    self.maxiter = maxiter
    self.zerotolerance = zerotolerance
    self.lpass = lpass
    self.hpass = hpass
    self.sfreq = sfreq

def fit(self, data):
    """ Fit data
    """

    print "First␣do␣stft"
    stft_ = stft(data, self.wsize, self.tstep)

    # bandpass filter
    if self.sfreq:
        freqs = stftfreq(self.wsize, self.sfreq)

        hpass, lpass = 0, len(freqs)
        if self.hpass:
            hpass = min(np.where(freqs >= self.hpass)[0])
        if self.lpass:
            lpass = max(np.where(freqs <= self.lpass)[0])

        self._freqs = freqs[hpass:lpass]
```

```
            stft_ = stft_[:, hpass:lpass, :]

        # store shape to retrieve it later
        self._stft_shape = stft_.shape

        # concatenate data
        data2d = self._concat(stft_)

        print "Whiten data"
        dewhitening, whitened = self._whiten(data2d)

        print "Do ICA"
        mixing_, ic_ = self._fastica(whitened)

        # sort according to objective value
        objectives = []
        for i in range(ic_.shape[0]):
            component = ic_[i, :]
            g_ = np.log(1 + np.abs(component)**2)
            objectives.append(np.mean(g_))
        indices = np.argsort(objectives)[::-1]

        sorted_ic = ic_[np.argsort(objectives), :]
        sorted_mixing = mixing_[:, np.argsort(objectives)]

        # store for retrieving
        self._mixing = sorted_mixing
        self._dewhitening = dewhitening
        self._source_stft = self._split(sorted_ic)

    @property
    def source_stft(self):
        return self._source_stft
```

```
@property
def freqs(self):
    return self._freqs

def component_in_sensor_space(self, idx):
    """
    """
    # get concatenated source stft
    data = self._concat(self._source_stft)

    # zero out other components
    data[:idx, :] = 0
    data[idx+1:, :] = 0

    # use mixing matrix to get to whitened sensor space
    data = np.dot(self._mixing, data)

    # dewhiten
    data = np.dot(self._dewhitening, data)

    # add the mean
    data += self._mean[:, np.newaxis]

    # split again and return
    return self._split(data)

def _fastica(self, data):
    """
    Complex fastica depicted from
    (Bingham and Hyvarinen, 2000)
    """
```

```
        if self.maxiter:
            maxiter = self.maxiter
        else:
            maxiter = max(200 * self.n_components, 2000)

        if self.converps:
            converps = self.converps
        else:
            converps = 1e-13

        x = data

        def sym_decorrelation(w_):
            return np.dot(w_, sqrtm(inv(np.dot(np.conj(w_.T), w_)))
                ↪ )

        # get decorrelated initial mixing matrix
        r_ = np.random.randn(self.n_components, self.n_components)
        i_ = np.random.randn(self.n_components, self.n_components)
        w_old = r_ + 1j * i_
        w_old = sym_decorrelation(w_old)

        for j in range(maxiter):

            # get new mixing matrix by updating columns one by one
            w_new = np.zeros((self.n_components,
                              self.n_components),
                             dtype=np.complex128)
            for i in range(w_old.shape[1]):
                y_ = np.dot(np.conj(w_old[:, i]).T, x)

                g_ = 1.0/(0.1 + np.abs(y_)**2)
                dg_ = -1.0 / (0.1 + np.abs(y_)**2)**2
```

```
            first = np.mean(x*np.conj(y_)*g_, axis=-1)
            second = np.mean(g_ + (np.abs(y_)**2)*dg_, axis=-1)
                ↪ *w_old[:, i]  # noqa
            w_new[:, i] = first - second

        # symmetrically decorrelate
        w_new = sym_decorrelation(w_new)

        # calculate convergence criterion
        criterion = (np.sum(np.abs(np.sum(w_new*np.conj(w_old),
            ↪  axis=1))) /
                     self.n_components)

        # show progress
        if j%30 == 0:
            print "Criterion:", str(1 - criterion),
            print "- Conveps:", str(conveps),
            print "- i:", str(j), "- maxiter:", str(maxiter)
            sys.stdout.flush()

        # check if converged
        if 1 - criterion < conveps:
            y_ = np.dot(np.conj(w_new.T), x)
            g_ = np.log(1 + np.abs(y_)**2)
            print '\nObjective values are: ' + str(np.mean(g_,
                ↪ axis=-1))
            print 'Convergence value: ' + str(1 - criterion)
            break

        # store old value
        w_old = w_new
```

```
        if j+1 == maxiter:
            raise Exception('ICA did not converge.')

        print 'ICA finished with ' + str(j+1) + ' iterations'

        return w_new, np.dot(np.conj(w_new).T, x)

    def _whiten(self, data):
        """
        Whiten data with PCA

        """
        # substract mean value from channels
        mean_ = data.mean(axis=-1)
        data -= mean_[:, np.newaxis]
        self._mean = mean_

        # calculate covariance matrix
        covmat = np.cov(data)

        # calculate eigenvectors and eigenvalues from covariance
            ↪ matrix
        eigw, eigv = np.linalg.eigh(covmat)

        # filter out components that are too small (or even
            ↪ negative)
        if not self.zerotolerance:
            self.zerotolerance = 1e-7

        valids = np.where(eigw/eigw[-1] > self.zerotolerance)[0]
        eigw = eigw[valids]
        eigv = eigv[:, valids]
```

```
        # adjust number of pca components
        n_components = self.n_components
        if not n_components:
            n_components = len(valids)
        elif n_components > len(valids):
            n_components = len(valids)
        self.n_components = n_components

        # sort in descending order and take only n_components of
        ↪ components
        eigw = eigw[::-1][0:n_components]
        eigv = eigv[:, ::-1][:, 0:n_components]

        # construct whitening matrix
        dsqrt = np.sqrt(eigw)
        dsqrtinv = 1.0/dsqrt
        whitening = np.dot(np.diag(dsqrtinv), np.conj(eigv.T))

        # whiten the data, note no transpose
        whitened = np.dot(whitening, data)

        # dewhitening matrix
        dewhitening = np.dot(eigv, np.diag(dsqrt))

        return dewhitening, whitened

    def _concat(self, data):
        """
        """
        fts = [data[:, :, idx] for idx in range(data.shape[2])]
        data2d = np.concatenate(fts, axis=1)
        return data2d
```

```
    def _split(self, data):
        """
        """
        parts = np.split(data, self._stft_shape[2], axis=1)

        xw = data.shape[0]
        yw = data.shape[1]/self._stft_shape[2]
        zw = self._stft_shape[2]

        splitted = np.empty((xw, yw, zw), dtype=data.dtype)
        for idx, part in enumerate(parts):
            splitted[:, :, idx] = part

        return splitted
```