Horizon 2020
European Union funding
for Research & Innovation

H2020 ICT-22-2014
COCOHA
Project No.: 644732

Cognitive
Control of a
Hearing Aid

*edited by Alain de Cheveigné, 13 September 2017*

# COCOHA report on EEG Preprocessing - v2

This document provides an overview of preprocessing techniques for EEG signals. It is part of Deliverable D2.2 of the COCOHA project, funded by the H2020 ICT programme of the European Union under grant number 644732. COCOHA aims to help hearing impaired persons so that they can deal with challenging noisy environments, by providing them with the means to steer sophisticated acoustic processing (such as microphone arrays) with control signals derived directly from the brain. This document is available at https://cocoha.org/cocoha-reports/. See also the COCOHA report on Alternative Technologies for Brain Signal Sensing.

## Executive summary

1. Preprocessing aims to improve the quality of EEG signals so that decoding algorithms can extract reliable control signals.
2. Raw EEG signals are contaminated by multiple artifacts from environmental noise sources (e.g. electromagnetic interference), the electrode-skin interface, muscular or ocular activity, and irrelevant brain activity.
3. A basic distinction can be made between *channel-specific* noise due for example to the electrode-skin interface or local muscle activity, and *channel-shared* noise, for example from irrelevant brain sources. Spatial filtering is more useful for the latter than the former.
4. Four strategies are available: (a) reduce noise at the source, (b) process signals to improve their quality, (c) design robust decoding algorithms, (d) train subjects to optimize their cortical signals. This document focuses on (b).
5. A wide panoply of methods is available from the literature, and new methods have developed within the COCOHA project.
6. EEG signals are typically measured on *multiple channels*, yielding multidimensional data to which linear methods such as spatial filtering can be applied (e.g. ICA or beamforming).
7. *Dimensionality* of the data is a critical factor. Artifacts are more likely to be separable within high-dimensional data (many channels) but this advantage may be mitigated by overfitting and noise.
8. The COCOHA project has developed several new methods. Of particular interest are methods that exploit the *nonstationarity* of the noise correlation structure. Several have been published, others are under development.
9. Preprocessing is crucial to the success of the COCOHA project.

## 1. The COCOHA context

The COCOHA project (http://cocoha.org/) funded by the EU H2020 initiative aims at developing a "smart" hearing aid in which acoustic processing is under the control of signals from the brain recorded by electroencephalography (EEG) or other means. With such a device a hearing impaired listener could focus attention on a particular sound source (for example a person speaking), and isolate it from noise and competing sources with the assistance of sophisticated acoustic processing (e.g. microphone arrays). Our intact auditory is quite adept at performing such a task, usually without our noticing, but this ability is reduced by impairment. The COCOHA project aims to restore this ability by artificial means.

A major impediment is that EEG signals are usually very noisy and contaminated by various artifacts that mask the desired cortical signals.  These may be removed either by preprocessing the EEG signals before extracting the control signals, or else by designing the control signal extraction algorithms such that they are insensitive to the presence of artifacts. The distinction between the two approaches is to some extent arbitrary, and here we adopt the viewpoint of preprocessing. Our focus is the COCOHA task, but much of what is said is applicable within the wider context of preprocessing of data from EEG or other modalities (such as MEG) for which noise and artifacts are an issue. These issues are very similar to those met for Brain Computer Interface (BCI) applications. In the COCOHA application, the device continuously monitors electrical activity from the brain for cues to the direction that the user wishes to attend.

A companion document, the COCOHA report on Alternative Technologies for Brain Signal Sensing, investigates the possibility of obtaining better signals at the origin.

## 2. The importance of preprocessing

With ~ 90 billion neurons and an order of magnitude more synapses, the human brain is an extremely complex machine (Herculano-Houzel 2012). EEG can record signals from anywhere between a handful of electrodes to a dense array of up to 1024 electrodes, so the number of "observable" signals is much smaller than the number of brain signals to observe, and the picture obtained is necessarily impoverished. Furthermore, current spread between cortical sources and electrodes implies that each electrode picks up multiple sources, whereas each source impinges upon multiple electrodes. The resulting many-to-many mapping between sources and electrodes greatly complicates the task of isolating useful activity, because useful signals are mixed with irrelevant cortical activity on every electrode. Finally, in addition to such irrelevant cortical activity, there are numerous non-cortical sources of noise and artifact that contaminate the signals, such as power-line noise (50 or 60 Hz), electromagnetic interference (e.g. from cell phones), skin-electrode contact noise, muscular artifacts, eye-blinks, etc.  The amplitude of unwanted signals can *greatly exceed* that of useful signals, in which case the data cannot be exploited without significant processing.

Typically, a control signal is derived from EEG by applying a "classifier" that defines a boundary within the space of measured data distinguishing between one outcome (e.g. *attend left*) and another outcome (e.g. *attend right*). The occurrence of an artifact can push the decision to the wrong side of that boundary, causing the device to make a decision that is contrary to the user's intention (e.g focus the acoustic processor to the left rather than the right).  Furthermore, the boundary applied by the classifier itself is usually "learned" from training data using various techniques taken from the field of Machine Learning (e.g. Lotte et al 2007, Wolpaw and Wolpaw 2012). The presence of artifacts in the data used for learning can lead to an incorrect or suboptimal boundary, for example if the artifacts were present more often for one class than for the other.

Preprocessing is thus a critical issue, both in the research and development phase for interpreting experimental data, and for the implementation of a workable device.

## 3. Sources of artifacts and noise

It is useful to understand the sources of artifact and noise, and the nature of their signals, so as to more effectively suppress or mitigate their effects. The distinction between artifact and noise itself is arbitrary (we use both words interchangeably), but one thing is clear: we should not expect 'noise' to be of the Gaussian and/or white sort familiar to many signal-processing practitioners. Apart from thermal noise which is spectrally white with Gaussian characteristics, most sources have non-stationary, non-white and/or non-Gaussian distributions, and there may also be correlations between different sources.

Noise sources can be divided into three classes: (1) environmental, such as power lines or electromagnetic sources, (2) instrumental, such as electrode/skin contact noise or quantization noise, and (3) physiological, including muscle artifacts, eye-blinks, cardiac signals, and irrelevant neural activity.

From a signal processing point of view it is useful make a different classification between: (a) channel-specific noise, such as electrode/skin contact noise or localized artifacts from shallow muscles proximal to an electrode, and (b) channel-shared noise, such as results from many environmental and physiological sources. The latter can often be attenuated with the help of a spatial filter (for example provided by a beamforming or ICA algorithm) whereas the former cannot.

### 3.1 Power line noise

Power line noise is a ubiquitous artifact. Electrical power is usually distributed as sinusoidally alternating currents at 50 or 60 Hz within power cables, common in most environments, that radiate both electrical and magnetic fields. EEG signals are tiny (~ 1-100 microvolts) and easily swamped by artifactual signals via capacitive coupling (electric field) and/or inductive coupling (magnetic field). The contaminating signal may include components at the fundamental frequency (50 or 60 Hz) as well as multiples of that frequency. The amplitude and phase of all of these components may fluctuate depending on shifts in power consumption (possibly remote within the power grid), and of course with movements of the subject within these fields. These artifacts are addressed by (a) careful design of circuits and shielding, (b) spectral filtering (e.g. a notch filter) and (c) regression or spatial filtering, see below.

### 3.2 Other electromagnetic sources

Electromagnetic waves can be emitted by various kinds of apparatus (video monitors, switching power supplies, computers, wifi stations, cell phones and cell phone relays, radio and TV transmitters etc.). This interference is often in a high frequency range, whereas EEG signals are usually filtered to a restricted range (e.g. < 500 Hz) by the hardware filter that precedes analog-to-digital conversion. EEG might thus be expected to be immune to such higher frequency noise, but two factors can thwart that expectation. One is that a high-amplitude signal may be insufficiently attenuated by the (finite) attenuation of the filter. Another is that the high-frequency signals may be *demodulated* (or *intermodulated* with other signals) due to non-linearities in the electronics (e.g. clipping) or at the skin-electrode contact. Such situations are very hard to diagnose because the effect (low frequency signal) bears little resemblance to its cause (high-frequency source). Motion or vibration within a magnetic field (e.g. Earth's) can induce potentials in the EEG leads, and in the presence of electrical charges (e.g. triboelectric), movement-induced variations of capacitance between electrodes and nearby objects can cause potential variations.

### 3.3 Sensor and electrode contact noise

The contact between electrode and skin is the site of multiple noise-generating phenomena (Huigen et al 2002; Hokajärvi 2012). For electrodes, most commonly used in the laboratory (Ag/AgCl with electrolytic gel), the noise is mainly produced by the skin/electrolyte interface. In the absence of motion, this noise has a low-pass spectral characteristic ($1/f^{\alpha}$ with $1.5 < \alpha < 2.0$; Huigen et al 2002), implying that the noise is dominated by slow variations (electrode drift). These are troublesome because they interfere with the analysis of relatively slow EEG patterns, and require the use of high-pass filters that can introduce other issues (see below). Sweat can also induce slow variations.

Relatively high amplitude artifacts can be generated by *motion* of the electrode/gel relative to the skin or deformation of the skin itself. These can be manifest as a high-amplitude transient, and/or a step of the resting potential. A particular form of motion artifact called "pulse" can arise from a blood vessel proximal to the electrode.

The noise from one electrode is usually uncorrelated with that of the other electrodes (it is "channel specific"). However, each electrode measurement is necessarily made *relative to another electrode* (reference). Noise at the reference electrode contact will on the contrary appear as correlated over signals measured on the other electrodes. Motion artifacts might also be correlated across several electrodes. Uncorrelated noise is problematic because it cannot be factored out using linear component analysis techniques such as ICA.

Electrode-skin noise is a major issue for an application such as ours, because the steps that can be taken to reduce it (electrolytic gels, skin abrasion) usually conflict with requirements of comfort and ergonomics. Less constraining alternatives (e.g. dry electrodes) tend to have higher noise. Dealing with this type of noise is an important task for preprocessing.

Additional sources of noise are thermal noise (at the electrode-gel interface and in the electronics) and quantization noise. They can usually be neglected relative than other sources, although they may need to be taken into account if (a) reduced quantization (e.g. 8-16 bits) is imposed by implementation constraints, (b) information is to be gathered from the higher frequency regions.

### 3.4 Muscle artifacts

Many small muscles are present under the skin, in particular on the head where EEG electrodes are attached. Muscle artifacts typically take the form of regular or irregularly-spaced spike trains, with a spectral composition dominated by relatively high frequencies (Goncharova et al 2003; Fatourechi et al 2007; McMenamin et al 2011; Ma et al 2012). Muscle artifacts are troublesome because they mask cortical signals in the gamma band (>20 Hz) and they may also depend on mental state or task, and thus masquerade as cortical correlates of those tasks. Artifacts from shallow muscle fibers close to an electrode may be specific to that fiber, whereas deeper muscle activity may be correlated across electrodes.

Muscle artifacts are a nuisance in experiments that try to measure cortical activity but, in a control system they are potentially of value if the user is able to learn to reliably and selectively control the muscles that produce them. A brain-control purist might object, but that is of little concern if the system works.

### 3.4 Ocular and cardiac artifacts

The eyeball acts like an electric dipole with an anterior positive pole (cornea) and posterior negative pole (retina). Eye ball movements (in particular eyeblinks) produce large deflections, principally in electrodes in frontal positions. These may mask cortical activity of interest, or even masquerade as cortical activity (if correlated with cognitive state). They can also be of use as control signals if the user is willing to have his or her eye movements enrolled for that purpose. Ocular control is a plausible option to control a hearing aid (Kidd et al 2013).

EEG signals may also be contaminated from potentials from the tongue (which is also polarized) as well as from heart activity. Cardiac, ocular and similar artifacts usually affect multiple electrodes, which makes them amenable to linear component analysis techniques. Signals from electrodes placed so as to pick up only cardiac activity (electrocardiography, ECG) or ocular activity (electrooculography, OEG) can be used to project out cardiac or ocular activity from the EEG using regression techniques.

### 3.5 Unwanted cortical activity

The brain is the theater of countless neural processes that all impinge on the EEG electrodes. Only a fraction of the measured signal reflects any single process, for example a cortical process indicative of attention. Most of the signal variance reflects the myriad other ongoing cortical processes. An example of a prominent contribution to EEG signals is known as "alpha" activity, a high-amplitude oscillatory signal that occurs in bursts, with a frequency in the 8-10 Hz region. There are multiple sources of alpha activity. Alpha amplitude from occipital sources increases when eyes are closed, whereas that from sources in other regions may *decrease* when cortical processing is engaged. The large amplitude and dimensionality (see below) of alpha may contribute to obscure other sources of interest, However, to the extent that characteristics of alpha activity are indicative of attential, it may also be harnessed to derive a control signal.

In addition to alpha, there are many other forms of ongoing EEG activity that must be suppressed or discounted if we wish to derive a reliable control signal. Cortical sources are usually deep enough to impinge on multiple electrodes, and thus are amenable to removal using linear spatial filtering techniques (e.g. ICA).

## 4. Strategies to suppress or mitigate artifacts and noise

There exists a large range of preprocessing tools reported in the literature, to which the COCOHA project has contributed.

### 4.1 Dimensionality

Before delving into specific strategies and tools, it is worth considering a very useful concept, that of *dimensionality* of the data. The physical process that produces the EEG signal is usually linear to a very good approximation, which means that the signal produced by two brain sources is the sum of the signals produced by each (additivity), and the amplitude of the measured signal scales with the amplitude of the source that produced it (proportionality). That being the case, it is useful to describe the measured signals (as well as the neural electrical activity that gives rise to them) as belonging to a *vector* space. A vector space is a set for which any sum of elements (or "points") also belongs to the space, as does the product of any element by a scalar.

The EEG signals recorded by an array of electrodes span a *vector space* that includes all weighted sums of these signals (such as might be produced by a spatial filter, or an analysis algorithm such as ICA). For N electrodes, the *dimensionality* of such a space is at most N. It can be smaller than the number of signals if the signals are linearly dependent, e.g. one signal equals a weighted sum of the others. The myriad sources of electrical activity within the brain (and elsewhere) span a vector space of much larger dimensionality. Since EEG signals are weighted sums of these signals, the space that they span is a *subspace* of the larger space

The concepts of subspace and dimensionality are very useful. Many analysis strategies can be understood as finding a subspace that spans most of the interfering noise and artifacts, so that brain activity of interest can be observed within the subspace orthogonal to it. Another way of expressing this, is to say that noise and artifact are *projected*, or *regressed* out of the data. This strategy can only work if interference and target live within distinct subspaces, which requires that there be enough dimensions to start with. In general, P distinct and uncorrelated noise sources span a subspace of dimension P, and for there to exist a subspace orthogonal to it containing activity of interest we must have $N>P$. Moreover, if there are Q distinct sources of interest, we must have $N>=P+Q$ As a rule of thumb: the greater the dimensionality of the data the better. The more electrodes, sensors, etc., the more interference sources can be projected out, and the more sources of interest can be resolved.

There are two caveats to this rule. One is the presence of *sensor-specific* noise (uncorrelated with other sensors), such that the space spanned by N electrode signals already contains N noise sources. That being the case, other sources can only be resolved approximately. The maximum number of sources that can be isolated depends on the "noise floor" determined by sensor-specific noise, so that once that number is reached adding more sensors will not be useful. The second caveat is that many methods involve parameters (such as regression coefficients) that must be *learned* from the data. The higher the dimensionality the greater the risk of *overfitting*. This too may impose a practical limit on the number of additional sensor channels to consider.

The number of dimensions of N-channel data is at most N, but it can be smaller in particular if each channel contains fewer than N samples of data. For M samples the dimensionality is the smaller of M and N. Serial correlations within the data (for example if it is low-pass filtered) may cause the effective dimensionality to be yet smaller. In general, signals that have high serial correlation tend to behave as if they were shorter, and have fewer degrees of freedom, than their number of samples would suggest.

## 4.2 Preprocessing approaches

### *4.2.1 Denoise or discount?*

Corruption by noise or artifact cause the data to be less reliable than if they were intact. *Denoising* involves processing to remove or attenuate the noise, whereas *discounting* involves marking invalid portions so that they do not influence the outcome of processing. For example, if an electrode is detached or has poor contact, its data may be completely unreliable, or if a strong electrical glitch affected all channels during a time interval, data values corresponding to that interval may need to be discounted.

Discounting may be associated with *interpolation* to restore the semblance that the data are complete. It is important to realize that interpolation does not restitute the information that was lost by masking by the artifact and/or discounting. Interpolation is convenient because standard processing algorithms to be used, but it may invalidate statistical analysis by reducing the amount of variance and/or the degrees of freedom in the data. For example, interpolating a missing channel as a weighted sum of its neighbors reduces data dimensionality by 1.

#### *4.2.2 Fourier filtering*

Noise and signal may have different spectral characteristics. For example, power-line noise is usually concentrated at 50 Hz (or 60Hz) and its harmonics, alpha activity is usually concentrated in the 8-12 Hz region, and electrode drift is mainly restricted to very low frequencies.

Filters are perhaps the most commonly used tool in our panoply. In addition to the anti-aliasing low-pass filter in hardware that precedes analogue-to-digital conversion, it is common to use a high pass filter (with cutoff frequency typically in the 0.1 Hz to 1 Hz range) to attenuate drift, and possibly a notch filter at the line frequency, or a low-pass filter (for example at 20 or 30Hz) to attenuate higher components judged non-informative and improve the smoothness of waveform plots.

Fourier filtering is extremely useful but entails several risks. An obvious concern is that the signal may extend to the region suppressed by the filter, and thus be distorted. This is best understood by noting that filtering involves *convolution* of the waveform by the impulse response of the filter, so that each sample of the filtered signal is a weighted sum of *several* samples of the original signal. The new sample thus reflects events within an extended time interval, and the temporal relation between the filtered waveform and events in the brain is thus blurred, which may be a problem when making inferences about the latency of a brain response relative to a stimulus, or the causal relations between events. Furthermore, if the filter's transfer function is narrow or has a sharp transition, its response to a transient event may be oscillatory (ringing) possibly leading to incorrect conclusions concerning the oscillatory nature of brain activity. High-pass filtering may convert a unipolar pulse into a multiphasic response, the positive and negative deflections of which are purely artifactual (i.e. reflecting only properties of the filter). The span of such temporal distortion equals the length of the filter's impulse response, and is generally more marked as the filter is spectrally narrow or sharp.

For these reasons, it is common to consider alternatives to Fourier filtering as described below.

#### *4.2.3 Spatial filtering*

Noise and signal may likewise have different *spatial* characteristics. Spatial filtering consists in replacing the signal on each channel by a weighted sum of signals on all other channels. With appropriate weights, spatial filtering can attenuate or suppress the contribution of one or several noise sources. Weights can be predetermined, or else derived automatically based on a data-driven algorithm. Linear spatial filtering is a powerful tool that makes full use of the multidimensional nature of data provided by EEG electrode arrays.

Examples of predetermined spatial filtering are *rereferencing*, where the signal from a particular electrode (for example mastoid), or the mean signal over electrodes, is subtracted from the signals of all electrodes, or a *Laplacian* filter, where the mean of the nearest neighbors is subtracted from each electrode. Examples of data-driven filtering are the filters determined by component-analysis techniques such as *ICA* or *beamforming*.

By extension, suppression of a reference signal (for example ECG or EOG) by regression on the data can be assimilated to a form of spatial filtering (the reference signal being treated as one particular spatial channel). A spatial filter can be understood as a linear transform in the space spanned by the sensor signals. A filter that suppresses a noise source can be understood as defining a *projection* on the subspace orthogonal to that noise source.

More on the design of optimal spatial filters below.

#### 4.2.3 Detrending

Particularly troublesome is the "drift" that arises at the electrode-skin interface, a slowly varying potential that shifts the "baseline" potential on each EEG electrode. It interferes with the analysis of slow potentials in the brain, and is a primary motivation for applying high-pass filtering to EEG. High-pass filtering with a low-frequency cutoff (typically 0.1 to 1 Hz) requires a long impulse response, potentially leading to extensive filter-induced distortions. Electrode drift is usually uncorrelated between electrodes, and thus is not amenable to spatial filtering techniques. Furthermore, the strong serial correlation reduces the degrees of freedom, promoting overfitting.

An alternative to high-pass filtering is *detrending*, in which a slowly-varying function (for example a linear ramp or low-order polynomial) is fit to the data, and then subtracted. Each sample of the detrended data is now function of *all* the data, raising concerns similar to those raised for filtering, but the constraint of a low-order fit limits the impact of waveform distortion.

In the presence of temporally-local glitches (for example electrode motion artifacts) detrending suffers a similar problem as met in high-pass filtering where a glitch can cause ringing of the filter. The fit is affected by the glitch, causing the trend to be imperfectly removed on the non-glitch portions, and in some situations, a trend can be introduced where none was present initially. A solution is *robust detrending*, in which samples too distant from the fit are simply discounted, so that the fit depends only on the non-glitch parts. The glitch itself is then addressed by other means (see below). Similar processing is harder to implement in a high-pass filter, so robust detrending is a tool of choice to address electrode drift.

#### 4.2.4 Sensor noise suppression

Sources that impinge on several sensors can be suppressed by forming a linear combination with appropriate coefficients so that they cancel out. This is what techniques such as ICA or beamforming achieve. A source that impinges on a single sensor or electrode (for example a local muscle artifact, or electrode-skin noise that was not suppressed by detrending) cannot be suppressed in this way. The only way to suppress that noise is to discard the channel, which is of course wasteful.

In other words, the added dimension that the channel offers to the EEG representation is squandered by its own noise, and if all channels carry such noise, then it is impossible to obtain a clean subspace by linear projection. This noise floor imposes a hard limit on our ability to extract weak brain activity from EEG. Without that channel-specific noise floor, we could partition the data into as many orthogonal subspaces as there are sensors, thus improving our chances of isolating activity of interest, however weak. With the noise floor, most analyses yield much fewer exploitable dimensions than sensors. Dealing with channel-specific noise is thus essential.

This issue is of prime importance within the COCOHA project, because usability constraints will likely impose a small number of channels, possibly with dry electrodes, and there may be additional artifacts due to movements. Building on our Sensor Noise Suppression algorithm (SNS) (de Cheveigné and Simon 2008) we developed the Sparse Time Artifact Removal algorithm (STAR) (de Cheveigné 2016) that identifies channel-specific glitches and interpolates them from the intact channels. Whereas the previous SNS algorithm applied a single linear transform to all the data to suppress noise, STAR applies a distinct linear transform restricted to the corrupted segments of each channel, maintaining the full dimensionality of the data. Work is ongoing to lift current limits on the applicability of the algorithm (currently it cannot handle the case where glitches occur simultaneously on several channels.

#### *4.2.4 Optimal data-driven spatial filters*

Linear spatial filtering allows the data to be projected into the null subspace of noise and artifacts. The signal-to-noise ratio (SNR) improvement is potentially very large, limited only by the noise floor imposed by irreducible channel-specific noise (see above). However this outcome depends crucially on the choice of the filter weights: a slight error in weights can allow noise components to leak into the cleaned signal.

A data-driven solution can be automatically tuned to the particular artifacts present in the data, and thus is potentially of better quality than a solution based on prior knowledge, particularly as the geometry of artifact sources and the properties of the propagation milieu may not be well known. The downside is that the solution may be prone to overfitting, which occurs if there is a mismatch between the structure of the data used to tune the filter, and the data to which it is applied.

There are many data-driven approaches to finding spatial coefficients. The well-known *Principal Component Analysis* (PCA) is sometimes proposed for denoising. Its effectiveness in that role is limited, but it serves as an important ingredient in other methods. *Independent Component Analysis* (ICA) allows data to be transformed into a set of components that are not only uncorrelated, but also statistically independent according to some measure of independence. Measures of independence fall in three main classes: non-Gaussianity, non-stationarity, non-whiteness (Cardoso 2001), and countless methods have been proposed that exploit various variants of these measures (Choi et al 2005; Parra et al 2005; Hyvarinen 2012). Many of these have been applied to EEG analysis and BCI applications (e.g. Cichocki, A. 2004; Delorme et al 2007, 2012 Nicola-Alonzo 2012). ICA produces a set of "independent" components which are then sorted (manually or by an automatic procedure) into noise and target categories. ICA is often successful and often used as a preprocessing tool. Downsides are (a) there is not a single authoritative algorithm but rather a wide range of methods, (b) several methods are stochastic and do not produce the same result on every trial, (c) computational costs, (d) the need for post-hoc component selection.

*Beamforming* finds a filter that minimizes power from all sources except those from a particular target direction, defined by its lead matrix. Beamforming has a long history first in radio antenna array processing, then in microphone array processing (e.g. Benesty 2007), and is a well-developed technique in brain signal processing (e.g.

Sekihara et al 2006; Grosse-Wenstrup et al 2009). While data-driven, the method requires knowledge of the lead matrix of the desired sources. Classically this is obtained from geometrical knowledge, but it can also be derived from other sources.

*Common Spatial Patterns* (CSP) is a powerful method commonly used in the BCI community, but strangely less in the EEG and brain imaging community. CSP finds a linear transform of the data that maximizes the variance ratio between two intervals (Koles et al 1990; Fukunaga et al 1970; Bashashati et al 2007; Nicola-Alonso et al 2012; Kawanabe et al 2014; Samek et al 2014). Our work has contributed to extending the applicability of this method in the wider framework of *Joint Decorrelation* (JD) or *Denoising Source Separation* (DSS) (de Cheveigné and Simon 2008; de Cheveigné et al 2010, 2012; de Cheveigné and Parra 2014). This has been applied in the COCOHA project to optimize the extraction of oscillatory components such as alpha that are a useful cue to spatial attention (de Cheveigné and Arzounian 2015).

A word of warning: data-driven solutions are designed to optimize a criterion (e.g. narrow-band power), and may appear misleadingly successful in that respect (Kriegeskorte et al 2009). When judging the effectiveness of the outcome it is important to use appropriate cross-validation techniques. This is related to the issue of overfitting in which the solution is tuned to the training data but not generalizable to unseen data.

#### *4.2.5 Time-varying analysis*

As stressed above, dimensionality of the noise subspace (and/or noise floor of the remaining signal subspace) sets the ultimate limit of our ability to isolate weak cortical activity. Isolating weak activity (by suppressing the many stronger sources) is paramount to our application. While machine-learning techniques (e.g. deep networks) may be useful to optimally exploit the available noise-corrupted data, it is obvious that radically better performance can be expected if the noise is removed.

The dimensionality of the noise space is determined by the number of distinct (uncorrelated) noise sources that are active within the analysis interval. However, the activity of many sources is *temporally sparse*, and so we can expect noise dimensionality to be smaller if analysis is conducted on shorter intervals. This idea is exploited by the STAR algorithm (de Cheveigné 2016) to suppress channel-specific sources of noise without reducing dimensionality (as would occur with e.g. ICA, JD or beamforming).

The thrust of our efforts in the COCOHA project is devoted to extending these ideas. We believe that smart time-varying analysis is the ultimate frontier towards a robust and reliable decoding device.

### 4.3 Online real-time preprocessing

Most algorithms (e.g. ICA, beamforming, JD, etc.) are formulated for batch processing, which is not useful in a device that must produce a control signal in real time. It is essential that the algorithms that we propose be realizable in real time. This constraint should not be imposed in the algorithm development phase, because it would slow the development of new ideas, but it remains a preoccupation. Fortunately most algorithms take as their starting point the *covariance matrix* of the data (possibly augmented with a range of time shifts), which is the mean over time of the cross-products between channels, which can be calculated according to a time-varying schedule and updated in real time. Real-time processing is supported by the COCOHA toolbox.

### 4.4 Toolboxes for preprocessing

Many preprocessing methods are implemented in publicly available software toolboxes (in Python, Matlab or C++). We maintain the NoiseTools toolbox (http://audition.ens.fr/adc/NoiseTools/) which offers general preprocessing tools, and the COCOHA toolbox (soon to be made public) oriented to brain decoding.

Many other toolboxes are available, including EEGLab (https://sccn.ucsd.edu/eeglab/), FieldTrip (http://www.fieldtriptoolbox.org/), SPM (http://www.fil.ion.ucl.ac.uk/spm/), BrainStorm (http://neuroimage.usc.edu/brainstorm/), Biosig (http://biosig.sourceforge.net), NutMEG (http://www.nitrc.org/projects/nutmeg/), MNE (http://martinos.org/mne/stable/index.html), OpenMEEG (http://openmeeg.github.io), SciKit (http://scikit-learn.org/stable/).

Some resources for realtime processing include: OpenEEG (http://openeeg.sourceforge.net/doc/), OpenVibe (http://openvibe.inria.fr), PureData (https://puredata.info), EEGSynth (https://github.com/eegsynth/eegsynth).

### Summary

Electrical signals recorded by EEG electrodes are extremely noisy, and this is a major obstacle for deriving robust and timely control signals. Preprocessing is necessary to remove and/or discount the noise and artifacts. Preprocessing is a major thrust of the efforts within the COCOHA project to design a brain-controlled hearing aid device that can well enough to really help people who need it.

### References


Bashashati, A., Fatourechi, M., Ward, R. K., & Birch, G. E. (2007). A survey of signal processing algorithms in brain–computer interfaces based on electrical brain signals. Journal of Neural Engineering, 4(2), R32–R57. Retrieved from http://stacks.iop.org/1741-2552/4/i=2/a=R03?key=crossref.cdbb14e4d6a8a03f7c5a86bc262cb3a9

Benesty, J., Chen, J., Huang, Y., & Dmochowski, J. (2007). On Microphone-Array Beamforming From a MIMO Acoustic Signal Processing Perspective. IEEE Transactions on Audio, Speech, and Language Processing, 15(3), 1053–1065. Retrieved from http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4100665

Bertrand, A (2011) "Applications and trends in wireless acoustic sensor networks: a signal processing perspective" Proc. of the IEEE Symposium on Communications and Vehicular Technology (SCVT)

Choi, S., Cichocki, A., Park, H.-M., & Lee, S.-Y. (2005). Blind Source Separation and Independent Component Analysis: A Review. Neural Information Processing - Letters and Reviews, 6, 1–57. Retrieved from papers2://publication/uuid/CC1F69E9-3AD0-4D03-B15B-DE37A5239D8C

Cichocki, A. (2004). Blind signal processing methods for analyzing multichannel brain signals. International Journal of Bioelectromagnetism, 6(1). Retrieved from papers2://publication/uuid/A49A2C45-5802-4B7E-8F26-0D9F37769471

D.Y. Levin, E.A.P. Habets and S. Gannot (2015) On the average directivity factor attainable with a beamformer incorporating null constraints, IEEE Signal Processing Letters, Vol. 22, No. 11, pp. 2122-2126.

de Cheveigné, A. (2010). Time-shift denoising source separation. Journal of Neuroscience Methods, 189(1), 113–120. Retrieved from http://linkinghub.elsevier.com/retrieve/pii/S0165027010001202

de Cheveigné, A. (2012). Quadratic component analysis. NeuroImage, 59(4), 3838–3844. Retrieved from http://linkinghub.elsevier.com/retrieve/pii/S1053811911012560
de Cheveigné, A. (2016). Sparse time artifact removal. Journal of Neuroscience Methods, 262, 14–20. Retrieved from http://dx.doi.org/10.1016/j.jneumeth.2016.01.005
de Cheveigné, A., & Arzounian, D. (2015). Scanning for oscillations. Journal of Neural Engineering, 12(6), 66020. Retrieved from http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2015JNEng..12f6020D&link_type=EJOURNAL
de Cheveigné, A., & Parra, L. C. (2014). Joint decorrelation, a versatile tool for multichannel data analysis. NeuroImage, 98, 487–505. Retrieved from http://linkinghub.elsevier.com/retrieve/pii/S1053811914004534
de Cheveigné, A., & Simon, J. Z. (2008). Denoising based on spatial filtering. Journal of Neuroscience Methods, 171(2), 331–339. Retrieved from http://linkinghub.elsevier.com/retrieve/pii/S0165027008002008
de Cheveigné, A., & Simon, J. Z. (2008). Sensor noise suppression. Journal of Neuroscience Methods, 168(1), 195–202. Retrieved from http://linkinghub.elsevier.com/retrieve/pii/S0165027007004621
Delorme, A., Palmer, J., Onton, J., Oostenveld, R., & Makeig, S. (2012). Independent EEG sources are dipolar. PLoS ONE, 7(2), e30135. Retrieved from http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=22355308&retmode=ref&cmd=prlinks
Delorme, A., Sejnowski, T., & Makeig, S. (2007). Enhanced detection of artifacts in EEG data using higher-order statistics and independent component analysis. NeuroImage, 34(4), 1443–1449. Retrieved from http://linkinghub.elsevier.com/retrieve/pii/S1053811906011098
F. Lotte, M. Congedo, A. Lécuyer, F. Lamarche, B. Arnaldi, "A Review of Classification Algorithms for EEG-based Brain-Computer Interfaces", Journal of Neural Engineering, 4, R1-R13, 2007.
Fatourechi M, Bashashati A, Ward RK, Birch GE (2007) EMG and EOG arti- facts in brain computer interface systems: A survey. Clinical Neurophysiol- ogy 118:480–494.
Fukunaga, K., & Koontz, W. L. G. (1970). Application of the Karhunen-Loeve expansion to feature selection and ordering. Computers, IEEE Transactions on, C-19(4), 311–318. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1671511
Goncharova II, McFarland DJ, Vaughan TM, Wolpaw JR (2003) EMG contami- nation of EEG: spectral and topographical characteristics. Clinical Neurophys- iology 114:1580–1593.
Grosse-Wentrup, M., Liefhold, C., Gramann, K., & Buss, M. (2009). Beamforming in noninvasive brain-computer interfaces. Biomedical Engineering, IEEE Transactions on, 56(4), 1209–1219.
Herculano-Houzel, S. The remarkable, yet not extraordinary, human brain as a scaled-up primate brain and its associated cost}, Proc Nat Acad Sci 109, Supplement 1, 10661-10668.
Hokajärvi I.A. Master's Thesis. Tampere University of Technology; Tampere, Finland: 2012. Electrode Contact Impedance and Biopotential Signal Quality.
Huigen, E.; Peper, A.; Grimbergen, C.A. Investigation into the origin of the noise of surface electrodes. Med. Biol. Eng. Comput. 2002, 40, 332–338.
Hyvarinen, A. (2012). Independent component analysis: recent advances. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 371(1984), 20110534. Retrieved from http://rsta.royalsocietypublishing.org/cgi/doi/10.1098/rsta.2011.0534
J Benesty, J Chen, Y Huang, J Dmochowski (2007) On microphone-array beamforming from a MIMO acoustic signal processing perspective, Audio, Speech, and Language Processing, IEEE Transactions on 15 (3), 1053-1065
J. Wolpaw and E. Wolpaw, Brain-computer interfaces: principles and practice. Oxford University Press, 2012
Kawanabe, M., Samek, W., Müller, K.-R., & Vidaurre, C. (2014). Robust Common Spatial Filters with a Maxmin Approach. Neural Computation, 26(2), 349–376. Retrieved from http://www.mitpressjournals.org/doi/abs/10.1162/NECO_a_00544

Kidd, G., Favrot, S., Desloge, J. G., Streeter, T. M., & Mason, C. R. (2013). Design and preliminary testing of a visually guided hearing aid. The Journal of the Acoustical Society of America, 133(3), EL202. Retrieved from http://scitation.aip.org/content/asa/journal/jasa/133/3/10.1121/1.4791710

Koles, Z. J., Lazar, M. S., & Zhou, S. Z. (1990). Spatial patterns underlying population differences in the background EEG. Brain Topography, 2(4), 275–284. Retrieved from http://www.springerlink.com/index/L378X271583V34P7.pdf

Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S. F., & Baker, C. I. (2009). Circular analysis in systems neuroscience: the dangers of double dipping. Nature Neuroscience, 12(5), 535–540. Retrieved from http://www.nature.com/doifinder/10.1038/nn.2303

Ma J, Tao P, Bayram S, Svetnik V (2012) Muscle artifacts in multichannel EEG: characteristics and reduction. Clinical Neurophysiology 123:1676–1686.

Nicolas-Alonso, L. F., & Gomez-Gil, J. (2012). Brain Computer Interfaces, a Review. Sensors, 12(12), 1211–1279. Retrieved from http://www.mdpi.com/1424-8220/12/2/1211/

Parra, L. C., Spence, C. D., Gerson, A. D., & Sajda, P. (2005). Recipes for the linear analysis of EEG. NeuroImage, 28(2), 326–341. Retrieved from http://linkinghub.elsevier.com/retrieve/pii/S1053811905003381

Samek, W., Kawanabe, M., & Muller, K. R. (2014). Divergence-based framework for common spatial patterns algorithms. IEEE Reviews in Biomedical Engineering, 7, 50–72. http://doi.org/10.1109/RBME.2013.2290621

Sekihara, K., Hild, K. E., & Nagarajan, S. S. (2006). A novel adaptive beamformer for MEG source reconstruction effective when large background brain activities exist. IEEE Transactions on Bio-Medical Engineering, 53, 1755–1764. Retrieved from papers2://publication/uuid/A2E55AF8-0083-426E-807B-625ABB179FA6

Whitham EM, Pope KJ, Fitzgibbon SP, Lewis T, Clark CR, Loveless S, Broberg M, Wallace A, DeLosAngeles D, Lillie P, Hardy A, Fronsko R, Pulbrook A, Willoughby JO (2007) Scalp electrical recording during paralysis: quantitative evidence that EEG frequencies above 20 Hz are contaminated by EMG. Clinical Neurophysiology 118:1877–1888.