Journal of Neuroscience Methods 262 (2016) 14–20

Contents lists available at ScienceDirect

Journal of Neuroscience Methods

journal homepage: www.elsevier.com/locate/jneumeth

Basic Neuroscience

Sparse time artifact removal

Alain de Cheveigné [a,b,c,*]

[a] Laboratoire des Systèmes Perceptifs, UMR 8248, CNRS, France
[b] Département d'Etudes Cognitives, Ecole Normale Supérieure, France
[c] UCL Ear Institute, United Kingdom

HIGHLIGHTS

- The STAR algorithm addresses channel-specific noise that is sparse in time.
- It removes electrode or sensor noise and certain forms of myogenic artifact.
- In contrast to other techniques, few data are lost and the dimensionality of the data is preserved.
- The STAR algorithm complements component analysis techniques such as ICA.

ARTICLE INFO

Article history:
Received 25 October 2015
Received in revised form
23 November 2015
Accepted 2 January 2016
Available online 8 January 2016

Keywords:
EEG
MEG
LFP
ECoG
Artifact
Myogenic
ICA
Sensor noise

ABSTRACT

Background: Muscle artifacts and electrode noise are an obstacle to interpretation of EEG and other electrophysiological signals. They are often channel-specific and do not fully benefit from component analysis techniques such as ICA, and their presence reduces the dimensionality needed by those techniques. Their high-frequency content may mask or masquerade as gamma band cortical activity.
New method: The sparse time artifact removal (STAR) algorithm removes artifacts that are sparse in space and time. The time axis is partitioned into an artifact-free and an artifact-contaminated part, and the correlation structure of the data is estimated from the covariance matrix of the artifact-free part. Artifacts are then corrected by projection of each channel onto the subspace spanned by the other channels.
Results: The method is evaluated with both simulated and real data, and found to be highly effective in removing or attenuating typical channel-specific artifacts.
Comparison with existing methods: In contrast to the widespread practice of trial removal or channel removal or interpolation, very few data are lost. In contrast to ICA or other linear techniques, processing is local in time and affects only the artifact part, so most of the data are identical to the unprocessed data and the full dimensionality of the data is preserved.
Conclusions: STAR complements other linear component analysis techniques, and can enhance their ability to discover weak sources of interest by increasing the number of effective noise-free channels.

1. Introduction

Among the many sources of noise and artifact that plague studies involving human and animal electrophysiology, some affect only one channel at a time. This paper addresses such *channel-specific* artifacts, leaving aside other types that impinge on several channels such as eyeblink, heartbeat or background neural activity. Signals recorded by multichannel recording techniques such as electroencephalography (EEG), magnetoencephalography (MEG), electrocorticography (ECoG), invasive electrode arrays or optical techniques are related to underlying sources by a linear mixing process:

$$x_j(t) = \sum_i s_i(t) u_{ij}, \quad (1)$$

where t is time, $s_i(t)$ are the source signals, and u_{ij} are mixing coefficients. We call "channel-specific" a source s_i for which the u_{ij} are zero for all channels j except one.

Channel-specific noise includes electrode contact noise, pulsation noise, and certain forms of muscle artifact in EEG, as well as sensor noise in MEG or other techniques. Electrode-tissue contact artifacts are usually temporally sparse, occurring as isolated events or bursts of events that affect one channel at a time.

* Correspondence to: Audition, DEC, ENS, 29 rue d'Ulm, 75230 Paris, France. Tel.: +33 144322672/+44 7912504027.
E-mail address: Alain.de.Cheveigne@ens.fr

http://dx.doi.org/10.1016/j.jneumeth.2016.01.005

Muscle artifacts may also be channel-specific if they are produced by motor units proximal to a single electrode (Yilmaz et al., 2014), although activity from deeper muscles may affect multiple electrodes. Such artifacts often have a spectrum that extends to high frequency, but they may also include low-frequency components that overlap with low-frequency cortical activity (Goncharova et al., 2003), so that spectral filtering is not sufficient to eliminate them. Electromyogenic noise is particularly troublesome as it can be confused with high-frequency cortical activity (Yuval-Greenberg et al., 2008; Whitham et al., 2007; Muthukumaraswamy, 2013; Yilmaz et al., 2014), particularly as muscle activity may correlate with cognitive state (Whitham et al., 2008; Yuval-Greenberg et al., 2008; McMenamin et al., 2011). Cephalic skin potential (Corby et al., 1974) may also covary with cognitive state, and contact noise may correlate with behavior.

A standard approach to dealing with spatio-temporally sparse artifacts is to discard either the offending channel, or the offending time interval or trial (Junghöfer et al., 2000). This entails loss of data, particularly if artifacts affect multiple channels and/or are widely distributed in time, and it also complicates analysis and interpretation stages, that need to be made tolerant to the missing data.

Another approach is to apply multichannel linear analysis techniques such as independent component analysis (ICA), beamforming, or joint diagonalization (JD) (de Cheveigné and Parra, 2014) to isolate noise components. Component signals $y_k(t)$ obtained by these methods are related to observations as:

$$y_k(t) = \sum_j x_j(t) w_{jk}, \tag{2}$$

where t is time and the w_{jk} are weights. The J observation channels span a space that contains all such linear combinations, and the *dimensionality* of the data is the number of dimensions of this space (J or less). Components belong to this space. Different methods (PCA, ICA, beamforming, etc.) differ in how they find the appropriate weights to apply to the data. ICA in particular has been proposed to remove artifacts including myogenic (Delorme et al., 2007; Ma et al., 2012; Crespo-Garcia et al., 2008).

The appeal of these linear techniques is that a noise source x_i can potentially be perfectly canceled: a component y_k is insensitive to source i as long as $\sum_j u_{ij} w_{jk} = 0$, where u_{ij} are the mixing coefficients and w_{jk} the component weights (Eqs. (1) and (2)). With high-dimensional data (lots of channels) there is considerable flexibility in satisfying this constraint, and the strength of analysis algorithms such as ICA lies in their ability to find such sets of weights. However, if a noise source is specific to a single channel j, it can only be cancelled by setting w_{jk} to zero for every component k, effectively discarding that channel. In this situation, component analysis offers little over the time-honored practice of discarding noisy channels.

Component analysis itself is vulnerable to channel-specific noise because it relies on the dimensionality of the data (determined by the number of channels) to resolve the various sources. If channels are discarded due to artifacts, analysis may be impaired, whereas if they are *not* discarded (due to lack of knowledge or the need to conserve enough channels), the artifact is injected into the extracted components via Eq. (2). The presence of artifacts may also interfere with the ability of the algorithm to find the optimal w_{jk}. For example an algorithm such as CSP (Koles et al., 1990), that searches for components that differ in power between two intervals, may lock on to an artifact that is present in one interval but not the other. Finally, the artifacts may interfere with the ability to estimate the topography associated with cortical activity, possibly compromising source modeling. Channel-specific noise thus limits the ability of linear methods to improve the signal to noise ratio (SNR) of weak brain activity.

These considerations lead us to focus on channel-specific noise, leaving other techniques such as ICA, JD, or beamforming to deal with noise sources that impinge on multiple electrodes. This is important scientifically, to obtain a more accurate picture of brain activity, and also for applications such as brain-computer interfaces (BCI), prediction of epileptic seizures, wearable brain-monitoring devices, and so on. The method described here is effective, fully automatic, and amenable to an online implementation for applications that involve realtime monitoring.

2. Methods

2.1. Signal model and assumptions

Each observation $x_j(t)$ is the sum of signals from multiple sources i within the brain or the environment (Eq. (1)). We make several restrictive assumptions. Each noise source $n_i(t)$ affects only one particular channel (*assumption 1*). Noise activity is temporally sparse so that artifacts on different channels do not temporally overlap (*assumption 2*), and for a significant proportion of time the data are artifact-free (*assumption 3*). Finally, we assume that in the absence of artifacts the data are *linearly dependent* such that each channel belongs to the subspace spanned by the other channels. In other words for each channel j there exist $a_{jj'}$ such that $x_j(t) = \sum_{j' \neq j} a_{jj'} x_{j'}(t)$ (*assumption 4*). This is plausible for neurogenic activity due to source-to-sensor mixing. Of course, many kinds of noise do *not* meet these assumptions; the focus here is on those that do. In real data, these assumptions will be met only as an *approximation*, for example because of non-stationarity of the brain and noise processes underlying the data. The quality of the outcome depends on the quality of the approximation.

2.2. The STAR algorithm

The algorithm proceeds in two phases. The first phase detects the presence of channel-specific artifacts, the second phase corrects them.

Phase 1. The covariance matrix of the data is estimated, and from it is calculated the matrix **A** that projects each channel on the subspace spanned by the other channels. The projection $\bar{x}_j(t)$ of channel j is the weighted sum of the channels $j' \neq j$ that best approximates $x_j(t)$. In the absence of an artifact we should have $\bar{x}_j(t) - x_j(t) = 0$ as a result of the linear dependence assumption, so a significant deviation indicates the presence of an artifact. Values of $\bar{x}_j(t) - x_j(t)$ are fit by a zero-mean Gaussian distribution, and values eccentric from this distribution (relative to a predefined threshold θ) are flagged as artifactual. This is repeated for all channels, and the union of eccentric time samples is labeled as artifact-contaminated. The covariance matrix is initially estimated from the entire data, and subsequently reestimated on the part of data labeled as artifact-free. A few iterations of this process lead to a stable partition of the time axis between artifact-free and artifact-contaminated parts.

Phase 2. The artifact-contaminated part is further divided according to which channel is most degraded at each time sample. For this purpose, a second eccentricity measure is calculated for each channel as the ratio of instantaneous power to power averaged over the artifact-free portion. The channel with the highest score at a given time sample "owns" that sample, and its data are replaced with its projection on the subspace spanned by the other channels, using projection coefficients calculated from the artifact-free part. Data replacement occurs only at the part corresponding to the artifact: at other times the data are left intact, so most of the data remain untouched by the processing.

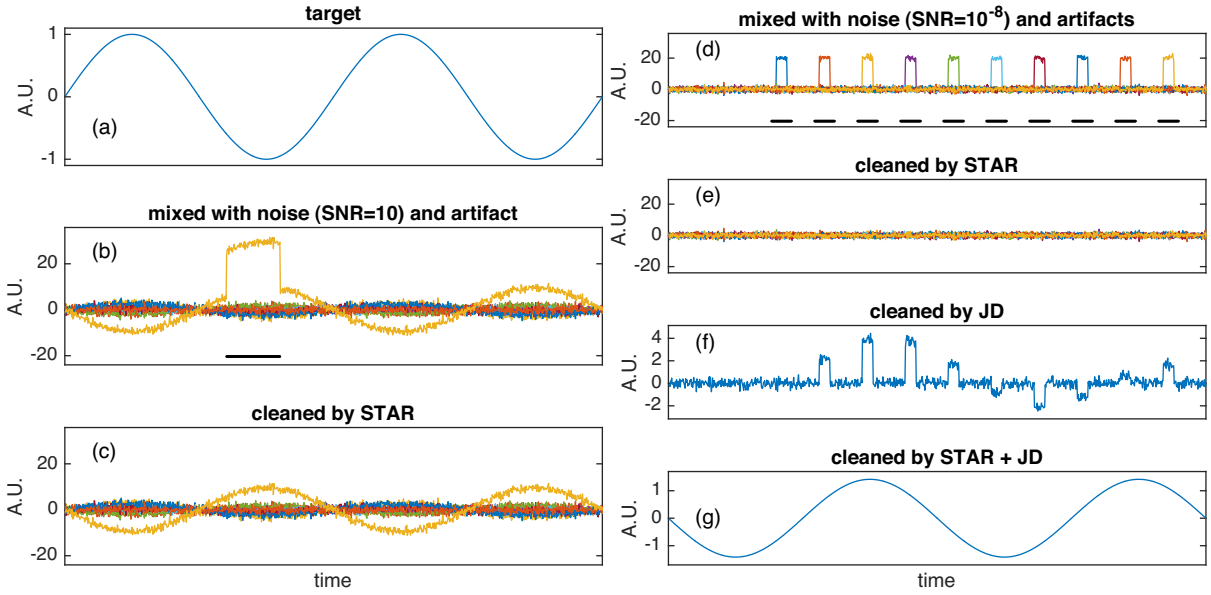

Fig. 1. Denoising simulated data. (a) Target signal. (b) 10-channel data including target, Gaussian noise (SNR = 10), and an artifact on one channel. The black bar indicates STAR's estimate of the "artifact-contaminated" part. (c) Data after removal of the artifact by the STAR algorithm. (d) 10-channel data including the same target and Gaussian noise at a much lower SNR (10^{-8}) and with transient artifacts on all channels. The black bars indicate STAR's estimate of the "artifact-contaminated" part. (e) Same after removal of artifacts by the STAR algorithm. (f) Result of applying the JD algorithm to the raw data in (d), with the aim of recovering the target. (g) Same, but the STAR algorithm was first applied before the JD algorithm.

2.3. Implementation details

A Matlab implementation of the algorithm is available in the NoiseTools toolbox (http://audition.ens.fr/adc/NoiseTools/). In the projection steps, to avoid issues with rank-deficient data, PCA is applied and PCs with power below a threshold are discarded. The algorithm operates on a sample-by-sample basis, each channel switching between an artifact-free state (no projection) and an artifact-contaminated state (projection). To limit the number of switches, the eccentricity measure is temporally smoothed by convolution with a short triangular window.

The algorithm thus has 4 parameters: the eccentricity threshold θ (the same value is used in phases 1 and 2), the size of the smoothing kernel in phase 2, the number of neighbors used to fix each channel, and the truncation threshold for the PCA. Of these, only the first is critical, as it determines the proportion of time samples that will be labeled as artifact-contaminated. An excessively low threshold leads to insufficient "artifact-free" data to reliably determine its correlation structure. For convenience, our implementation is designed to automatically increment the threshold (by a factor 1.1) if this occurs.

To save computation and reduce overfitting, channels may be projected on a subset of neighbouring channels, rather than on the full set of channels. On the assumption that cortical sources of interest should affect groups of physically close channels, this subset should suffice. 'Neighborhood' can be quantified based on geometrical proximity, or by measuring between-channel correlation in the raw data. To improve noise rejection it may be useful to apply STAR repeatedly, so that each iteration removes artifacts missed by the previous iteration.

2.4. Relaxing assumption 2

Assumption 2 is often violated for myogenic noise that may occur in bursts affecting multiple motor units (Yilmaz et al., 2014). We can relax it to allow several channels $[j_C]$ to be contaminated, as long as we modify assumption (4) to assume that each of the contaminated channels $[j_C]$ can be expressed as the weighted sum of the remaining non-contaminated channels. The algorithm is modified to project each of $[j_C]$ on the subspace spanned by the non-contaminated channels. This modification is computationally expensive (the number of cases to consider is exponential in the number of $[j_C]$), and in our implementation it is only approximated, see below.

3. Results

3.1. Simulated data

A sinusoidal target and several Gaussian noise sources were mixed together to simulate multichannel data, and channel-specific artifacts were added. The sinusoidal target signal (Fig. 1(a)) was mixed into the 10-channel data via a 1×10 mixing matrix with random coefficients, together with 8 independent Gaussian noise sources mixed via a 8×10 random mixing matrix. Coefficients were chosen to ensure SNR = 10, and a pulse-like artifact was added to one channel (Fig. 1(b)). Applying STAR effectively suppresses the artifact (Fig. 1(c)), apparently without distorting the signal. The data remain contaminated by the other noise sources, although at this SNR the target still emerges. Fig. 1(d) shows a similar mixture with SNR = 10^{-8} and one artifact on each channel, and Fig. 1(e) shows the result of applying STAR. The black bars in Fig. 1(b, d) indicate STAR's estimate of the artifact-contaminated part. The artifacts are removed, but at this SNR the target is invisible.

A linear algorithm such as JD (de Cheveigné and Parra, 2014) can be used to extract a weak target such as Fig. 1(a) despite a very unfavorable SNR. However this requires the dimensionality of the noise to be smaller than that of the data, a requirement that fails in the presence of the channel-specific artifacts in addition to the Gaussian noise as in Fig. 1(d). Indeed, applying JD to the artifact-contaminated data fails to recover the target (Fig. 1(f)). However, if STAR is first applied to remove the artifacts (Fig. 1(e)), JD can successfully recover the target (Fig. 1(g)). This example shows that STAR can effectively remove channel-specific artifacts, that it can do so for artifacts affecting multiple channels, and that such removal

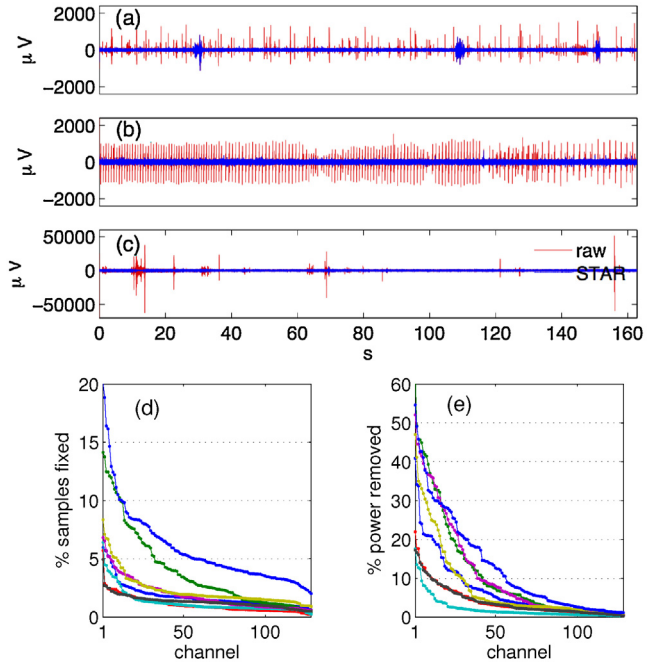

Fig. 2. 128-channel EEG data. (a–c) Examples of single-channel signals from a 128-channel EEG recording before (red) and after (blue) applying STAR. (d) Percentage of samples removed for each channel (sorted by decreasing percentage), for each subject. (e) Proportion of power removed for each channel (sorted by decreasing power), for each subject. (For interpretation of the references to color in this legend, the reader is referred to the web version of the article.)

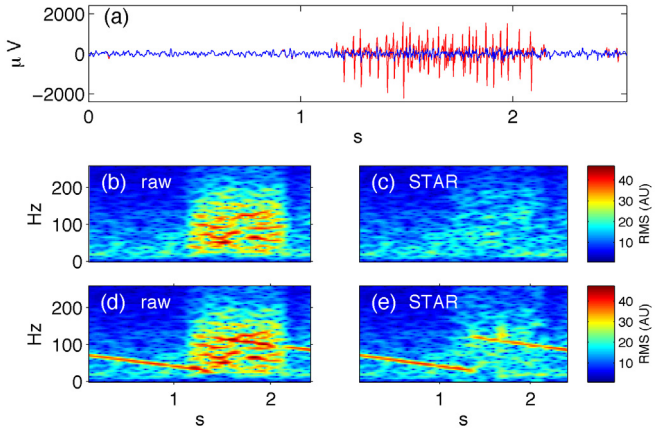

Fig. 3. 128-channel EEG data. (a) One channel with muscle artifact before (red) and after (blue) processing. (b, c) Spectrograms of the signal before and after applying STAR. (d, e) Same, with superimposed chirp signal, illustrating that processing can reveal a signal otherwise masked by the artifact. (For interpretation of the references to color in this legend, the reader is referred to the web version of the article.)

may be needed for a component analysis technique such as JD to succeed.

3.2. EEG data

A first example involves data recorded with a 128-channel EEG system (Biosemi). The data sampled at 512 Hz were high-pass filtered with 20 Hz cutoff (2nd order Butterworth) to emphasize the gamma region that is relatively susceptible to artifacts. The STAR algorithm was applied with threshold parameter $\theta = 2$ to suppress channel-specific artifacts. The size of the smoothing window applied to the eccentricity measure was 19 samples. Fig. 2(a–c) shows three example signals from individual electrodes before (red) and after (blue) processing. In Fig. 2(a) the data are contaminated by what seems to be muscle activity. This is suppressed by STAR (the three remaining bursts reflect activity correlated across sensors and therefore not removed). In Fig. 2(b) the data are contaminated by a pulsatile noise with a period of approximately 1 s, most likely a "pulsation artifact" due to a vein underlying the electrode (Niedermeyer and Lopez da Silva, 2005). In Fig. 2(c) the data are contaminated by a series of glitches of high amplitude (note the scale) that are equally well suppressed.

Data from this study were available for 8 subjects, about 1.5–2.5 h each. Fig. 2(d) shows the percentage of samples corrected by STAR for each channel (channels sorted by decreasing percentage) and subject, cumulated over the recording session. The percentage varies across subjects and channels, reflecting a different prevalence of channel-specific artifacts between subjects, and for a given subject, between channels. For the worst subject the percentage is below 5% for most channels, implying that >95% of the data of those channels were untouched by the processing. For good subjects the percentage of intact samples is >98% for almost all channels. These numbers suggest that the *negative* impact of applying STAR, if any, is likely to be minimal. Fig. 2(e) shows the percentage of power removed by STAR for each channel (channels sorted by decreasing percentage) and each subject (individual lines). Again the percentage varies across subjects and channels. For the worst subject, >20% of power was removed for >30 channels, and >55% for the worst channel. For other subjects the percentages are smaller, but still significant for a subset of channels. To summarize, STAR is both safe and of potentially significant benefit.

Fig. 3(a) shows one channel contaminated by a burst of pulsatile activity, likely muscle artifact, before (red) and after (blue) applying STAR. The spectrogram of these data is shown in Fig. 3 before (b), and after (c) processing. Artifactual activity dominates the higher spectral region, but this is alleviated by processing. Fig. 2(d) and (e) illustrates how this processing could potentially reveal high frequency activity (here a synthesized narrowband chirp) otherwise masked by the artifact.

A second example uses data recorded from three closely-spaced EEG electrodes within the ear canal of a subject (Kidmose et al., 2012). The rationale for recording from within the ear is unobtrusiveness (for hearing aid-related applications) and the hope that muscle artifacts will be weaker, however the close spacing between electrodes is expected to result in signals highly correlated across channels. Fig. 4(a) shows a short segment of signals from the three electrodes. Contrary to expectations, one channel (green) is affected by channel-specific glitches of possible muscular origin. Applying STAR removes these artifacts resulting in signals that are much more similar across channels, as expected from collocated electrodes (Fig. 4(b)). Another way to attenuate channel-specific artifacts is to average across channels, but the benefit here is relatively small (Fig. 4(c), red). However applying STAR before averaging greatly increases this benefit (blue). The deflection near 0.5 s is not removed because it is shared across sensors. The effect of artifact reduction of the 3-channel data is reflected in its PCA spectrum (Fig. 4(d)). Cortical signals from closely spaced electrodes should be highly correlated, implying a rapidly decreasing PCA spectrum, and this expectation is better met after applying STAR. This example shows that STAR can be effective even for a small number of channels.

3.3. Simulated target embedded in real EEG

This example involves a simulated repetitive target (half-sinusoid pulse) superimposed on the same 3-channel EEG signal as in the previous example. The advantage of using a simulated target is that the ground truth is known, the advantage of using real EEG as noise is that it is representative of real EEG noise. The

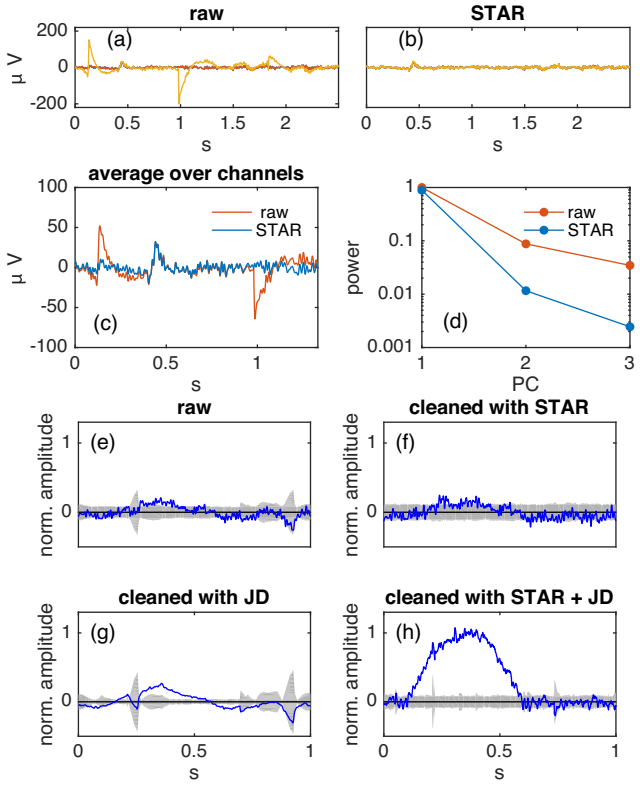

Fig. 4. Three-channel EEG. (a) Raw data. (b) Same after artifact removal by STAR. (c) Average over channels of raw (red) or processed (blue) data. (d) Eigenvalue spectrum of raw (red) and processed (blue) data. The smaller power of PCs 2 and 3 indicates that the signals are closer to being identical across channels, as expected for three collocated electrodes. (e, f) Simulated target (290 repetitions of a half-sinusoid) embedded in 3-channel EEG noise (SNR = 0.03). Blue is average over trials, gray band indicates ±2SD of a bootstrap resampling of the mean. (e) Average over trials of the "best" electrode. (f) Same as (e) after applying the STAR algorithm to suppress artifacts. (g) Same as (e) after applying the JD algorithm to enhance the target. (h) Same as (e) after applying first STAR then JD. The values plotted are normalized to equate the variance over trials in each condition, thus the greater amplitude of the mean (blue) in (h) reflects the better SNR of the recovered signal. (For interpretation of the references to color in this legend, the reader is referred to the web version of the article.)

target pulse was repeated 290 times and added to the second channel of the EEG with SNR = 0.03. Fig. 4(e) shows the average over trials of the second EEG channel. The gray band represents ±2SD of a bootstrap resampling of the mean. The presence of artifacts is reflected by the irregular width of the gray band. In principle, the SNR of such multichannel data can be enhanced by the JD algorithm (de Cheveigné and Parra, 2014) with the appropriate bias function to enhance repeatable components. However in this case the benefit is limited (Fig. 4(g)) presumably because all channels are artifact-contaminated so that there is no artifact-free subspace. Applying the STAR algorithm attenuates these artifacts (Fig. 4(f)) and JD produces a well-isolated target signal (Fig. 4(h)). This example shows again that removing channel-specific artifacts with STAR can improve the effectiveness of component analysis.

4. Discussion

The STAR algorithm exploits the between-channel correlation of the underlying data to identify channel-specific artifacts and repair them. Figuratively, it "surfs" the correlation structure of the underlying data, selectively shaving off the parts that do not fit. Channel-specific artifacts are ubiquitous in EEG and other recording techniques. Removing them is usually tedious (if done manually) or complex and unreliable (if done automatically), and a reliable automatic method to suppress them is thus useful.

4.1. Comparison with other approaches

The most common approach to channel-specific artifacts is to discard the offending channel(s), or else the offending samples or trials (Picton et al., 2000). Valid data are lost as a result, and the strategy fails if too many channels and/or samples are affected. In contrast, STAR entails minimal data loss, because only a small proportion of samples are discarded. Data are rejected on a sample-by-sample basis, in contrast for example to SCADS (Junghöfer et al., 2000) that removes larger chunks (trials) of the affected channel, or AFOP (Boudet et al., 2012) that operates on a relatively long window (e.g. 20 s).

As in several other approaches, missing channels are interpolated spatially. In STAR the interpolation is based on an estimate of the correlation structure gathered over the artifact-free part, in contrast to methods such as SCADS that apply spherical spline interpolation to reconstruct the missing data (Perrin et al., 1989). STAR thus requires no geometrical information (it is purely data-driven) although it may benefit from using a proximity map to restrict the number of channels upon which each channel is projected. STAR uses the multivariate correlation structure of the data, in contrast to wavelet or empirical mode basis (Safieddine et al., 2012) that use the *within-channel* autocorrelation structure to interpolate missing values.

A very different approach is to use component analysis to form linear combinations of channels as in Eq. (1), and project the components that correspond to artifacts out of the data. There are many techniques available (ICA, beamforming, JD, etc.) that differ in the way they determine the weights w_{jk}. In ICA (Delorme et al., 2007; Ma et al., 2012; Crespo-Garcia et al., 2008) a measure of statistical independence (e.g. kurtosis) is used to constrain the weights. Channel-specific noise is usually both kurtotic (lots of zero or extreme values) and independent between channels, so this procedure is likely to isolate components that contain channel-specific noise. In beamforming (Grosse-Wentrup et al., 2009) a spatial filter is used to isolate sources, either desired (to extract them) or undesired (to suppress them), on the basis of their spatial position. In CSP (Koles et al., 1990) or JD (de Cheveigné and Parra, 2014) the spatial filter is instead designed to maximize or minimize some desirable or undesirable trait of the data.

Regardless of the method, component analysis faces a fundamental limitation: it cannot resolve more sources than dimensions in the data. STAR is not subject to this limitation, as illustrated by the example in Fig. 1(d–g) where artifacts were present on all channels, in addition to target and noise sources, but the target could nonetheless be separated. Neither ICA, nor JD, nor beamforming can deal with this situation because the number of sources to resolve exceeds the number of channels. The reason why STAR succeeds is that it operates locally in time (so that the solution applied to fix each sample need not consider other samples), and leaves most of the data intact (the dimensionality of the cleaned data is that of the data before artifacts were added). In contrast, techniques such as ICA apply a single linear transform to all the data, each artifact removed reduces the data dimensionality, and there may not be enough dimensions to allow them all to be removed. As stressed in Section 1, an artifact specific to channel j can only be removed by setting all the w_{jk} to zero in Eq. (2), i.e. discarding that electrode, and therefore if artifacts are present on many channels the dimensionality is greatly reduced. Empirically, the effectiveness of ICA as a tool to reduce myogenic artifacts has been questioned because ICs may be found to contain both artifact and neural activity (McMenamin et al., 2010, 2011), presumably as a consequence of the large number of distinct generators of muscle artifact.

Similar to STAR, the sensor noise suppression (SNS) method (de Cheveigné and Simon, 2008b) addresses sensor-specific noise by replacing each channel by its projection on the subspace spanned by the other channels. However with SNS this operation is applied uniformly across the time axis, and thus shares the same drawbacks of other component analysis techniques.

The AFOP method of Boudet et al. (2012) applies different solutions to different parts of the time axis (or time-frequency plane). Artifacts within a limited spectrotemporal window are isolated using a variant of the common spatial pattern (CSP) algorithm (Koles et al., 1990), and the data over that window are fixed by projecting them out. AFOP can remove any activity specific to that window, not just channel-specific, in contrast to STAR which only removes channel-specific artifacts. However as noted above, AFOP operates at a coarser temporal granularity.

The focus of STAR on channel-specific artifacts could be seen as a limitation. However this restriction allows it to avoid removing activity shared across channels, of deeper and possibly neural origin. Most importantly, STAR is compatible with other techniques (ICA or other) that *can* deal with channel-shared artifacts, increasing their effectiveness by relieving them of the burden of channel-specific artifacts. Indeed, the primary motivation for developing STAR was to push the limits of these noise reduction techniques, including our own previous efforts (de Cheveigné and Simon, 2007, 2008a,b; de Cheveigné, 2010, 2012; de Cheveigné and Parra, 2014).

An feature of STAR is that it is automatic, in contrast to traditional artifact rejection that requires manual intervention, or to ICA that requires inspection of components to decide which to remove. It is relatively parameter-free, in contrast to methods such as FASTER (Nolan et al., 2010) that involve multiple parameters. It is also conceptually simple and computationally fast.

To summarize the main features that distinguish STAR from other approaches: (a) it is automatic and requires little tuning, (b) dimensionality is not reduced and apart from the artifact portions the data are untouched, (c) the method addresses only channel-specific artifacts, leaving other forms of artifact to other approaches, that it complements and can enhance.

4.2. Caveats and cautions

Assumption 4 requires that the dimensionality of the activity underlying the data (e.g. number of significant neural sources) be less than the number of channels. As there are potentially billions of neural sources this can only be an approximation.

The correlation structure is assumed to be stationary so that the projection parameters estimated from the artifact-free part can be applied in the artifact-contaminated part. A departure from this assumption (for example if additional neural sources coincided with muscle artifact) would reduce effectiveness.

The calculation of the projection matrix, that allows each channel to be replaced by its projection on the other channels, is subject to overfitting if there is a large number of channels or insufficient data. It is useful to limit the projection to a subset of neighboring channels (10 seems to work well).

Assumptions 2 and 3 require that channel-specific artifacts be local in time, which excludes slow channel-specific fluctuations. It may be useful to remove such slow fluctuations by high-pass filtering the data before applying STAR (or split the data into low- and high-pass streams, process the latter, then recombine).

For each channel, the algorithm switches between an artifact-free state, for which the data are untouched, and an artifact-contaminated state for which the data are replaced by their projection. The transition between states may produce a step, and this may paradoxically lead to increased high-frequency noise if the transitions are numerous and/or the steps are large. Smoothing the eccentricity measure reduces the number of switches, and in practice switching seems rarely a problem unless the data contain a lot of low-frequency power (see previous point).

Assumption 2, that artifacts should occur on only one channel at a time, is often violated for myogenic activity that may occur in bursts that affect several channels simultaneously. This has two consequences. First, the algorithm (in its basic form) only fixes one channel at a given time point, and so artifacts on the other channels may remain. Second, projection on artifact-contaminated channels may contaminate the channel being fixed. The extension described in Section 2.4 addresses the problem, although its applicability is limited by the computational cost of dealing with the many combinations of artifact-contaminated channels (exponential in their number). The NoiseTools toolbox implements an approximation that appears to be effective. STAR may be applied repeatedly, so that each new pass removes artifacts missed by previous passes.

4.3. Applicability to electromyogenic activity

There is evidence that much EEG power beyond 20 Hz is muscular in origin: when muscular activity is silenced by injection of a paralysing drug, power in that band is greatly diminished (Whitham et al., 2007). The fact that muscular activity covaries with mental state makes it a potential confound in cognitive studies (Whitham et al., 2008; Yuval-Greenberg et al., 2008; Shackman et al., 2009). The STAR algorithm offers a partial solution, in that it removes the type of myogenic activity that is local to one electrode (Ma et al., 2012). This constitutes a subset of myogenic activity: activity that is instead correlated across electrodes (e.g. from deeper or distant muscles) must be removed by other means (McMenamin et al., 2010; Yilmaz et al., 2014). Empirically, the amount of reduction observed for each artifact burst is variable, ranging from 100% to a few tens of percent (e.g. Fig. 3). Artifacts that are correlated across electrodes (e.g. from ocular or deep muscles) may be addressed by component analysis such as ICA or JD, to which STAR is complementary in that it avoids squandering dimensions useful to those methods. The STAR algorithm is thus of potential benefit in studies of cortical activity at higher frequencies.

4.4. Applicability to real-time and BCI applications

The algorithm can operate online in real time. After an initial bootstrap phase (to get an estimate of the artifact-free covariance matrix), the algorithm switches between two states: (1) no artifact, the covariance matrix is updated, (2) artifact, the offending channels are corrected. Updating of the covariance matrix estimate may occur for example according to a leaky integrator mechanism, to allow for slow variations of the correlation structure.

Applications such as BCI involving a wearable EEG device are limited by artifacts related to movement, including contact and muscle artifacts (Fatourechi et al., 2007; McFarland et al., 2005; Reis et al., 2014) that can for example disrupt an otherwise effective technique such as CSP (Grosse-Wentrup et al., 2009). Manual correction or offline processing are obviously ruled out, and STAR thus is of potential benefit for these applications.

5. Conclusions

The sparse time artifact reduction (STAR) algorithm deals effectively with channel-specific artifacts that are common in EEG, including electrode contact noise and certain types of muscle artifact. It operates locally in time, correcting only the samples affected by the artifact, and leaving the remainder intact, and thus has minimal negative impact on the underlying data. Removal of an artifact does not reduce the dimensionality of the data, in contrast to channel rejection or component analysis techniques such as ICA. The

STAR algorithm can complement those techniques and enhance their effectiveness by relieving them of the burden of dealing with channel specific artifacts. The STAR algorithm operates automatically, requires few parameters, is computationally efficient, and can be implemented to perform online processing in real time.

Acknowledgements

This work was supported by the EU H2020-ICT grant 644732 (COCOHA), and grants ANR-10-LABX-0087 IEC and ANR-10-IDEX-0001-02 PSL*. Ed Lalor and Giovanni Liberto contributed the 128-channel EEG data, and Thomas Lunner and Carina Graverson contributed the 3-channel EEG data.

References

Boudet S, Peyrodie L, Forzy G, Pinti A, Toumi H, Gallois P. Improvements of adaptive filtering by optimal projection to filter different artifact types on long duration EEG recordings. Comput Methods Prog Biomed 2012];108:234–49.

Corby JC, Roth WT, Kopell BS. Prevalence and methods of control of the cephalic skin potential EEG artifact. Psychophysiology 1974];11:350–60.

Crespo-Garcia M, Atienza M, Cantero JL. Muscle artifact removal from human sleep EEG by using independent component analysis. Ann Biomed Eng 2008];36:467–75.

de Cheveigné A. Time-shift denoising source separation. J Neurosci Methods 2010];189:113–20.

de Cheveigné A. Quadratic component analysis. NeuroImage 2012];59:3838–44.

de Cheveigné A, Parra LC. Joint decorrelation, a versatile tool for multichannel data analysis. NeuroImage 2014];98:487–505.

de Cheveigné A, Simon JZ. Denoising based on time-shift PCA. J Neurosci Methods 2007];165:297–305.

de Cheveigné A, Simon JZ. Denoising based on spatial filtering. J Neurosci Methods 2008a];171:331–9.

de Cheveigné A, Simon JZ. Sensor noise suppression. J Neurosci Methods 2008b];168:195–202.

Delorme A, Sejnowski T, Makeig S. Enhanced detection of artifacts in EEG data using higher-order statistics and independent component analysis. NeuroImage 2007];34:1443–9.

Fatourechi M, Bashashati A, Ward RK, Birch GE. EMG and EOG artifacts in brain computer interface systems: a survey. Clin Neurophysiol 2007];118:480–94.

Goncharova II, McFarland DJ, Vaughan TM, Wolpaw JR. EMG contamination of EEG: spectral and topographical characteristics. Clin Neurophysiol 2003];114:1580–93.

Grosse-Wentrup M, Liefhold C, Gramann K, Buss M. Beamforming in noninvasive brain–computer interfaces. IEEE Trans Biomed Eng 2009];56:1209–19.

Junghöfer M, Elbert T, Tucker DM. Statistical control of artifacts in dense array EEG/MEG studies. Psychophysiology 2000];37:523–32.

Kidmose P, Looney D, Mandic DP. Auditory evoked responses from Ear-EEG recordings. In: International Conference of the IEEE Engineering in Medicine and Biology Society. Proceedings 2012; 2012]. p. 586–9.

Koles Z, Lazar M, Zhou S. Spatial patterns underlying population differences in the background EEG. Brain Topogr 1990];2:275–84.

Ma J, Tao P, Bayram S, Svetnik V. Muscle artifacts in multichannel EEG: characteristics and reduction. Clin Neurophysiol 2012];123:1676–86.

McFarland DJ, Sarnacki WA, Vaughan TM, Wolpaw JR. Brain–computer interface (BCI) operation: signal and noise during early training sessions. Clin Neurophysiol 2005];116:56–62.

McMenamin BW, Shackman AJ, Maxwell JS, Bachhuber DRW, Koppenhaver AM, Greischar LL, Davidson RJ. Validation of ICA-based myogenic artifact correction for scalp and source-localized EEG. NeuroImage 2010];49:2416–32.

McMenamin BW, Shackman AJ, Greischar LL, Davidson RJ. Electromyogenic artifacts and electroencephalographic inferences revisited. NeuroImage 2011];54:4–9.

Muthukumaraswamy SD. High-frequency brain activity and muscle artifacts in MEG/EEG: a review and recommendations. Front Hum Neurosci 2013];7:138.

Niedermeyer E, Lopez da Silva F. Electroencephalography: basic principles, clinical applications and related fields. Lippincott Williams & Wilkins; 2005].

Nolan H, Whelan R, Reilly RB. FASTER: fully automated statistical thresholding for EEG artifact rejection. J Neurosci Methods 2010];192:152–62.

Perrin F, Pernier J, Bertrand O, Echallier JF. Spherical splines for scalp potential and current density mapping. Electroencephalogr Clin Neurophysiol 1989];72:184–7.

Picton TW, Bentin S, Berg P, Donchin E, Hillyard SA, Johnson R, Miller GA, Ritter W, Ruchkin DS, Rugg MD, Taylor MJ. Guidelines for using human event-related potentials to study cognition: recording standards and publication criteria. Psychophysiology 2000];37:127–52.

Reis PMR, Hebenstreit F, Gabsteiger F, von Tscharner V, Lochmann M. Methodological aspects of EEG and body dynamics measurements during motion. Front Hum Neurosci 2014];8:156.

Safieddine D, Kachenoura A, Albera L, Birot G, Karfoul A, Pasnicu A, Biraben A, Wendling F, Senhadji L, Merlet I. Removal of muscle artifact from EEG data: comparison between stochastic (ICA and CCA) and deterministic (EMD and wavelet-based) approaches. Eurasip J Adv Signal Process 2012];2012:127.

Shackman AJ, McMenamin BW, Slagter HA, Maxwell JS, Greischar LL, Davidson RJ. Electromyogenic artifacts and electroencephalographic inferences. Brain Topogr 2009];22:7–12.

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. Scalp electrical recording during paralysis: quantitative evidence that EEG frequencies above 20 Hz are contaminated by EMG. Clin Neurophysiol 2007];118:1877–88.

Whitham EM, Lewis T, Pope KJ, Fitzgibbon SP, Clark CR, Loveless S, DeLosAngeles D, Wallace AK, Broberg M, Willoughby JO. Thinking activates EMG in scalp electrical recordings. Clin Neurophysiol 2008];119:1166–75.

Yilmaz G, Ungan P, Sebik O, Uginčius P, Türker KS. Interference of tonic muscle activity on the EEG: a single motor unit study. Front Hum Neurosci 2014];8:504.

Yuval-Greenberg S, Tomer O, Keren AS, Nelken I, Deouell LY. Transient induced gamma-band response in EEG as a manifestation of miniature saccades. Neuron 2008];58:429–41.