Filters: When, Why, and How (Not) to Use Them

Alain de Cheveigné, Israel Nelken

Filters: when, why, and how (not) to use them

Alain de Cheveigné (1, 2, 3), Israel Nelken (4)

(1) Laboratoire des Systèmes Perceptifs, UMR 8248, CNRS, France.
(2) Département d'Etudes Cognitives, Ecole Normale Supérieure, PSL, France.
(3) UCL Ear Institute, United Kingdom.
(4) Edmond and Lily Safra Center for Brain Sciences and the Silberman Institute of Life Sciences, Hebrew University Jerusalem, Israel.

LEAD CONTACT / CORRESPONDING AUTHOR:
Alain de Cheveigné, Audition, DEC, ENS, 29 rue d'Ulm, 75230, Paris, France, Alain.de.Cheveigne@ens.fr, phone 0033144322672, 00447912504027.

Emails:
Alain.de.Cheveigne@ens.fr israel.nelken@mail.huji.ac.il
approximate number of characters (with spaces): 79000

Summary

Filters are commonly used to reduce noise and improve data quality. Filter theory is part of a scientist's training, yet the impact of filters on interpreting data is not always fully appreciated. This paper reviews the issue, explains what is a filter, what problems are to be expected when using them, how to choose the right filter, or how to avoid filtering by using alternative tools. Time-frequency analysis shares some of the same problems that filters have, particularly in the case of wavelet transforms. We recommend reporting filter characteristics with sufficient details, including a plot of the impulse or step response as an inset.

Keywords

filter, artifact, Fourier analysis, time-frequency representation, ringing, causality, impulse response, oscillations

1. Introduction

One of the major challenges of brain science is that measurements are contaminated by noise and artifacts. These may include environmental noise, instrumental noise, or signal sources within the body that are not of interest in the context of the experiment ("physiological noise"). The presence of noise can mask the target signal, or interfere with its analysis. However, if signal and interference occupy different spectral regions, it may be possible to improve the signal-to-noise ratio (SNR) by applying a *filter* to the data.

For example, a DC component or slow fluctuation may be removed with a high-pass filter, power line components may be attenuated by a notch filter at 50 Hz or 60 Hz, and unwanted high-frequency components may be removed by "smoothing" the data with a low-pass filter. Filtering takes advantage of the difference between spectra of noise and target to improve SNR, attenuating the data more in the spectral regions dominated by noise, and less in those dominated by the target.

Filters are found at many stages along the measurement-to-publication pipeline (Fig. 1). The measuring rig or amplifier may include a high-pass filter and possibly a notch filter, the analog-to-digital (AD) converter is preceded by a low-pass antialiasing filter, preprocessing may rely on some combination of high-pass, lowpass and notch filters, data analysis may include bandpass or time-frequency analysis, and so on. Filters are ubiquitous in brain data measurement and analysis.

[Figure 1 about here]

The improvement in SNR offered by the filter is welcome, but filtering affects also the target signal in ways that are sometimes surprising. Obviously, any components of the target signal that fall within the stop band of the filter are lost. For example applying a 50 Hz notch filter to remove power line artifact might also remove brain activity within the 50 Hz region. The experimenter who blindly relies on the filtered signal is blind to features suppressed by the filter.

Harder to appreciate are the distortions undergone by the target. Such distortions depend on the frequency characteristics of the filter, including both amplitude and phase characteristics (which are often not reported). The output of a filter is obtained by *convolution* of its input with the impulse response of the filter, which is a fancy way of saying that each sample of the output is a weighted sum of several samples of the input. Each sample therefore depends on a whole segment of the input, spread over time. Temporal features of the input are smeared in the output, and conversely "features" may appear in the output that were not present in the input to the filter.

We first explain what is a filter in detail, and how filters are involved in data analysis. Then we review the main issues that can arise, and make suggestions on how to fix them. Importantly, similar issues occur also in time-frequency analyses, such as spectrograms and wavelet transforms, which are based on a collection of filters (a filterbank). Finally, we list a number of recommendations that may help investigators identify and minimize issues related to the use of filters, and we suggest ways to report them so that readers can make the best use of the information that they read. In this paper, "filter" refers to the familiar one-dimensional convolutional filter (e.g. high-pass or band-pass) applicable to a single-channel waveform, as opposed to "spatial filters" applicable to multichannel data.

2. What is a filter?

For many of us, a filter is "a thing that modifies the spectral content of a signal". For the purposes of this paper, however, we need something more precise. A filter is an operation that produces each sample of the output waveform y as a weighted sum of several samples of the input waveform x. For a digital filter:

$$y(t) = \sum_{n=0}^{N} h(n)x(t-n) \tag{1}$$

where t is the analysis point in time, and $h(n), n = 0, \dots, N$ is the impulse response. This operation is called convolution.

[Figure 2 about here]

We expect the reader to fall into one of three categories: (a) those who understand and feel comfortable with this definition, (b) those who mentally transpose it to the frequency domain where they feel more comfortable, and (c) those who remain mind-boggled. Categories (b) and (c) both need assistance, and that is what this section is about. Category (b) need assistance because a frequency-domain account is incomplete unless phase is taken into account, but doing so is mentally hard and often not so illuminating. It is often easier to reason in the time domain.

For the mind-boggled, the one important idea to retain is that every sample of the output depends on *multiple* samples of the input, as illustrated in Fig. 2 (top). Conversely, each sample $x(t)$ of the input impacts several samples $y(t+n)$ of the output (Fig. 2, bottom). As a result, the signal that is being filtered is smeared along the temporal axis, and temporal relations between filtered and original waveforms are blurred. For example, the latency between a sensory stimulus and a brain response, a straightforward notion, becomes less well defined when that brain response is filtered.

The exact way in which the output of a filter differs from its input depends upon the filter, i.e. the values $h(n)$ of the impulse response. Some filters may smooth the input waveform, others may enhance fast variations. There is a considerable body of theory, methods, and lore on how best to design and implement a filter for the needs of an application.

Expert readers will add that a filter is a *linear system*, that $h(n)$ is not expected to change over time (*linear time-invariant system*), that in addition to *causal* filters described by Eq. 1, there are *acausal* filters for which the series $h(n)$ includes also negative indices (gray lines in Fig. 2), that N may be finite (*finite impulse response, FIR*) or infinite (*infinite impulse response, IIR*). IIR filters are often derived from standard analogue filter designs (e.g. Butterworth or elliptic).

Essentially everything we discuss below is true for these more general notions of filtering. Expert readers will also recognize that Eq. 1 can be substituted by the simpler equation $Y(\omega) = H(\omega)X(\omega)$ involving the Fourier transforms of $x(t)$, $y(t)$ and $h(n)$, that neatly describes the effects of filtering in the frequency domain as a product of two complex functions, the transfer function of the filter $H(\omega)$ and the Fourier transform of the input, $X(\omega)$. The *magnitude transfer function* $|H(\omega)|$ quantifies the amount of attenuation at each frequency ω.

A special mention should be made of *acausal* filters. These are filters for which each sample of the output depends also on future samples of the input, i.e. we must modify Eq. 1 to include negative indices $n = -N', \dots, -1$. All physical systems must be causal (the future cannot influence the past) so this filter cannot represent a physical system, nor could it be implemented in a real time processing device. However, for offline data analysis we can take samples from anywhere in the data set, so in that context acausal filters are realizable. In particular, it is common to use *zero phase* filters, for which the impulse response is symmetrical relative to zero. The Matlab function `filtfilt` applies the same filter to the data twice, forward and backward, effectively implementing a zero-phase filter.

While acausal filters are easy to apply, interpreting their output requires special care. An important goal of neuroscience is to determine causal relations, for example between a stimulus and brain activity, or between one brain event and another, and we must take care that these relations are not confused by an acausal stage in the data analysis.

For an *IIR* filter, the output depends on all samples from the start of the data, previous samples being treated as 0. If the IIR filter is acausal it can also depend on all samples until the end of the data, samples beyond the end being treated as 0. This is also the case when filters are implemented in the *Fourier domain*: each output sample $y(t)$ depends potentially on all input samples $x(t)$ that are used to compute the Fourier transform, i.e. every sample within the analysis window.

[Figure 3 about here]

Figure 3 illustrates four common types of filter: low-pass, high-pass, band-pass, and notch (or band-reject). The upper plots show the magnitude transfer function (on a log-log scale) and the bottom plots show the impulse response of each filter. For high-pass and notch filters, the impulse response includes a one-sample impulse ("Dirac") of amplitude much greater than the rest (plotted here using a split ordinate). For each filter two versions are shown, one with with shallow (blue) and the other with steep (red) frequency transitions. Note that a filter with a steep transition in the frequency domain tends to have an impulse response that is extended in the time domain.

Also important to note is that different impulse responses can yield the same magnitude transfer function. Figure 4 (left) shows four impulse responses that all share the same magnitude frequency characteristic (low-pass, similar to that shown in Fig. 3) but differ in their phase characteristics (plotted on the right). Magnitude and phase together fully specify a filter (as does the impulse response). Among all the filters that yield the same magnitude frequency response, one is remarkable in that it is causal and has *minimum phase* over all frequency (thick blue). Another is remarkable in that it has *zero phase* over all frequency (thick green). It is acausal.

[Figure 4 about here]

3. Uses of filters

Antialiasing.. Ubiquitous, if rarely noticed, is the hardware "antialiasing" filter that precedes analog-to-digital conversion within the measuring apparatus. Data processing nowadays is almost invariably done in the digital domain, and this requires signals to be sampled at discrete points in time so as to be converted to a digital representation. Only values at the sampling points are retained by the sampling process, and thus the digital representation is ambiguous: the same set of numbers might conceivably reflect a different raw signal. The ambiguity vanishes if the raw signal obeys certain conditions, the best known of which is given by the *sampling theorem*: if the original signal's spectrum contains no power beyond the Nyquist frequency (one half the sampling rate) then it can be perfectly reconstructed from the samples. The antialiasing filter aims to enforce this condition ("Nyquist condition"). A hardware antialiasing filter is usually applied before sampling, and a software antialiasing filter may later be applied if the sampled data are further downsampled or resampled.

Smoothing / low-pass filtering.. Phenomena of interest often obey slow dynamics. In that case, high-frequency variance can safely be attributed to irrelevant noise fluctuations and attenuated by low-pass filtering. Smoothing is also often used to make data plots visually more palatable, or to give more emphasis on longer-term trends than on fine details.

High-pass filtering to remove drift and trends.. Some recording modalities such as electroencephalography (EEG) or magnetoencephalography (MEG) are susceptible to DC shifts and slow drift potentials or fields, upon which ride the faster signals of interest (Huigen et al., 2002; Kappenman and Luck, 2010; Vanhatalo et al., 2005). Likewise, in extracellular recordings, spikes of single neurons ride on slower events, such as negative deflections of the local field potential (LFP) that often precede spikes, or the larger and slower drifts due to the development of junction potentials between the electrode tip and the brain tissue. High-pass filters are the standard tool to remove such slow components prior to data analysis. A hardware high-pass filter might also be included in the measurement apparatus to remove DC components prior to conversion so as to make best use of the limited range of the digital representation. This is the meaning of "ac coupling" on an oscilloscope: it consists of the application of a high-pass filter - often implemented as a mere capacitor - to the signal. Amplifiers for recording extracellular brain activity are usually AC coupled.

Notch filtering.. Electrophysiological signals are often plagued with power line noise (50 or 60 Hz and harmonics) coupled electrically or magnetically with the recording circuits. While such noise is best eliminated at its source by careful equipment design and shielding, this is not always successful, nor is it applicable to data already gathered. Notch filtering is often used to mitigate such power line noise. Additional notches may be placed at harmonics if needed.

Band-pass filtering.. It has become traditional to interpret brain activity as coming from frequency bands with names such as *alpha*, *beta*, *theta*, etc., and data analysis often involves applying one or more band-pass filters to isolate particular bands, although the consensus is incomplete as to the boundary frequencies or the type of filter to apply.

Time-frequency analysis.. One prominent application of filtering is *time-frequency* (TF) analysis. A TF representation can be viewed as the time-varying magnitude of the data at the outputs of a filterbank. A filterbank is an array of filters that differ over a range of parameter values (e.g. center frequencies and/or bandwidths). The indices of the filters constitute the frequency axis, while the time series of their output magnitude unfolds along the time axis of the TF representation. The time-varying magnitude is obtained by applying a non-linear transform to the filter output, such as half-wave rectification or squaring, possibly followed by a power or logarithmic transform. The time-varying phase in each channel may also be represented.

4. How do filters affect brain data?

The answer to this question depends on the data and on the filter. In this section we review a number of archetypical "events" that might occur within a time series of brain activity, and look at how they are affected by commonly used filters.

Impulse or spike.. Brain events that are temporally localized, for example a neuronal "spike", can be modeled as one or a few impulses. It is obvious from Eq. 1 that such events must be less localized once filtered, as summarized schematically in Fig. 5. The response is *spread over time*, implying that the temporal location of the event is less well defined. It is *delayed* if the filter is causal. The delay may be avoided by choosing a zero-phase filter (green), but the response is then *acausal*. If the impulse response has multiple modes, these may appear misleadingly as multiple *spurious events*, confusing the analysis.

[Figure 5 about here]

The nature and extent of these effects depends on the filter, and can be judged by looking at its impulse response. Figure 6 shows impulse responses of a selection of commonly used filters (others were shown in Fig. 3). The left-hand plot shows the time course of the impulse response, and the right-hand plot displays the logarithm of its absolute value using a color scale, to better reveal the low-amplitude tail. The first three examples (A-C) correspond to low-pass filters with the same nominal cutoff (10

Hz). The next two (D-E) are low pass filters with nominal cutoff 20 Hz. The following three (F-H) are band-pass filters. Two of the filters are zero-phase (B and E, in green), the others are causal (blue).

The response of the first two filters is relatively short and unimodal, that of the others is more extended and includes excursions of both signs. The temporal span is greater for filters of high order (compare F and G) and for lower frequency parameters (compare C and D). Bandpass filters have relatively extended impulse responses, particularly if the band is narrow or the slopes of the transfer function steep. The oscillatory response of a filter to an impulse-like input is informally called 'ringing', and may occur in all filter types (lowpass, bandpass, highpass and so on).

Of course, real brain events differ from an infinitely narrow unipolar impulse, for example they have finite width, and the response to such events will thus differ somewhat from the ideal impulse response. As a rule of thumb, features of the impulse response that are *wider* than the event are recognizable in the response of the filter to the event. Features that are *narrower* (for example the one-sample impulse at the beginning of the impulse response of the high-pass and notch filters in Fig. 3) may appear smoothed.

[Figure 6 about here]

Step.. Certain brain events can be modeled as a step function, for example the steady-state pedestal that may follow the onset of a stimulus (Picton et al., 1978; Lammertmann and Lütkenhöner, 2001; Southwell et al., 2017). Figure 7 illustrates the various ways a step can be affected by filtering: the step may be *smoothed* and spread over time, implying that its temporal location is less well defined, and it may be *delayed* if the filter is causal. Multiple *spurious events* may appear, some of which may occur before the event if the filter is acausal.

[Figure 7 about here]

The nature of these effects depends on the filter and can be inferred from its *step response* (integral over time of the impulse response). Step responses of typical filters are shown in Fig. 8. The sharp transition within the waveform is smoothed by a lowpass filter (A-B) and delayed relative to the event if the filter is causal (A), or else it starts before the event if the filter is acausal (B). The steady-state pedestal is lost for a high-pass (C-E) or bandpass (F-H) filter. The response may include spurious excursions, some of which precede the event if the filter is acausal. The response may be markedly oscillatory (ringing) (F-H), and it may extend over a remarkably long duration if the filter has a narrow transfer function.

Of course, actual step-like brain events differ from an ideal step. As a rule of thumb, features of the step response that are *wider* than the event onset will be recognizable in the output, whereas features that are *narrower* will appear smoothed. Note that a response of opposite polarity would be triggered by the *offset* of a pedestal.

[Figure 8 about here]

Oscillatory pulse. Some activity within the brain is clearly oscillatory (Buzsáki, 2006; da Silva, 2013). A burst of oscillatory activity can be modeled as a sinusoidal pulse. As Fig. 9 shows, the time course of such a pulse is affected by filtering: it is always *smoothed* and spread over time, it may be *delayed* if the filter is causal, or else start earlier than the event if the filter is acausal. These effects are all the more pronounced as the filter has a narrow passband (as one might want to use to increase the SNR of such oscillatory activity).

For a notch filter tuned to reject the pulse frequency, ringing artifacts occur at both onset and offset. If the filter is acausal, these artifacts may both precede and follow onset and offset events. For a notch filter tuned to reject power line components (50 or 60 Hz), such effects might also be triggered by fluctuations in amplitude or phase. They might also conceivably affect the shape of a short narrow-band gamma brain response in that frequency region (Fries et al., 2008; Saleem et al., 2017).

[Figure 9 about here]

5. What can go wrong?

The use of filters raises many concerns, some serious, others merely inconvenient. It is important to understand them, and to report enough details that the reader too fully understands them. An obvious concern is *loss of useful information* suppressed together with the noise. Slightly less obvious is the *distortion* of the temporal features of the target: peaks or transitions may be smoothed, steps may turn into pulses, and *artifactual features* may appear. Most insidious, however, is the *blurring of temporal or causal relations* between features within the signal, or between the signal and external events such as stimuli. This section reviews a gallery of situations in which filtering may give rise to annoying or surprising results.

Loss of information.. This is an obvious gripe: information in frequency ranges rejected by the filter is lost. High-pass filtering may mask slow fluctuations of brain potential, whether spontaneous or stimulus-evoked (Picton et al., 1978; Lammertmann and Lütkenhöner, 2001; Vanhatalo et al., 2005; Southwell et al., 2017). Low-pass filtering may mask high-frequency activity (e.g. gamma or high-gamma bands), or useful information about the shape of certain responses (Cole and Voytek, 2017; Lozano-Soldevilla, 2018). A notch filter might interfere with narrowband gamma activity that happens to coincide with the notch frequency (Fries et al., 2008; Saleem et al., 2017). A bandpass filter may reduce the distinction between shapes of spikes emitted by different neurons and picked up by an extracellular microelectrode, degrading the quality of spike sorting.

Artifactual features. Slightly less obvious is the *distortion* of the temporal features of the target: peaks or transitions may be smoothed, steps may turn into pulses, and so on. *Artifactual features* may emerge, such as response peaks, or oscillations (”ringing”) created de novo by the filter in response to some feature of the target or noise signal. Figure 10 shows the response to a step of a high-pass filter (Butterworth order 8) of various cutoff frequencies. The response includes multiple excursions of both polarities (”positivities” and ”negativities”) that may have no obvious counterpart in the brain signal. Disturbingly, the latencies of some fall in the range of standard ERP response features (schematized as lines in Figure 10).

[Figure 10 about here]

The morphology of these artifacts depends on both the filter and the brain activity, as further illustrated in Fig. 11. An investigator, or a reader, might wrongly be tempted to assign to the multipolar deflections of the filter response a sequence of distinct physiological processes. Similar issues have been pointed out with respect to spike waveform morphology from extracellular recordings (Quian Quiroga, 2009; Molden et al., 2013).

Spurious oscillations.

[Figure 11 about here]

Oscillatory phenomena play an important role in the brain (Buzsáki, 2006; da Silva, 2013), and many response patterns are interpreted as reflecting oscillatory activity (Zoefel and VanRullen, 2017; Meyer, 2017; Singer, 2018), although in some cases this interpretation has been questioned (Yeung et al., 2004; Yuval-Greenberg et al., 2008; Jones, 2016; van Ede et al., 2018).

Non-oscillatory inputs (e.g. an impulse or step) can trigger a filter response with distinctly oscillatory features. Figure 12 shows the response of a 8-11 Hz bandpass filter (such as might be used to enhance alpha activity relative to background noise) to several inputs, including a 10 Hz sinusoidal pulse (top) and two configurations of impulses. Visually, the responses to the non-oscillatory impulse pairs are, if anything, more convincingly oscillatory than the response to the oscillatory input!

[Figure 12 about here]

Oscillations tend to occur with a frequency close to a filter cutoff, and to be more salient for filters with a high order. They can occur for any filter with a sharp cutoff in the frequency domain, and are particularly salient for band-pass filters, as high-pass and low-pass cutoffs are close and may interact. Furthermore, if the pass band is narrow, the investigator might be tempted to choose a filter with steep cutoffs, resulting in a long impulse response with prolonged ringing.

Masking or reintroduction of artifacts. Cognitive neuroscientists are alert to potential artifacts, for example muscular activity that differs between conditions due to different levels of effort. Muscular artifacts are most prominent in the gamma range (where they emerge from the 1/f background), and thus low-pass filtering is often indicated to eliminate them. Indeed, visually, there is little in the low-pass filtered signal to suggest muscle artifacts. Low-frequency correlates are nonetheless present (muscle spikes are wideband) and could potentially induce a statistically significant difference between conditions. Filtering masks this problem (if there is one).

Conditions that require different levels of effort might also differ in the number of eye-blinks that they induce. Subjects are often encouraged to blink between trials, so as avoid contaminating data within the trials. However if high-pass or band-pass filtering is applied to the data before cutting them into epochs, the filter response to the blink may extend into the epoch, again inducing a statistically significant difference between conditions. For a causal filter each epoch is contaminated by any blinks that precede it, for an acausal filter it may also be contaminated by any blinks that follow it.

Temporal blurring, delay, causality. The most subtle effect of filtering is the blurring of temporal relationships, which can interfere with the comparison between brain measurements and stimulation or behavior, or between recordings at different recording sites, or between different frequency bands. Temporal or causal relationships between events are less clear when looking at filtered data. The problem is mild if the filter impulse response is short relative to the phenomena being measured, but such is not always the case. Impulse responses of commonly-used filters may extend over hundreds of milliseconds (Fig. 6) whereas important stages of neural processing may occur over shorter time scales.

The time course of sensory processing is often inferred either from the latency of the *peak* response to stimulation, or of the point at which the response emerges from background noise. A causal filter introduces a systematic bias in the first measure (towards a longer latency), and an acausal filter a bias in the second measure (towards a shorter latency). The early part of an acausal filter response might misleadingly masquerade as an early brain response, or as the correlate of a predictive mechanism.

Similar issues arise for Temporal Response Functions (TRF) obtained by fitting stimulus and response data with a linear model (Lalor and Foxe, 2010; Ding and Simon, 2013; O'Sullivan et al., 2015; Crosse et al., 2016). TRF analysis has become popular as a tool to characterize the response to continuous stimuli such as speech or environmental sound. Features of the TRF (e.g. peaks) are sometimes interpreted as reflecting particular brain processes, and inferences are made about their anatomical localization based on their latency. If, as is common, the brain data are filtered to restrict the analysis to a frequency range where the response is expected to best follow the stimulus (e.g. 1-10 Hz), the estimated TRF will approximate the *real TRF filtered with the impulse response of the filter*. To illustrate this point, a simulated "stimulus" consisting of Gaussian white noise was processed with a simulated "TRF" consisting of a half-sinusoidal pulse of duration 50 ms (Fig. 13 black line) to obtain simulated brain data. These data were then filtered with a bandpass filter and the TRF estimated using the mTRF toolbox (Crosse et al., 2016). Figure 13 (blue line) shows the TRF estimate. The green line is the estimate when the same filter was applied in both directions using `filtfilt`. In both cases, the shape of the estimated TRF differs from that of the real TRF. The potential effect of filtering on TRFs is rarely discussed, and filters used to preprocess the data prior to TRF analysis are often not fully described.

[Figure 13 about here]

Time-frequency analysis. Time-Frequency (TF) analysis is usually seen as a data analysis rather than filtering tool. Nonetheless, filters are involved "under the hood" and TF representations are vulnerable to similar problems as noted for filters.

Time-frequency (TF) representations (e.g. spectrograms) are obtained by applying short-term spectral analysis to the data with a short analysis window that slides in time. At each time point the analysis yields a spectrum and these spectra are concatenated to form the two-dimensional TF representation. Each pixel in the 2-D representation is indexed by time (abscissa) and analysis frequency (ordinate). The representation usually displays some transform of amplitude or power (Fig. 14 B, C), but it is also possible to plot phase (Fig. 14 D, E). Importantly, the computation of a TF representation can be equivalently formulated in terms of filtering, using one filter (or two related filters) for each frequency. Thus, everything we said about filters holds for TF representations as well.

In a standard "short-term Fourier transform" (STFT) spectrogram, the size of the analysis window is the same for all frequencies (Fig. 14 B). In contrast, in a wavelet spectrogram this parameter varies with frequency, for example such that each analysis window spans the same number of cycles (Fig. 14 C), i.e. it is longer at low frequencies and shorter at high frequencies.

The value of the TF representation at the analysis time point reflects all signal values within the analysis window. Conversely each signal value impacts TF values over a range of analysis time points. The overall alignment between data values and TF values depends on the convention chosen to assign a time index to the analysis value. TF samples can be aligned with the *end* of the analysis window, corresponding to a causal analysis, or more commonly with the *center* of the analysis window, corresponding to an acausal analysis. TF features are thus either delayed relative to events within the data (causal analysis) or else they partly reflect future events (acausal analysis).

Figures 14 (B, C) show TF magnitude representations in response to a pulse-shaped input signal. The temporally-localized event at $t = 0$ affects the spectrogram over a range of time points spanning the event (for example ± 0.25 s in Fig. 14 B). Equivalently, the value of the spectrogram at $t = 0$ can "see" all signal values within a range of time points spanning that instant.

[Figure 14 about here]

Figures 14 (D, E) show TF phase representations in response to the same pulse-shaped input signal. The color of each pixel represents the phase estimate (calculated over the analysis window) for that time and frequency channel, in response to the pulse at $t = 0$. Phase is defined only for non-zero magnitude, i.e. only when the pulse falls within the analysis window. The event at $t = 0$ affects the phase estimate of analyses made over a range of time points spanning the event. Equivalently, the phase estimate obtained at $t = 0$ is affected by all events within that range, some of which occur *later* than the analysis point. This blurred, non-causal relation between data and TF analysis can lead to misleading conclusions.

As an example of such a misleading conclusion, suppose that we wish to establish whether the phase of brain oscillations preceding a stimulus predicts the brain or behavioral response to that stimulus. TF analysis seems to be the right tool for that purpose. Indeed, using it we observe that phase within some frequency band (e.g. alpha) measured just before the trial is systematically biased towards a particular value on successful trials. From this we conclude that oscillatory phase preceding stimulation determines the response. Unfortunately, that conclusion is not warranted if the analysis window overlaps the stimulus-evoked sensory or behavioral motor response. The interesting conclusion (response dependency on prior phase) can only be made if that more trivial possibility is ruled out, for example by using causal TF analysis (Zoefel and Heil, 2012). Similar issues may arise in analyses of cross-spectral coupling. These issues may be harder to spot if wavelet analysis is involved, because the span of the analysis window varies with the frequency channel.

Notch filter artifacts. A narrow notch filter works well to remove narrowband interference that is stationary (e.g. 50 or 60 Hz line power). However, notch filtering may be less effective if the interference is not stationary. Amplitude fluctuations may occur if the subject moves, and for MEG the phase may fluctuate with changes in load in the tri-phase power network from which originates the interference. Artifacts can also be triggered by large-amplitude glitches (Kıraç et al., 2015). Notch filtering is ineffective in removing interference close to the ends of the data (Fig. 9, bottom), and thus should not be applied to epoched data.

Inadequate antialiasing.. Effects of the antialiasing filter are rarely noticed or objectionable. More serious may be a *lack* of sufficient antialiasing. Figure 15 (left) shows the power spectrum of a sample of data from a MEG system with shallow (or missing) antialiasing. As common in MEG, there are salient power line components at 60 Hz and harmonics (black arrows), but also many additional narrowband components that likely reflect aliasing of sources with frequencies beyond the Nyquist frequency (250 Hz). Possible sources include higher harmonics of 60 Hz, or high-frequency interference from computer screens, switching power supplies, etc. The frequency of the artifact cannot be known for sure. For example, the spectral line at 200 Hz (red arrow) could be the aliased 300 Hz harmonic of the power line interference, or it could have some other origin. This example underlines the importance of an adequate antialiasing filter.

In contrast, Fig. 15 (left) shows the response to a sharp change in sensor state of one channel of a different MEG system with a particularly steep antialiasing filter (8th order elliptic filter with 120 dB rejection and 0.1 ripple in pass-band) (Oswal et al., 2016). The data show a prominent oscillatory pattern that is likely not present in the magnetic field measured by the device.

[Figure 15 about here]

6. How to fix it ?

Report full filter specs. This should go without saying. In each case, the problem is compounded if the reader can't form an opinion about possible effects of filtering. Filter type, order, frequency parameters, and whether it was applied in one direction or both (`filtfilt`) should be reported. Include a plot of the impulse response (and/or step response) as an inset to one of the plots. When reporting the *order* of a filter, be aware that for FIR filters this refers to the duration of the impulse response in samples (minus one), whereas for IIR filters implemented recursively it refers to the largest delay in the difference equation that produces each new output sample as a function of past input and output samples. The order of an IIR filter is usually small (e.g. 2-16), whereas that of an FIR is often large (e.g. 100-1000). The plot thickens when an IIR filter is *approximated* by an FIR filter. In that case, both numbers should be reported. An "order-512 Butterworth filter" is an unusual beast.

Antialiasing artifacts.. Antialiasing artifacts are rarely an issue. In the event that they are, consider first whether antialiasing is needed. If you are certain that the original data contain no power beyond the Nyquist frequency, omit the filter and live dangerously. If instead there are high-frequency sources of large amplitude, you might want to verify that the antialiasing filter attenuates them sufficiently before sampling. Note that, because of the aliasing, the frequency of those sources cannot be inferred with confidence from the sampled data. A wide-band oscilloscope or spectrum analyzer might be of use applied to the data before sampling. To reduce temporal smear and ringing, consider using an antialiasing filter with a lower cutoff and shallower slope. In the case of downsampling or resampling of a signal that is already sampled, consider alternatives such as interpolation (e.g. linear, cubic or spline). The artifact of Fig. 15 (right) can be removed as described by de Cheveigné and Arzounian (2018).

Low-pass. First, ask whether the aim is to *smooth* the temporal waveform, for example to enhance the clarity of a plot, or whether it is to ensure attenuation of high-frequency power (for example preceding downsampling). If the former, consider using a simple smoothing kernel, for example square, triangular, or Gaussian (Fig. 6 a, b). Such kernels have a limited and well-defined temporal extent, and no negative portions so they do not produce ringing. They tend however to have poor spectral properties. Conversely, if temporal distortion is of no importance, the filter can be optimized based only on its frequency response properties (Widmann et al., 2015).

If data are recorded on multiple channels (e.g. local field potentials, EEG, or MEG), spatial filters may be applied to remove noise sources with a spatial signature different from the target sources. The appropriate filters can be found based on prior knowledge or using data-driven algorithms (e.g. Parra et al., 2005; de Cheveigné and Parra, 2014).

High-pass. If the high-pass filter is required merely to remove a constant DC offset, consider subtracting the overall mean instead. If there is also a slow trend, consider *detrending* rather than high-pass filtering. Detrending involves fitting a function (slowly-varying so as to fit the trend but not faster patterns) to the data and then subtracting the fit. Suitable functions include low-order polynomials. Like filtering, detrending is sensitive to temporally localized events such as glitches, however these can be addressed by *robust detrending* (de Cheveigné and Arzounian, 2018).

If the slow trend signal can be estimated independently from the measurement that it contaminates, consider using regression techniques to factor it out (Vrba and Robinson, 2001; de Cheveigné and Simon, 2007). Even when this is impossible, if the data are multichannel, consider using a component-analysis technique to factor it out, as has also been suggested to obtain distortion-free extracellular spike waveforms (Molden et al., 2013).

If all else fails, and high-pass filtering must be used, pay particular attention to its possible effects on the morphology of responses. If the initial portion of the data (duration on the order of $1/f_c$ where f_c is the cutoff frequency) is on average far from zero, it may be useful to subtract the average over that portion, so as to minimize the filter response to the implicit initial step (the filter treats the input data as being preceded by zeros). If the data are to be cut into epochs (e.g. to excise responses to repeated stimuli), it is usually best to filter the continuous data first. Be aware that artifacts from out-of-epoch events (e.g. eye blinks) may extend to within the epoch.

Band-pass. Consider whether a band-pass filter is really needed, as the potential for artifactual patterns is great. If band-pass filtering must be applied (for example to improve signal-to-noise ratio to assist a component-analysis technique), consider filters with relatively shallow slopes, and cutoff frequencies distant from the activity of interest. Be on the lookout for artifactual results due to the filtering.

Notch. Notch filtering is usually motivated by the desire to suppress line noise (50 Hz or 60 Hz and harmonics). Of course, the best approach is to eliminate that noise at the source by careful design of the setup, but this is not always feasible. As an alternative to filtering, it may be possible to measure the line noise on one or more reference channels and regress them out of the data (Vrba and Robinson, 2001). If the data are multichannel, consider using component analysis to isolate the line noise components and regress them out (Delorme et al., 2012; de Cheveigné and Parra, 2014; de Cheveigné and Arzounian, 2015).

If the high-frequency region is not of interest, a simple expedient is to apply a boxcar smoothing kernel of size 1/50 Hz (or 1/60 Hz as appropriate). This simple low-pass filter has zeros at the line frequency and all its harmonics, and thus perfectly cancels line noise. The mild loss of temporal resolution (on the order of 20 ms) might be deemed acceptable. If the sampling rate differs from a multiple of the line frequency, the appropriate kernel can be implemented using interpolation (see de Cheveigné and Arzounian, 2018, for details).

Time-frequency analysis. If the patterns of interest can be interpreted in the time domain, eschew TF analysis. If the data are multichannel, and the aim is to increase the signal-to-noise ratio of narrow-band or stimulus-induced activity, consider component analysis techniques that can boost SNR of narrow-band signals (Nikulin et al., 2011; de Cheveigné and Arzounian, 2015).

If TF analysis must be applied, consider using fixed kernel-size analysis (e.g. DFT) rather than, or in addition to, wavelet analysis, so that temporal bias and smearing are uniform across the frequency axis. Consider using relatively short analysis windows to reduce temporal bias and/or smearing. Weigh carefully the choice between causal analysis (temporal bias but no causality issues) and acausal analysis (no temporal bias but risk of misleading causal relations). In every case, be alert for potential artifacts. One should be particularly concerned if an interesting effect only emerges with a particular analysis method.

7. Horror scenarios

This section imagines scenarios in which filtering effects might affect the science. Some are mildly embarrassing, others might keep a scientist awake at night.

Missed observation.. Researcher **A** applies a high-pass filter to data recorded over a long period and fails to notice the existence of infra-slow brain activity (as reported by Vanhatalo et al., 2005). Researcher **B** applies a low-pass filter and fails to notice that a certain oscillatory activity is not sinusoidal (as reported by Cole and Voytek, 2017). It is frustrating to miss part of the phenomena one set out to study.

Bias from eye movements.. Following a scenario hinted at in Sect. 5, researcher **C** runs a study in which some conditions are more demanding than others. Subjects are instructed to blink only between trials, but because acausal high-pass (or band-pass) filtering is applied to the data, each blink triggers a filter response that extends into the trial, resulting in a significant difference between conditions. Researcher **D** runs studies that create miniature eye movements (microsaccades) that differ between conditions. Microsaccades introduce so-called spike potentials, transients with a time course of a few tens of milliseconds, which after TF analysis boost energy in the gamma band selectively in some conditions rather than others (Yuval-Greenberg et al., 2008). In both cases ocular activity masquerades as brain activity.

Distorted observation.. Researcher **E** records brain responses to stimulation, applies a high-pass filter to attenuate a pesky slow drift, and fails to notice that the brain response actually consisted of a sustained pedestal. Instead, a series of positive and negative peaks is observed and interpreted as reflecting a succession of processing stages in the brain. In a milder version of this scenario, the brain response does include such peaks, but the filter affects their position, leading to incorrect inferences concerning brain processing latencies.

Flawed replication.. Researcher **F** replicates Researcher **E**'s experiments, using the same filters and generating the same artifacts. Results are consistent, giving weight to the conclusion that they are real.

Faulty communication.. Researcher **G**, who is filter-savvy, reads his/her colleague's papers and suspects something is amiss, but cannot draw firm conclusions because methods were not described in full. He/she re-runs the experiments with careful methods, and finds results that invalidate the previous studies. The paper is not published because the study does not offer new results.

Proliferation of "new" results.. Other researchers run further studies using analogous stimuli, but using different analysis parameters. New patterns of results are found that are interpreted as new discoveries, whereas the actual brain response (in this hypothetical scenario) is the same.

Oscillations?. Researcher **H** knows that with the right kind of preprocessing, multiple layers of oscillatory activity can be found hidden within brain signals, and is confident that the analysis is revealing them. Researcher **I** suspects that these oscillations reflect filter ringing, but finds it hard to counter **H**'s arguments (Fourier's theorem says that any signal is indeed a compound of oscillations). **I** remains worried because the observed oscillations depend on the choice of filter, but **H** is not: different filters extract different parts of the data, each with its own oscillatory nature. The debate mobilizes a good proportion of their energy.

Biased time-frequency analysis.. Researcher **K** uses time-frequency analysis to test the hypothesis that the phase of ongoing brain oscillations modulates perceptual sensitivity. To avoid contamination by the sensory or behavioral response, the analysis is carefully restricted to the data preceding stimulation. However the analysis window, centered on the analysis point, extends far enough to include the sensory or behavioral response, biasing the distribution of measured phase. **K** concludes (incorrectly in this hypothetical scenario) that the hypothesis is correct. In a variant of this scenario, **L** uses time frequency analysis to test the hypothesis that brain activity is durably entrained by a rhythmical stimulus. The analysis is applied to the data beyond the stimulus offset, but the analysis window overlaps with the stimulus-evoked response, again biasing the phase distrtibution. **L** concludes (again incorrectly in this scenario) that the hypothesis was correct.

8. Discussion

A filter has one purpose, improve SNR, and two effects: improve SNR and distort the signal. Many investigators consider only the first and neglect the second. The filtered data are the sum of the filtered target signal and the filtered noise, and thus one can focus separately on these two effects (Fig. 16). Here we focused on target distortion.

[Figure 16 about here]

Issues related to distortion have been raised before, in particular distortion due to low-pass filtering (VanRullen, 2011; Rousselet, 2012; Widmann and Schröger, 2012), high-pass filtering (Kappenman and Luck, 2010; Acunzo et al., 2012; Tanner et al., 2015, 2016; Widmann et al., 2015; Lopez-Calderon and Luck, 2014) and band-pass filtering (Yeung et al., 2004), in the context of EEG and MEG and also extracellular recordings (Quian Quiroga, 2009; Molden et al., 2013; Yael and Bar-Gad, 2017). They are also discussed in textbooks and guidelines (Picton et al., 2000; Nunes and Srinivasan, 2006; Gross et al., 2013; Keil et al., 2014; Luck, 2014; Puce and Hämäläinen, 2017; Cohen, 2014, 2017).

How serious are these issues?. They can be recapitulated as follows. First, the *loss of information* in spectral regions suppressed by the filter. This problem is straightforward and does need elaboration. Second, the *distortion of response waveforms* and the *emergence of spurious features*. This is certainly a concern if spurious features (e.g. delayed excursions, or ringing) misleadingly suggest brain activity that is not there. Third, the *blurring of temporal relations*, in particular violation of *causality*. This too is a concern given the importance of response latency in inferring the sequence of neural events, or the anatomical stage at which they occur. Fourth, *the non-uniqueness of phenomenological descriptions*: the same event can take very diverse shapes depending on the analysis. This can interfere with comparisons between studies, and can lead to redundant reports of the same phenomenon under different guise. Fifth, the *lack of details* required by a knowledgeable reader to infer the processing involved. Rather than an issue with filtering per se, the issue is with sloppy practice in reporting methodological details of filtering and TF analysis.

Cutoff frequencies may be reported, but not the type of filter, its order, or whether it was applied in a single pass or both ways. As illustrated in Figs. 6 and 8, the cutoff frequency of a filter is not sufficient to characterize its impulse or step response, information that is needed to guess how it might have impacted a reported response. Failure to report details can be due to space limits (sometimes misguidedly imposed by journals), incomplete knowledge (e.g. proprietary or poorly documented software), reluctance to appear pedantic by reporting mundane trivia, or lack of understanding that this information is important.

The issue of non-uniqueness is not often raised. Non-uniqueness refers to the fact that analysis of the same phenomenon can give rise to different descriptions depending on the analysis parameters, making it hard to compare across studies. It is sometimes recommended that parameters should be adjusted to the task at hand, rather than use default values proposed by the software (Widmann and Schröger, 2012). Optimizing data analysis is laudable, but it carries the risks of "cherry-picking" or "double-dipping" (Kriegeskorte et al., 2009).

Quid frequency and phase?. Filter design has developed sophisticated methods to optimize the frequency response to maximize rejection, minimize ripple, and/or obtain the steepest possible transition between pass and stop bands. Engineers and scientists trained in those methods tend to choose a filter based on these properties, with less attention to their time-domain counterpart. It is not always clear that this emphasis is justified. For example, a band-pass filter with steep slopes might be motivated by the desire to "keep the delta band distinct from the theta band", but given that there is little theoretical or phenomenological evidence for a clear boundary between bands, this is should perhaps not be a primary goal.

For any given magnitude response, there are multiple filters with different phase responses. Of particular interest are *zero phase* filters with minimal waveform distortion and no delay (but that are unfortunately acausal), and *minimum phase* filters with greater waveform distortion but that are causal. The choice between these phase characteristics (or others) depends on whether one wishes to favour causality, overall delay, or waveform distortion, knowing that it is impossible to favour all. Some authors recommend causal filters (Rousselet, 2012), others linear phase or acausal (Widmann and Schröger, 2012). Some studies report using simple filters (e.g. low order Butterworth), others sophisticated designs (e.g. Chebyshev or elliptic) or even "brickwall" filters implemented in the Fourier domain.

A crucial point that we strive to make in this paper is that *no* choice of filter can avoid temporal distortion, as *any* filter entails scrambling of the temporal axis (Fig. 2). Given that a filter with steep slopes in the frequency domain entails a long impulse response (a problem of particular importance when using brickwall filters in the Fourier domain), it may be worth relaxing spectral criteria so as to optimize temporal properties.

Causality, again.. As mentioned earlier, for an acausal filter the output depends on input values that occur later in the future. No physical system can have this behaviour. Offline analysis allows us greater flexibility to align the analysis arbitrarily with respect to the data, but we must be clear about what this implies. If we wish to relate the "brain response" to other events within the brain or the world (e.g. stimuli or behaviour), acausal filtering implies that that this response might depend on signal samples that occur *after* those events, indeed, a violation of causality.

9. Recommendations

Document.. This should go without saying, but many (most?) papers provide incomplete information about the filters employed, a situation exacerbated by the insistence of some journals on limiting the space devoted to methods. Data analysis decisions, however suboptimal, can be justified, incomplete reporting cannot. The reader needs this information to infer the brain signal from the patterns reported.

To authors: provide full specifications of the filters applied to the data. A simple plot of the impulse response (or step response) as an insert can be very helpful. To editors and reviewers: demand this information. To journals: avoid requirements that discourage proper documentation. To equipment manufacturers: provide full specifications of any hardware filters.

Know your filters.. Make sure that you know the exact filters that are involved in your data recording and analysis. This may require delving into the documentation (or even the code) of your analysis software (e.g. EEGlab, FieldTrip, SPM. etc.). Plot the impulse response and/or the step response and paste it on the wall in front of your desk. If several filters are cascaded, plot the response of their cascade. If specs are lacking, figure out how to deliver a pulse (and/or step) to the recording device and plot the resulting response. If you are using TF analysis, do you know exactly what kernels were employed? Are they causal and thus likely to introduce latency? Are they instead acausal (e.g. zero-phase) and thus likely to confuse causal relations? Are they wavelets, in which case temporal spread and latency might differ across frequency bands? All this should be known.

Know your noise.. The main purpose of a filter is to attenuate noise. What is that noise, where does it come from? Might it be possible to mitigate it at the source? Some experimenters speak of their rig as if it were inhabited by gremlins. This deserves little patience: how can one understand the brain if we can't find the source of line noise in the rig? It may not be possible to suppress the noise (e.g. turn off myogenic, cardiac, ocular or alpha activity, tramways in the street, etc.) but at least the source should be understood. Given that signal and noise both impact the results, understanding a noise process merits as much effort as understanding a brain process.

Eliminate noise at the source.. No need for a filter if there is nothing to attenuate. To get rid of line noise: banish power cables from the vicinity of the setup, use lights fed with filtered DC, apply proper shielding (electrostatic coupling), avoid loops (magnetic coupling), avoid ground loops (ensure that ground cables and shields have a star topology with no loops), etc. To eliminate high-frequency noise: banish computer screens, fluorescent lights, equipment with switching power supplies, cell phones, etc. If need, apply Faraday shielding. To minimize slow drifts in EEG: follow appropriate procedures when applying the electrodes, and keep the subjects cool. To minimize alpha components: ensure that subjects keep their eyes open, give them a task to keep them alert, and so on. Textbooks (e.g. Luck, 2014; Cohen, 2014) and guidelines can offer many such suggestions.

Ensure that you have adequate antialiasing.. Antialiasing filters in recording equipment are not always well documented. In some situations they might prove insufficient if there is high amplitude noise with a frequency beyond the Nyquist rate (for example from a computer screen, fluorescent light, or cell phone). A similar issue may arise when downsampling digital data: does the low-pass filter suffice to ensure that aliased components are negligible? This may require checking the data and/or software at hand (at the time of writing, Matlab's `resample` sets the low-pass cutoff *at* Nyquist rather than below, which is inadequate).

Consider alternatives to filtering.. Consider *detrending* (in particular robust detrending) as an alternative to high-pass filtering (Bigdely-Shamlo et al., 2015; de Cheveigné and Arzounian, 2018). Consider using an independent reference signal measurement that picks up only noise, and use *regression* techniques to factor out the noise (Vrba and Robinson, 2001; de Cheveigné and Simon, 2007; Molden et al., 2013). Consider component analysis techniques to design a *spatial filter* that factors out the noise (Parra et al., 2005; Delorme et al., 2012; de Cheveigné and Parra, 2014).

Choose the right filter.. If filter we must, a prime consideration is whether to optimize the time domain (minimal distortion of the waveform) or the frequency domain (optimal frequency response), the two being at loggerheads. Taking the example of a low-pass filter, if our goal is to smooth the waveform to enhance the visual clarity of a plot, or locate a peak with less jitter, then a simple box-car smoothing kernel (rectangular impulse response) may be sufficient, with minimal temporal blurring. The poor frequency response of such a low-pass filter is of little import. If instead the focus is on spectral features (e.g. frequency-following response, or narrowband oscillations), we may wish to optimize the spectral properties of the filter at the expense of greater temporal smearing. If the focus is on spectrotemporal features, then the choice of filter(s) necessarily involves a tradeoff between the two (Cohen, 2014).

Simulate.. It is hard to fully predict the impact of filtering, particularly if multiple stages are cascaded. A simple expedient is to simulate the situation using a known target signal (e.g. an idealized evoked response) and known noise (e.g. EEG data from an unrelated recording). The effect of filtering can then be evaluated separately on each, given that the filtered sum is the sum of the filtered parts (Fig. 16).

The synthetic target signal could be an impulse or step (to visualize canonical response properties), or a signal similar to a typically-observed response (to see how processing might affect it), or a signal constructed to mimic the observed response after filtering (to help infer true patterns from observations). Observing the response to the target tells us how it is distorted, observing the response to the noise tells us how well it is attenuated and what artifactual patterns to expect. Comparing the two tells us whether our observation is helped (or hindered) by filtering.

Be paranoid.. Is the effect of interest only visible for a particular type of filter, or a particular variety of TF analysis? Consider whether it might depend on an artifact of that filter or analysis. Do your conclusions involve temporal or causal dependencies between events in the EEG and events in the world? Make sure that you fully understand how they might be affected by filtering or the TF analysis.

Go with the zeitgeist.. This is in counterpoint to the previous recommendations. One cannot ignore that many studies, past and present, employ filters in ways that we describe as problematic. Those results cannot be discarded, and one may need to use similar methods oneself to allow comparisons, and place new results within the context of prior knowledge. Many researchers and laboratories have well established methodologies that may need to be adhered to for consistency. If such is the case, go for it, but don't forget to fully document, and do call the reader's attention to potential issues.

Conclusion

Filters are ubiquitous in electrophysiology and neuroscience and are an important part of the methodology of any study. Their role is to suppress noise and enhance target activity, but they may have deleterious effects that the investigator should be aware of. When reporting results, it is important to provide enough details so that the reader too can be aware of these potential effects. In some cases there exist alternatives to filtering that are worth considering, in others a filter cannot be avoided. In every case, care must taken to fully understand and report the potential effects of filtering on the patterns reported.

Acknowledgements

The first author was supported by the EU H2020-ICT grant 644732 (COCOHA), and grants ANR-10-LABX-0087 IEC and ANR-10-IDEX-0001-02 PSL*. The second author was supported by grant 390/12 from the Israel Science Foundation (ISF) and by ERC advanced grant GA-340063 (project RATLAND). Giovanni Di Liberto offered useful comments on previous manuscripts.

Author Contributions

Conceptualization and writing: AdC and IN.

Declaration of Interests

None.

References

Acunzo DJ, MacKenzie G, van Rossum MCW (2012) Systematic biases in early ERP and ERF components as a result of high-pass filtering. *Journal of Neuroscience Methods* 209:212–218.

Bigdely-Shamlo N, Mullen T, Kothe C, Su KM, Robbins KA (2015) The PREP pipeline: standardized preprocessing for large-scale EEG analysis. *Frontiers in Neuroinformatics* 9:16.

Buzsáki G (2006) *Rhythms Of The Brain* Oxford Univ. Press.

Cohen MX (2017) Rigor and replication in time-frequency analyses of cognitive electrophysiology data. *International Journal of Psychophysiology* 111:80–87.

Cohen MX (2014) *Analyzing Neural Time Series Data - Theory and practice* MIT Press.

Cole SR, Voytek B (2017) Brain Oscillations and the Importance of Waveform Shape. *Trends in Cognitive Sciences* 21:137–149.

Crosse MJ, Di Liberto GM, Bednar A, Lalor EC (2016) The Multivariate Temporal Response Function (mTRF) Toolbox: A MATLAB Toolbox for Relating Neural Signals to Continuous Stimuli. *Frontiers in Human Neuroscience* 10:1–14.

da Silva FL (2013) EEG and MEG: Relevance to Neuroscience. *Neuron* 80:1112–1128.

de Cheveigné A, Arzounian D (2015) Scanning for oscillations. *Journal of Neural Engineering* 12:66020.

de Cheveigné A, Arzounian D (2018) Robust detrending, rereferencing, outilier detection, and inpainting of multichannel data. *NeuroImage* 172:903–912.

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

de Cheveigné A, Simon JZ (2007) Denoising based on time-shift PCA. *Journal of Neuroscience Methods* 165:297–305.

Delorme A, Palmer J, Onton J, Oostenveld R, Makeig S (2012) Independent EEG sources are dipolar. *PLoS ONE* 7:e30135.

Ding N, Simon JZ (2013) Adaptive Temporal Encoding Leads to a Background-Insensitive Cortical Representation of Speech. *Journal of Neuroscience* 33:5728–5735.

Fries P, Scheeringa R, Oostenveld R (2008) Finding Gamma. *Neuron* 58:303–305.

Gross J, Baillet S, Barnes GR, Henson RN, Hillebrand A, Jensen O, Jerbi K, Litvak V, Maess B, Oostenveld R, Parkkonen L, Taylor JR, van Wassenhove V, Wibral M, Schoffelen JM (2013) Good practice for conducting and reporting MEG research. *NeuroImage* 65:349–363.

Huigen E, Peper a, Grimbergen Ca (2002) Investigation into the origin of the noise of surface electrodes. *Medical & biological engineering & computing* 40:332–338.

Jones SR (2016) When brain rhythms aren't rhythmic': implication for their mechanisms and meaning. *Current Opinion in Neurobiology* 40:72–80.

Kappenman ES, Luck SJ (2010) The effects of electrode impedance on data quality and statistical significance in ERP recordings. *Psychophysiology* 47:888–904.

Keil A, Debener S, Gratton G, Junghöfer M, Kappenman ES, Luck SJ, Luu P, Miller GA, Yee CM (2014) Committee report: Publication guidelines and recommendations for studies using electroencephalography and magnetoencephalography. *Psychophysiology* 51:1–21.

Kıraç L, Vollmar C, Rémi J, Loesch A, Noachtar S (2015) Notch filter artefact mimicking high frequency oscillation in epilepsy. *Clinical Neurophysiology* 127.

Kriegeskorte N, Simmons WK, Bellgowan PSF, Baker CI (2009) Circular analysis in systems neuroscience: the dangers of double dipping. *Nature Neuroscience* 12:535–540.

Lalor EC, Foxe JJ (2010) Neural responses to uninterrupted natural speech can be extracted with precise temporal resolution. *European Journal of Neuroscience* 31:189–193.

Lammertmann C, Lütkenhöner B (2001) Near-DC magnetic fields following a periodic presentation of long-duration tonebursts. *Clinical Neurophysiology* 112:499–513.

Lopez-Calderon J, Luck SJ (2014) ERPLAB: an open-source toolbox for the analysis of event-related potentials. *Frontiers in Human Neuroscience* 8:1–14.

Lozano-Soldevilla D (2018) Nonsinusoidal neuronal oscillations: bug or feature? *Journal of Neurophysiology* 119:1595–1598.

Luck S (2014) *An introduction to the event-related potential technique* MIT press.

Meyer L (2017) The neural oscillations of speech processing and language comprehension: State of the art and emerging mechanisms. *European Journal of Neuroscience* 48:2609–2621.

Molden S, Moldestad O, Storm JF (2013) Estimating extracellular spike waveforms from CA1 pyramidal cells with multichannel electrodes. *PLoS ONE* 8.

Nikulin VV, Nolte G, Curio G (2011) A novel method for reliable and fast extraction of neuronal EEG/MEG oscillations on the basis of spatio-spectral decomposition. *NeuroImage* 55:1528–1535.

Nunes P, Srinivasan R (2006) *Electric Fields of the Brain: The Neurophysics of EEG* Oxford University Press.

O'Sullivan JA, Shamma SA, Lalor EC (2015) Evidence for Neural Computations of Temporal Coherence in an Auditory Scene and Their Enhancement during Active Listening. *Journal of Neuroscience* 35:7256–7263.

Oswal A, Jha A, Neal S, Reid A, Bradbury D, Aston P, Limousin P, Foltynie T, Zrinzo L, Brown P, Litvak V (2016) Analysis of simultaneous MEG and intracranial LFP recordings during Deep Brain Stimulation: A protocol and experimental validation. *Journal of Neuroscience Methods* 261:29–46.

Parra LC, Spence CD, Gerson AD, Sajda P (2005) Recipes for the linear analysis of EEG. *Neuroimage* 28:326–341.

Picton TW, Bentin S, Berg P, Donchin E, Hillyard SA, Johnson R, Miller GA, Ritter W, Ruchkin DS, Rugg MD, Taylor MJ (2000) Guidelines for using human event-related potentials to study cognition: recording standards and publication criteria. *Psychophysiology* 37:127–152.

Picton TW, Woods DL, Proulx GB (1978) Human auditory sustained potentials. I. The nature of the response. *Electroencephalography and Clinical Neurophysiology* 45:186–197.

Puce A, Hämäläinen MS (2017) A review of issues related to data acquisition and analysis in EEG/MEG studies. *Brain Sciences* 7.

Quian Quiroga R (2009) What is the real shape of extracellular spikes? *Journal of Neuroscience Methods* 177:194–198.

Rousselet GA (2012) Does filtering preclude us from studying ERP time-courses? *Frontiers in Psychology* 3:1–9.

Saleem AB, Lien AD, Krumin M, Haider B, Rosón MR, Ayaz A, Reinhold K, Busse L, Carandini M, Harris KD, Carandini M (2017) Subcortical Source and Modulation of the Narrowband Gamma Oscillation in Mouse Visual Cortex. *Neuron* 93:315–322.

Singer W (2018) Neuronal oscillations: Unavoidable and useful? *European Journal of Neuroscience* 48:2389–2398.

Southwell R, Baumann A, Gal C, Barascud N, Friston K, Chait M (2017) Is predictability salient? a study of attentional capture by auditory patterns. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 372:doi:10.1098/rstb.2016.0105.

Tanner D, Morgan-Short K, Luck SJ (2015) How inappropriate high-pass filters can produce artifactual effects and incorrect conclusions in ERP studies of language and cognition. *Psychophysiology* 52:997–1009.

Tanner D, Norton JJ, Morgan-Short K, Luck SJ (2016) On high-pass filter artifacts (they're real) and baseline correction (it's a good idea) in ERP/ERMF analysis. *Journal of Neuroscience Methods* 266:166–170.

van Ede F, Quinn AJ, Woolrich MW, Nobre AC (2018) Neural Oscillations: Sustained Rhythms or Transient Burst-Events? *Trends in Neurosciences* 41:415–417.

Vanhatalo S, Voipio J, Kaila K (2005) Full-band EEG (FbEEG): An emerging standard in electroencephalography. *Clinical Neurophysiology* 116:1–8.

VanRullen R (2011) Four common conceptual fallacies in mapping the time course of recognition. *Frontiers in Psychology* 2:1–6.

Vrba J, Robinson SE (2001) Signal Processing in Magnetoencephalography. *Methods* 25:249–271.

Widmann A, Schröger E (2012) Filter effects and filter artifacts in the analysis of electrophysiological data. *Frontiers in Psychology* 3:1–5.

Widmann A, Schröger E, Maess B (2015) Digital filter design for electrophysiological data - a practical approach. *Journal of Neuroscience Methods* 250:34–46.

Yael D, Bar-Gad I (2017) Filter based phase distortions in extracellular spikes. *PLoS ONE* 12:1–13.

Yeung N, Bogacz R, Holroyd CB, Cohen JD (2004) Detection of synchronized oscillations in the electroencephalogram: An evaluation of methods. *Psychophysiology* 41:822–832.

Yuval-Greenberg S, Tomer O, Keren AS, Nelken I, Deouell LY (2008) Transient Induced Gamma-Band Response in EEG as a Manifestation of Miniature Saccades. *Neuron* 58:429–441.

Zoefel B, Heil P (2012) Detection of Near-Threshold Sounds is Independent of EEG Phase in Common Frequency Bands. *Frontiers in Psychology* 4:262.

Zoefel B, VanRullen R (2017) Oscillatory mechanisms of stimulus processing and selection in the visual and auditory systems: State-of-the-art, speculations and suggestions. *Frontiers in Neuroscience* 11:1–13.

Captions

Figure 1: A typical recording-to-publication data pipeline, showing where filters are applied. Filters are analog in the first stages (green) and digitial in subsequent stages (red). The recording rig *might include a high-pass filter (implicit in the case of AC coupling), and perhaps also a notch filter to attenuate line frequency power. The* analog-to-digital converter *is preceded by a low-pass antialiasing filter. In data* preprocessing *it is common to apply a high-pass filter to remove slow drift components, and a low-pass filter to attenuate noise (often spread over the entire spectrum), or to avoid antialiasing when the data are down-sampled.* Data analysis *might involve bandpass filtering (for example to isolate a standard frequency band such as "alpha" or "gamma") or time-frequency analysis.* Data display *or plotting might call for additional smoothing (low-pass filtering).*

Figure 2: Filtering. Top: each sample of the output y is the sum of samples of the input x weighted by the impulse response h. For a causal filter, only past or present samples of the input make a contribution (black). For an acausal filter, future samples too can contribute (gray). Bottom: another way of describing this process is that each sample of the input x affects multiple samples of the output y, with a weight determined by the impulse response h.

Figure 3: Typical magnitude transfer function shapes (top) and the associated impulse responses (bottom). The low-pass filter attenuates high frequencies, the high-pass attenuates low frequencies, the band-pass attenuates out-of band frequencies, the notch attenuates a narrow band of freqencies. The steeper the transition in the frequency domain, the more extended the impulse response (red). The steepness of the transitions depends on the type and order of the filter. Low-pass, high-pass and band-pass are Butterworth filters of order 4 and 16, notch filters are second-order filters with Q factors (ratio of bandwidth to centre frequency) of 1 and 10. Impulse responses for high-pass and notch include a high amplitude impulse, plotted here with a break.

Figure 4: Left: impulse responses that all yield the same magnitude transfer function (low-pass, similar to that shown in Fig. 3 top left). The filters in blue are causal, those in green acausal. The filter in thick green is zero-phase. All examples are implemented as a cascade of two Butterworth lowpass filters of order 8 and cutoff 10 Hz. Thick blue: both impulse responses are convolved. Blue dashed: same but the result is delayed. Green dashed: same as thick blue but time-reversed. Thick green: one impulse response is time-reversed then convolved with the other. Right: corresponding phase responses. In subsequent figures, causal filters are plotted in blue, acausal in green.

Figure 5: Effects of filtering on a temporally localized event (impulse or spike). The response is spread over time (i.e. no longer precisely localized in time), and delayed if the filter is causal. The overall delay is eliminated if the filter is zero-phase (green), but the response is then acausal, i.e. part of it occurs before *the event. The response may include multiple spurious 'features' due to filter ringing.*

Figure 6: Impulse responses of typical filters. The left panel shows the impulse response time series, the right plot shows the log of its absolute value (truncated at -60 dB) on a color scale to reveal the low-amplitude tail. A: boxcar of duration 50 ms, B: same, applied using `filtfilt` *, C: Butterworth lowpass, cutoff 10 Hz, order 8, D: same as C, cutoff 20 Hz, E: same as D, applied using* `filtfilt`*, F: Butterworth bandpass, 10-20 Hz, order 8, G: same, order 2, H: same as G, 10-12 Hz. Filters plotted in green are acausal.*

Figure 7: Effects of filtering on a step-like event. The response may be smoothed (low-pass filter), and delayed (causal filter). The steady-state part may be lost (high-pass filter), and spurious 'features' may appear, some of which may occur before the event (acausal filter, green).

Figure 8: Step responses of typical filters. The left panel shows the step response time series, the right plot shows the log of its absolute value (truncated at -60 dB) on a color scale to reveal the low-amplitude tail. A: Low-pass Butterworth order 8, cutoff 10 Hz, B: same, applied using `filtfilt`*, C: High-pass Butterworth order 2, cutoff 10 Hz, D: same, order 8, E: same as D, applied using* `filtfilt`*, F: Band-pass Butterworth order 2, 10-20 Hz, G: same, 10-12 Hz, H: same as G, applied using* `filtfilt`*. Filters in green are acausal.*

Figure 9: Effects of filtering on a sinusoidal pulse. The pulse is widened and delayed (causal flter) by a bandpass filter. The delay is avoided with a zero-phase filter, but the response then starts before the event. The pattern of distortion may be more complex (here a bandpass with cutoff slightly below the pulse frequency). For a notch filter tuned to the frequency of the pulse, the suppression may be delayed and there may be a rebound artifact after the pulse. If the filter is acausal a 'rebound' artifact may also occur before *onset and offset events.*

Figure 10: High-pass filter response to a step (Butterworth order 8, cutoff as indicated in legend). The lines at the bottom of each graph indicate temporal intervals within which certain widely-reported ERP features are expected (Berenscheiftspotential BP, N1 or N100, MMN, P300, P600). These were selected for illustrative purposes; numerous other features have been reported in this range.

Figure 11: Examples of artifactual features that can arise due to high-pass filtering. A: step response of a zero-phase filter (Butterworth order 8 applied using `filtfilt`), cutoff as indicated in legend. B: response to a pulse of a Butterworth order 8 filter, cutoff as indicated in legend. C: response to a smoothed step of a Butterworth filter of order 8 and cutoff 1 Hz, step transition duration as indicated in legend. D: response to a pulse of a Butterworth filter of order 8 and cutoff 1Hz, pulse duration as indicated in legend. The lines at the bottom of each graph are as defined in Fig. 10.

Figure 12: Response of a bandpass filter (8-11 Hz, Butterworth order 8 applied with `filtfilt`) to a 10 Hz sinusoidal pulse (top) and to an input consisting of two impulses (middle and bottom).

Figure 13: Temporal Response Function estimated from simulated stimulus-response data. Black: "true" TRF. Thick blue: TRF estimated using response data that has been filtered by a causal filter (Butterworth bandpass 1-10 Hz, order 4+4). Green: same with acausal filter (Matlab's `filtfilt`).

Figure 14: Time frequency power spectral density plots in response to a Dirac pulse (A). B: STFT-based spectrogram with an analysis window of duration 0.5 s. C: wavelet-based spectrogram with an analysis window of duration 7 cycles. D and E: phase plots corresponding to B and C.

Figure 15: Left: power spectrum of data from a MEG system with a shallow (or missing) antialiasing filter. Peaks at 60, 120, 180 and 240 Hz (black arrows) probably reflect power line harmonics, but the origin of the other peaks is mysterious. They might be the result of aliasing of high-frequency sources within the environment, or of higher harmonics of 60 Hz. For example the peak at 200 Hz (red arrow) might result from aliasing of the fifth harmonic of 60 Hz. Right: ringing artifact in a MEG system with a particulary steep antialiasing filter. The magnetic field change was a sharp step.

Figure 16: Linear operations can be swapped. The filtered noisy signal is the superposition of the filtered signal and the filtered noise.