

Systematic biases in early ERP and ERF components as a result of high-pass filtering

David J. Acunzo

Graham MacKenzie

Mark C.W. van Rossum

Corresponding author: David J. Acunzo

Email: david.acunzo@ed.ac.uk

Phone: +44 131 650 3391

Mail: Room S35, 7 George Square, Edinburgh, EH8 9JZ, UK

Authors' affiliations:

David J. Acunzo:

Neuroinformatics Doctoral Training Centre

Institute for Adaptive and Neural Computation

School of Informatics

University of Edinburgh, UK

Graham MacKenzie:

Psychology Department

School of Philosophy, Psychology and Language Sciences

University of Edinburgh, UK

Mark van Rossum:

Institute for Adaptive and Neural Computation

School of Informatics

University of Edinburgh, UK

Abstract

The event-related potential (ERP) and event-related field (ERF) techniques provide valuable insights into the time course of processes in the brain. Researchers commonly filter the data to increase the signal-to-noise ratio because the neural signals are typically weak. However, filtering may distort the data, leading to false results. Using our own EEG data, we show that acausal high-pass filtering can generate a systematic bias easily leading to misinterpretations of neural activity. In particular, we show that the early ERP component C1 is very sensitive to such effects. Moreover, we found that about half of the papers reporting modulations in the C1 range used a high-pass digital filter cut-off above the recommended maximum of 0.1 Hz. Additionally, among 185 relevant peer-reviewed publications, 80 used cutoffs above 0.1 Hz. As a consequence, part of the ERP/ERF literature may need to be re-analyzed. We provide guidelines on how to minimize filtering artifacts.

Keywords (up to 6)

ERP, ERF, high-pass filtering, data processing, C1

Highlights (3 to 5 bullet points, 85 char. per point max.)

- We evaluate artifacts due to high-pass filtering used to preprocess ERP/ERF data
- High-pass filtering can lead to systematic biases between conditions
- About 40% of surveyed papers used methods that can lead to such biases
- Guidelines to minimize high-pass filtering artifacts are provided

1. Introduction

Event-related potential and event-related fields (ERP/ERF), generated by averaging electro- and magneto-encephalograph (EEG/MEG) signals respectively, provide unique insights into human brain processes with unrivalled time resolution. Because the signal is weak and noisy, raw data typically requires several pre-processing steps, including filtering and removal or attenuation of artifacts, before epoching and averaging over trials and participants. Filtering is an efficient way to increase the signal-to-noise ratio by removing frequency bands that mainly contain non-neural or irrelevant information. Both high EEG frequencies (typically above 30 or 40 Hz), as well as low frequencies (typically below 1 Hz or less) are often filtered out. There has been a recent interest in the effects of filtering on ERP signals and their interpretation (van Rullen, 2011; Rousselet, 2012), focusing mainly on low-pass filtering. The present paper focuses on high pass filtering and on its effects on early components.

Low frequency signals can be of non-neural origin, such as electrodermal activity, drying or chemical stabilization of the electrolyte due to thermal changes and contact with the skin (see e.g. Hennigausen et al., 1993 and Tallgren et al., 2005), but can be minimized (Vanhatalo, Voipio, & Kaila, 2005). In addition, the neural signal itself contains drifts and low frequencies that can have cognitive significance (Grey Walter et al., 1964; Fitzgerald et al., 2001; Monto, Palva, Voipio, & Palva, 2008; Palva & Palva, 2012; Vanhatalo, Voipio, & Kaila, 2005; Pastor et al., 2008) but might not be relevant for the study.

Low frequency signals are removed with a high-pass filter. High-pass filtering can be understood as a way to force the average signal to be zero within a time window of a certain length, thus eliminating slowly varying components. The higher the cut-off frequency (the frequency at which a 3 dB attenuation is attained), the shorter the time window, and the

shorter the signal is allowed to depart from zero. As we shall demonstrate, this can lead to significant distortions of the data.

The layout of this paper is as follows. After discussing some principles behind commonly used filters, we apply different filters to our own EEG data. We demonstrate that high-pass filtering with an excessively high cut-off frequency can introduce systematic distortions to the signal and can lead to false results and interpretations. Although a well-known textbook warns of the dangers of filters and recommends a maximum cut-off value of 0.1 Hz in high-pass filters (Luck, 2005), a review of published papers reveals that many studies do not conform to this recommendation, which may have led to false conclusions, in particular concerning early modulations.

2. Filtering alters the shape of the signal

Although filtering can improve the signal-to-noise ratio, it can also distort the signal in an unwanted manner. A commonly known distortion introduced by filtering, often described in EEG/MEG data processing software manuals, is phase delay. Phase delay shifts the frequency components in time, which is undesirable given the core importance of the event timing in ERP/ERFs. Phase delay is particularly an issue for causal filters, for which the output at a given time only depends on past and present, but not future, input. Linear frequency-dependent phase delay, meaning that all frequencies will be delayed by the same amount, can be implemented with causal Finite Impulse Response (FIR) filters. This minimizes distortion of the overall shape of the signal, but generates large delays up to hundreds of milliseconds. Infinite Impulse Response (IIR) high-pass filters, such as Butterworth or Ellipsoid filters, commonly implemented in EEG/MEG software packages, can achieve comparable filtering performance to FIR filters with less computational resources

and with less delay (see e.g. Lynn, 1989). However, their non-linear phase response can generate strong distortions of the signal.

Because of the drawbacks of causal filtering and with the advent of computerized data processing and storage, acausal filtering has gained popularity. As the output of an acausal filter depends both on the past and future input, acausal filters are applied offline on stored data. The advantage of acausal filters is that they can be constructed to have no phase delay at all, in which case they are called zero phase-shift filters. Commonly, these acausal filters are implemented with a causal filter run twice over the data: once forwards and once backwards. Apart from doubling the order of the filter, the backwards pass counterbalances any delays that the forward run introduces. Because of the absence of delay and the reduced distortions induced by forward-backward filters, guidelines and software manuals more or less explicitly advise their use (see e.g. Picton et al., 2000).

However, acausal filters can also be a cause for aberrant misinterpretations in the timing of neural processes. Figure 1A shows the effect of a causal Butterworth filter applied forwards, and applied both forwards and backwards for two artificial example signals. Due to the non-linearity of the phase response of the Butterworth filter, the shape of the signal is significantly distorted in the causal case (top row). Note however that all the distortions happen after signal onset. In the forward-backward case, distortions are remarkably reduced compared to the causal case, but the signal is distorted more than one second prior to signal onset (bottom row).

----- Figure 1 about here -----

These distortions can lead to misinterpretation, in particular in the typical situations where one studies the onset of a particular component or the onset of a divergence between two waveforms. Figure 1B illustrates this point, showing two signals that are initially the same, but differ after one second (top). When a causal filter is applied, the timing of the

divergence of the two signals is preserved (middle). However, after applying an acausal (zero phase-shift) filter, differences between the two signals are observable where they were identical before filtering, even before $t = 0$, when the two signals onset (bottom). The difference in the later part contaminates the early part of the waveform.

With a causal filter, a difference in the unfiltered waveforms may lead to differences later in time, but cannot affect earlier timepoints, supporting inferences that the effect started at this time *at the latest*. However, with an acausal filter, distortions are spread both forward *and* backward in time. In this case, it is not possible to state when the effect starts: it can only be concluded that the waveforms differed at some point in time.

3. Effects of filtering on actual EEG data

To illustrate the problematic effects of filtering in practice, we submitted some of our own data to various high-pass filters. The original purpose of the experiment was to investigate the effects of spatial attention and facial expression on the C1 component of the visual ERP. The C1 component is characterized by a widespread centro-parietal negativity (positivity) peaking before 100 ms, evoked by presenting stimuli in the upper (lower) hemifield and using an average mastoid reference (Clark, Fan, & Hillyard, 1995). The results of this experiment and its detailed interpretation will be presented elsewhere (Acunzo et al., in preparation). Using these data, we show how high-pass filters can affect the shape of a real waveform, and how subsequent components can add systematic biases to earlier components and lead to erroneous interpretations. In addition, we show that filter parameters are critical for the proper interpretation of early components, and more particularly the C1 component.

3.1. Methods

3.1.1. Stimuli and Procedure

Stimuli and procedure will be described in more detail elsewhere (Acunzo et al., in preparation). All procedures were approved by the Psychology Department Ethical Committee at the University of Edinburgh. Twenty-four right-handed participants were first presented with an arrow near the fixation point, pointing left or right, for 200 ms. After an interstimulus interval of 750 ms, a facial stimulus was presented for 300 ms on the side congruent or incongruent to the cue arrow, presented on the upper visual hemifield to elicit a negative C1 (Jeffreys & Axford, 1972; Clark, Fan, & Hillyard, 1995). Facial stimuli showed three expressions: neutral, fearful, and happy. The experiment consisted of 1120 trials per participant, with 352 trials for the Happy facial expression condition, and 768 trials equally divided between the Neutral and Fearful conditions. Participants were instructed to press a button when detecting a happy face at the congruent location only (176 trials) only, while fixating on the fixation point.

3.1.2. Data Acquisition

EEG was recorded using a BioSemi Active-Two system (BioSemi B.V., Amsterdam, Netherlands), which has a DC coupled amplifier. The activity at 64 Ag-AgCl scalp electrodes following the location and label of the extended 10-20 system (Jasper, 1958), along with 4 EOG electrodes (above and below the right eye, and on the outer canthi) and 2 mastoid electrodes, was digitized on 24 bits with a sample rate of 1024 Hz.

3.1.3. Data Processing

Pre-processing was performed using the EEGLAB toolbox (Delorme & Makeig, 2004) under Matlab and custom scripts. The EEG signal was first re-referenced to the average mastoids, and low-pass filtered with a cut-off value of $f_c = 40$ Hz, using the default FIR filter implemented in EEGLAB: a least square linear-phase filter of order 75 with a transition bandwidth (the range of frequencies between the bandcut and the bandpass) of 6 Hz, ran forward and backward. Though it should be kept in mind that low-pass filters may

also induce artifacts (vanRullen, 2011), the low amplitude of the higher frequencies in the EEG signal reduces the risk of serious alteration of the waveform. Rousselet (2012) found little artifactual effect of low-pass filtering on real EEG data. Furthermore, as this preliminary filtering step is done for all subsequent high-pass filtering conditions, it does not alter our conclusions on high-pass filtering. The continuous data were then resampled to 256 Hz. High-pass filters with cut-offs f_c of 0.05, 0.1, 0.5, and 1 Hz were then applied to study the effects of filtering. Testing lower cut-off values was found to be unnecessary as little effect was observable up to 0.1 Hz. We used the default EEGLAB parameters for this filtering step as well: the filters used were least square linear phase FIR high-pass filters ran forward and backward. Their transition bandwidth was $0.15 \times f_c$ and their order $3 \times \text{fix}\left(\frac{f_s}{f_c}\right)$, with f_s the sampling frequency (256 Hz), and fix the function that rounds downwards to nearest integer. For each of the filtered sets of raw data, epoching was performed using facial stimulus onset time as time origin, and each channel was baselined using the 100 ms interval preceding stimulus onset. To remove artifactual epochs, a semi-automatic procedure was ran on the non high-pass filtered data, labelling epochs containing EOG data beyond 70 μV in absolute value. The data were visually inspected and epochs containing artifacts were removed. The same trials were used for all filtering conditions. On average, 936 correctly answered and artifact-free trials per participant contributed to the grand-average ERP.

3.2. Data Analysis

Voltage from electrodes P1, P2, CP1 and CP2 was averaged to generate the ERPs. For the C1 component, mean amplitude of the 50-100 ms interval was calculated. For the purpose of this paper, only results from the Congruent attentional condition are presented here. A one-way ANOVA with Emotion (Neutral, Fearful, Happy) was performed for each filter cut-off value. We tested for normality using a Kolmogorov-Smirnov test, and for sphericity using a

Mauchly test. Normality hypotheses were confirmed. When indicated, a Greenhouse-Geisser (GG) correction was applied to compensate for sphericity violations.

4. Results

Figure 2 shows grand-average ERPs for each emotional condition (Neutral, Fearful, Happy) averaged over all participants and the electrodes CP1, CP2, P1 and P2. Figure 2A shows the waveforms without high-pass filtering. Figure 2C zooms in on the waveform between -100 and 500 ms, containing the C1. After reaching a negative peak at 100 ms (as expected with stimuli presented in the upper visual hemifield, see Clark, Fan, & Hillyard, 1995), the waveforms exhibit a globally positive-going excursion, peaking at around 500 ms. The amplitude of this deflection is largest in the Happy condition, which corresponds to target trials during which participants had to press a button.

High-pass filtering strongly affected the shape of the waveforms. As the cut-off frequency of the filter is increased, the late positive component disappears (Fig. 2B, E, F). But simultaneously, the early part of the waveform is pushed downwards. This effect is similar to the distortion observed in Figure 1A (bottom row). Only when a low cut-off frequency of 0.1 Hz was used, the waveform was not drastically affected (Fig. 2D).

In parallel with its impact on the shape of the ERP, scalp topographies (averaged across latency periods 50-100 ms) are affected by filtering (Fig. 2C-F). As a consequence, high-pass filtering may alter results for studies using topography information, such as microstate analysis or source reconstruction.

----- Figure 2 about here -----

Next, we examined the apparent modulation of the C1 amplitude with experimental condition. Figure 3A shows that the modulation changes strongly as a function of the cut-off value. Significance tests (Fig. 3B) show that, in our experiment, reliable effects arise as the

filter cut-off frequency increases, at 0.5 Hz ($F_{2,46} = 9.38$, $p < 10^{-3}$) and 1 Hz ($F_{2,46} = 9.53$, $p < 10^{-3}$), while there is no significant effect for lower cut-off values.

----- Figure 3 about here -----

To better understand the mechanisms behind these effects, we tested directly whether the C1 effect observed at high cut-offs was due to a subsequent late component that contaminates C1 through the non-causality of the filter used. We measured the amplitude of the late positive complex (LPC), which is known to be affected by facial expression and target recognition (Eimer & Holmes, 2002; Picton, 1992; Rozenkrants & Polich, 2008). To quantify the LPC, we averaged the waveform between 400 and 500 ms. Consistent with the literature, the LPC showed a larger positive excursion for the Happy faces target than for the other conditions ($F_{1,17,27,01} = 34.05$, $p < 10^{-5}$, GG corrected). We defined ΔLPC_{HN} (and ΔLPC_{FN}) as the change in LPC amplitude in the Happy versus Neutral conditions (resp. Fearful versus Neutral), in the *unfiltered* condition. Similarly, $\Delta C1_{HN}$ (and $\Delta C1_{FN}$) is the difference of C1 amplitude between the Happy and Neutral conditions (resp. Fearful and Neutral) in the *unfiltered* condition, and $\Delta C1_{NH}^{f_c=1}$ and $\Delta C1_{FH}^{f_c=1}$ in the *filtered* condition with a 1 Hz cut-off. We calculated these values for each participant.

To evaluate the link between the LPC and the C1 amplitude, we correlated ΔLPC with $\delta C1$, defined as the change in $\Delta C1$ caused by the filtering ($\delta C1 = \Delta C1^{f_c=1} - \Delta C1$). $\delta C1$ is positive for most participants, as the C1 effect is larger with the high-pass filter than without (see e.g. Fig. 2C versus 2F). Figure 4 shows the correlation between the Neutral/Happy modulation ΔLPC_{HN} and the change in C1 effect $\delta C1_{HN}$, and similarly for the Neutral/Fearful modulation (ΔLPC_{FN} and $\delta C1_{FN}$). We found a strong correlation in both cases ($\rho_{HN} = 0.95$, $p < 10^{-8}$, and $\rho_{FN} = 0.71$, $p = 10^{-4}$). The first correlation reflects the contamination by the large late positive amplitude in the Happy condition that can be observed in Figure 2. The correlation in the Neutral/Fearful case is noteworthy and shows

subtler effects. Although there is no statistically significant LPC effect ($t_{23} = 1.01, p = 0.32$), a clear correlation exists on the individual level, contributing to the significant C1 Neutral/Fearful effect in the filtered waveform ($t_{23} = 2.06, p = 0.05$). This shows that portions of the grand-average waveform that do not show any effect can still constitute a subtle contributor to distortions leading to misinterpretation of the data.

----- Figure 5 about here -----

To ensure that the C1 effect observed in the filtered waveform was indeed due to a subsequent component, we used a causal filter (4th order high-pass Butterworth filter with 1 Hz cut-off). While the waveform also appears distorted when compared with the waveform without high-pass filtering (see Fig. 5), we failed to find a significant C1 effect ($F_{2,46} < 1$). Together, these results confirm that the C1 effect in our data is due to subsequent components, in particular the LPC, and is completely artifactual.

5. Discussion

In summary we showed that high-pass filtering of the raw EEG data can distort the resulting ERP waveforms, and induce systematic biases between conditions. For this reason, high-pass filtering should be used parsimoniously in EEG/MEG data, and the cut-off frequency should be kept as low as possible under normal circumstances. Our analysis on real EEG data shows how a late slow component can induce a systematic bias in an earlier component when using commonly used zero phase-shift filters, which can lead to erroneous interpretations. We found a highly statistically significant modulation of C1 by stimulus condition, but showed that it was completely artifactual.

The recent years have been characterized by an increasing interest in the very early visual components, happening before 100 ms after stimulus onset, and thought to be associated with activity in the striate (Jeffreys & Axford, 1972; Clark, Fan, & Hillyard,

1995), and perhaps extrastriate visual cortices (Foxe & Simpson, 2002; Ales, Yates & Norcia, 2010). When peaking, these components are often labeled C1 and M90 for the event-related potential (ERP) and event-related field (ERF) waveforms respectively. These early processes, and in particular, the C1 component, were thought to be only modulated by the physical characteristics of the triggering stimulus, but immune to endogenous modulations, and more specifically attention (Martínez et al. 1999; Noesselt et al. 2002; Di Russo et al. 2003; Hillyard et al. 1998). More recently, however, C1 modulations by attention and attentional load (Khoe et al. 2005; Wu et al. 2005; Kelly et al. 2008; Fu et al. 2009; Rauss et al. 2009; Fu et al. 2010a), perceptual learning (Pourtois et al., 2008), anxiety (Eldar et al., 2010), and emotional stimuli or faces (Pourtois et al., 2004; Stolarova et al., 2006; West et al., 2011) were reported. Similarly, very early ERF modulations by attention (e.g. Poghoshyan & Ioannides, 2008; Ioannides & Poghoshyan, 2012) and facial expression (e.g. Batty et al., 2009; Bayle & Taylor, 2009) have been reported. (For a review on early modulations, see Rauss, 2011.) While these results are extremely exciting, as they suggest that our brain processes and modulates visual information more quickly than is generally thought, difficulties of replication (e.g. Santesso et al., 2008; Fu et al., 2010) make it hard to draw conclusions on what mechanisms are at play. One possible reason for such inconsistencies is that experimental paradigms and methods vary widely.

In light of our findings, we examined the high-pass filter cut-off value used in these studies. We found that out of the 13 ERP/ERF attention studies finding a very early effect, 5 used a cut-off higher than 0.5 Hz. Similarly, 6 out of 10 studies looking at early effects of facial expression used a cut-off higher than 0.5 Hz. We found only one study reporting no early effect while using a high cut-off.

The inconsistency of results may therefore be a combination of a filtering artifact, publication bias, and genuine early effects. Future research in this area should seek to

elucidate the conditions under which very early components are modulated by taking care to minimize possible biases induced by high-pass filtering (see Summary and Guidelines).

Additionally, we inspected papers from *Journal of Neuroscience*, *Cerebral Cortex*, *Neuroimage*, *Human Brain Mapping*, and the *Journal of Cognitive Neuroscience*, published or accepted between January 2011 and March 2012. We only included papers using ERP or ERF amplitudes either directly, to find neural sources, or for Independent Component Analysis, and ignored studies interested in restricted frequency bands (event-related oscillations, event-related synchronizations/desynchronizations, time-frequency analyses, and power analyses). Although the issues presently discussed are also relevant for intra-cranial EEG, we only focused on data measured from scalp measurements. Our results are consistent with those reported by Rousselet (2012): out of the 185 scrutinized studies satisfying our criteria, we found that 80 (43%) of them used a cut-off above 0.1 Hz. Half of those, i.e. 40 studies (21 %), used a cut-off of 1 Hz or higher. The proportion of studies using a cut-off higher than 0.1 Hz and using MEG was comparable to those using EEG (18/40, or 45% vs 65/155, or 42%). The vast majority of these reports do not specify whether a causal or acausal filter was used, but the few which do used a zero phase-shift filter. Additionally, most papers do not specify why such a high cut-off value was used, and those who do invoked the classical reasons of signal-to-noise ratio increase and slow trends removal.

While the potential dangers of filtering were already described in Luck (2005), and while the majority of laboratories use digital high-pass filtering knowingly and parsimoniously, it appears that it is necessary to change the practices of more than a third of the works, and to possibly reinterpret a non-negligible portion of the literature. However, we do not claim that the conclusions of all reports using a high cut-off frequency are erroneous. The nature and extent of biases induced by filtering depend on the shape of the waveform, the

type of filter used, and the specific component studied. Rather, these studies should be regarded keeping the possibility of a filter-induced bias in mind.

7. Summary and Guidelines

The following guidelines should be considered to minimize artifacts due to high-pass filtering:

- Only use high-pass offline filtering if necessary. Visually inspect the data beforehand to judge if too much drift is present. Only if the data are indeed noisy and shows a large amount of drift, consider applying a high-pass filter.
- To choose the cut-off of the high-pass filter, we reiterate the guideline stated in Luck (2005): set the high-pass filter cut-off value to 0.01 Hz by default. For less docile participants whose data may be noisier, such as children or certain kinds of patients, consider the possibility of a higher cut-off, such as 0.05 or 0.1 Hz maximum.
- If, for any reason, a filter with a higher cut-off is applied, check the results against data filtered with a lower cut-off. Try to understand any discrepancy in the results and make sure they are not due to biases induced by filtering (e.g. with analyses as used in the present paper).
- To check the distortion induced by a filter, generate a grand average waveform with and without having applied a filter on the raw data, and ensure that the overall morphology is not affected.
- Zero-phase shift filters do not delay the signal but can generate distortions backward in time. If one is interested in the earliest moment when an effect occurs, a causal filter is preferred. If one is interested in the timing of a peak, a zero phase-shift filter should be preferred.

- Finally, as stated in Picton et al. (2000), the nature of the filter used should be specified. In particular, a vast majority of the reports omit to mention the causality of the filter, which, as we saw, may be critical for the interpretation of the data.

Acknowledgments

We thank Baris Demiral for technical assistance, and Simon Kelly for bringing this issue of filtering to our attention.

Funding

This research was supported by the Doctoral Training Centre in Neuroinformatics, itself co-funded by the EPSRC, MRC and BBSRC.

References

- Ales, J. M., Yates, J. L., & Norcia, A. M. (2010). V1 is not uniquely identified by polarity reversals of responses to upper and lower visual field stimuli. *NeuroImage*, *52*(4), 1401–1409.
- Bayle, D. J., & Taylor, M. J. (2009). Attention inhibition of early cortical activation to fearful faces. *Brain Research*, *1313*, 113–123.
- Delorme, A., & Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of neuroscience methods*, *134*(1), 9–21.
- Di Russo, F., Martínez, A., & Hillyard, S. A. (2003). Source Analysis of Event-related Cortical Activity during Visuo-spatial Attention. *Cerebral Cortex*, *13*, 486–499.
- Eimer, M., & Holmes, A. (2002). An ERP study on the time course of emotional face processing. *Neuroreport*, *13*(4), 427–431.

- Eldar, S., Yankelevitch, R., Lamy, D., & Bar-Haim, Y. (2010). Enhanced neural reactivity and selective attention to threat in anxiety. *Biological Psychology*, *85*(2), 252–257.
doi:10.1016/j.biopsycho.2010.07.010
- Fitzgerald, R. D., Lamm, C., Oczenski, W., Stimpfl, T., Vycudilik, W., & Bauer, H. (2001). Direct current auditory evoked potentials during wakefulness, anesthesia, and emergence from anesthesia. *Anesthesia and analgesia*, *92*(1), 154–160.
- Foxe, J. J., & Simpson, G. (2002). Flow of activation from V1 to frontal cortex in humans. *Experimental Brain Research*, *142*(1), 139–150.
- Fu, S., Huang, Y., Luo, Y.-J., Wang, Y., Fedota, J. R., Greenwood, P. M., & Parasuraman, R. (2009). Perceptual load interacts with involuntary attention at early processing stages: Event-related potential studies. *NeuroImage*, *48*(1), 191–199. Elsevier Inc.
doi:10.1016/j.neuroimage.2009.06.028
- Fu, S., Fedota, J. R., Greenwood, P. M., & Parasuraman, R. (2010). Early interaction between perceptual load and involuntary attention: An event-related potential study. *Neuroscience Letters*, *468*(1), 68–71. doi:10.1016/j.neulet.2009.10.065
- Fu, S., Fedota, J. R., Greenwood, P. M., & Parasuraman, R. (2010). Dissociation of visual C1 and P1 components as a function of attentional load: An event-related potential study. *Biological Psychology*, *85*(1), 171–178. doi:10.1016/j.biopsycho.2010.06.008
- Grey Walter, W., Cooper, R., Aldridge, V. J., McCallum, W. C., & Winter, A. L. (1964). Contingent Negative Variation: An Electric Sign of Sensorimotor Association and Expectancy in the Human Brain. *Nature*, *203*, 380–384.
- Hennighausen, E., Heil, M., & Rosler, F. (1993). A correction method for DC drift artifacts. *Electroencephalography and clinical neurophysiology*, *86*, 199–204.
- Hillyard, S. A., Vogel, E. K., & Luck, S. J. (1998). Sensory Gain Control (Amplification) as a Mechanism of Selective Attention: Electro-Physiological and Neuroimaging Evidence.

- Philosophical Transactions: Biological Sciences*, 353(1373), 1257–1270.
- Ioannides, A. A., & Poghosyan, V. (2012). Spatiotemporal dynamics of early spatial and category-specific attentional modulations. *NeuroImage*, 60(3), 1638–1651.
- Jasper, H. A. (1958). The ten-twenty system of the international federation. *Electroencephalography and clinical neurophysiology*, 10, 371–375.
- Jeffreys, D. A., & Axford, J. (1972). Source locations of pattern-specific components of human visual evoked potentials. I. Component of Striate Cortical Origin. *Experimental Brain Research*, 16(1), 1–21.
- Kelly, S. P., Gomez-Ramirez, M., & Foxe, J. J. (2008). Spatial Attention Modulates Initial Afferent Activity in Human Primary Visual Cortex. *Cerebral Cortex*, 18(11), 2629–2636. doi:10.1093/cercor/bhn022
- Khoe, W., Mitchell, J. F., Reynolds, J. H., & Hillyard, S. A. (2005). Exogenous attentional selection of transparent superimposed surfaces modulates early event-related potentials. *Vision Research*, 45(24), 3004–3014. doi:10.1016/j.visres.2005.04.021
- Lehmann, D., & Skrandies, W. (1980). Reference-free identification of components of checkerboard-evoked multichannel potential fields. *Electroencephalography and clinical neurophysiology*, 48(6), 609–621.
- Lynn, P. A. (1989). *An Introduction to the Analysis and Processing of Signals* (Third Edition.). New York: Hemisphere Publishing Corporation.
- Luck, Steven J. (2005). *An Introduction to the Event-Related Potential Technique*. Cambridge, Massachusetts: The MIT Press.
- Martínez, A. et al. (1999). Involvement of striate and extrastriate visual cortical areas in spatial attention. *Nature Neuroscience*, 2(4), 364–369.
- Monto, S., Palva, S., Voipio, J., & Palva, J. M. (2008). Very Slow EEG Fluctuations Predict the Dynamics of Stimulus Detection and Oscillation Amplitudes in Humans. *Journal of*

- neuroscience*, 28(33), 8268–8272. doi:10.1523/JNEUROSCI.1910-08.2008
- Noesselt, T. et al. (2002). Delayed Striate Cortical Activation during Spatial Attention. *Neuron*, 35, 575–587.
- Palva, J. M., & Palva, S. (2012). Infra-slow fluctuations in electrophysiological recordings, blood-oxygenation-level-dependent signals, and psychophysical time series. *NeuroImage*, 1–11. Elsevier Inc. doi:10.1016/j.neuroimage.2012.02.060
- Pastor, M. C., Bradley, M. M., Löw, A., Versace, F., Moltó, J., & Lang, P. J. (2008). Affective picture perception: Emotion, context, and the late positive potential. *Brain Research*, 1189, 145–151. doi:10.1016/j.brainres.2007.10.072
- Picton, T. W. (1992). The P300 wave of the human event-related potential. *Journal of clinical neurophysiology : official publication of the American Electroencephalographic Society*, 9(4), 456–479.
- Picton, T. W., Bentin, S., Berg, P., Donchin, E., Hillyard, S. A., Johnson, R., Miller, G. A., et al. (2000). Guidelines for using human event-related potentials to study cognition: recording standards and publication criteria. *Psychophysiology*, 37(2), 127–152.
- Poghosyan, V., & Ioannides, A. A. (2008). Attention Modulates Earliest Responses in the Primary Auditory and Visual Cortices. *Neuron*, 58(5), 802–813.
- Pourtois, G., Grandjean, D., Sander, D., & Vuilleumier, P. (2004). Electrophysiological Correlates of Rapid Spatial Orienting Towards Fearful Faces. *Cerebral cortex (New York, NY : 1991)*, 14(6), 619–633. doi:10.1093/cercor/bhh023
- Pourtois, G., Rauss, K. S., Vuilleumier, P., & Schwartz, S. (2008). Effects of perceptual learning on primary visual cortex activity in humans. *Vision Research*, 48(1), 55–62. doi:10.1016/j.visres.2007.10.027
- Rauss, K. S., Pourtois, G., Vuilleumier, P., & Schwartz, S. (2009). Attentional load modifies early activity in human primary visual cortex. *Human Brain Mapping*, 30(5), 1723–1733.

doi:10.1002/hbm.v30:5

- Rauss, K. S., Schwartz, S., & Pourtois, G. (2011). Top-down effects on early visual processing in humans: A predictive coding framework. *Neuroscience and Biobehavioral Reviews*, 35(5), 1237–1253. doi:10.1016/j.neubiorev.2010.12.011
- Rousselet, G. A. (2012). Does filtering preclude us from studying ERP time-courses? *Frontiers in Psychology*, 3, 131. doi:10.3389/fpsyg.2012.00131
- Rozenkrants, B., & Polich, J. (2008). Affective ERP processing in a visual oddball task: Arousal, valence, and gender. *Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology*, 6. doi:10.1016/j.clinph.2008.07.213
- Santesso, D. L., Meuret, A. E., Hofmann, S. G., Mueller, E. M., Ratner, K. G., Roesch, E. B., & Pizzagalli, D. A. (2008). Electrophysiological correlates of spatial orienting towards angry faces: a source localization study. *Neuropsychologia*, 46(5), 1338–1348. doi:10.1016/j.neuropsychologia.2007.12.013
- Stolarova, M., Keil, A., & Moratti, S. (2006). Modulation of the C1 visual event-related component by conditioned stimuli: evidence for sensory plasticity in early affective perception. *Cerebral Cortex*, 16(6), 876–887. doi:10.1093/cercor/bhj031
- Tallgren, P., Vanhatalo, S., Kaila, K., & Voipio, J. (2005). Evaluation of commercially available electrodes and gels for recording of slow EEG potentials. *Clinical Neurophysiology*, 116(4), 799–806. doi:10.1016/j.clinph.2004.10.001
- Teder-Sälejärvi, W. A., McDonald, J. J., Di Russo, F., & Hillyard, S. A. (2002). An analysis of audio-visual crossmodal integration by means of event-related potential (ERP) recordings. *Brain research. Cognitive brain research*, 14(1), 106–114.
- Vanhatalo, S., Voipio, J., & Kaila, K. (2005). Full-band EEG (FbEEG): an emerging standard in electroencephalography. *Clinical Neurophysiology*, 116(1), 1–8.

doi:10.1016/j.clinph.2004.09.015

VanRullen, R. (2011). Four common conceptual fallacies in mapping the time course of recognition. *Frontiers in Psychology*, 2, 365. doi:10.3389/fpsyg.2011.00365/abstract

West, G. L., Anderson, A. K., Ferber, S., & Pratt, J. (2011). Electrophysiological evidence for biased competition in V1 for fear expressions. *Journal of Cognitive Neuroscience*, 23(11), 3410–3418. doi:10.1162/jocn.2011.21605

Wu, Y., Chen, J., & Han, S. (2005). Neural mechanisms of attentional modulation of perceptual grouping by collinearity. *Neuroreport*, 16(6), 567–570.

Figures

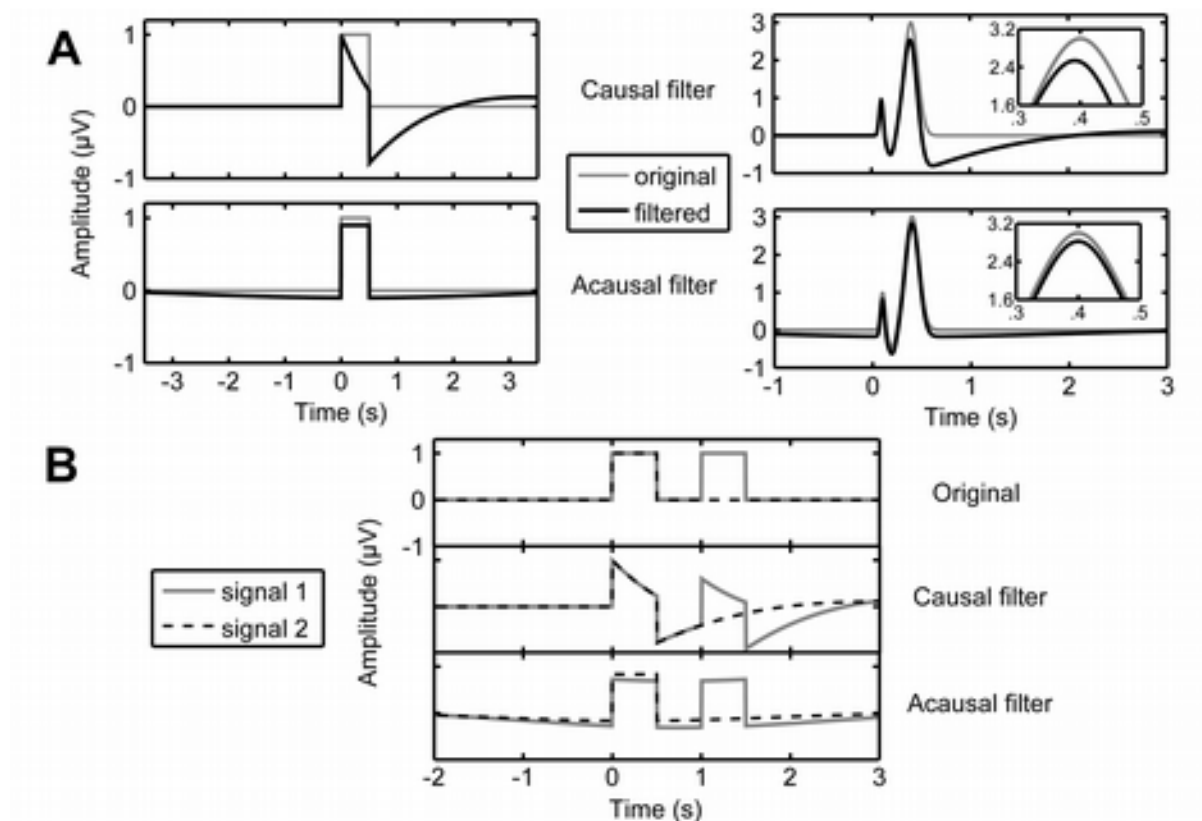


Figure 1. The effect of causal and acausal high-pass filtering on artificial signals. **(A):** A boxcar (left) and an artificial waveform (right) were filtered with one forward pass of the

filter (top), and with a forward and backward pass (bottom). The zero-phase filter preserves the shape of the signal better than the causal filter, and, as opposed to the causal filter, preserves the latency of the third peak of the artificial waveform (see detail). However, it induces distortion *before* the onset of the signal. **(B)**: Two artificial signals differing only between 1 and 1.5 s (top). When filtered with the causal filter (middle), the two signals differ from $t = 1$ s, like the original signals. The zero phase-shift filter (bottom), however, introduces differences even before the onset of both signals at $t = 0$. The filter was a 3rd order Butterworth filter of cut-off frequency 0.5 Hz.

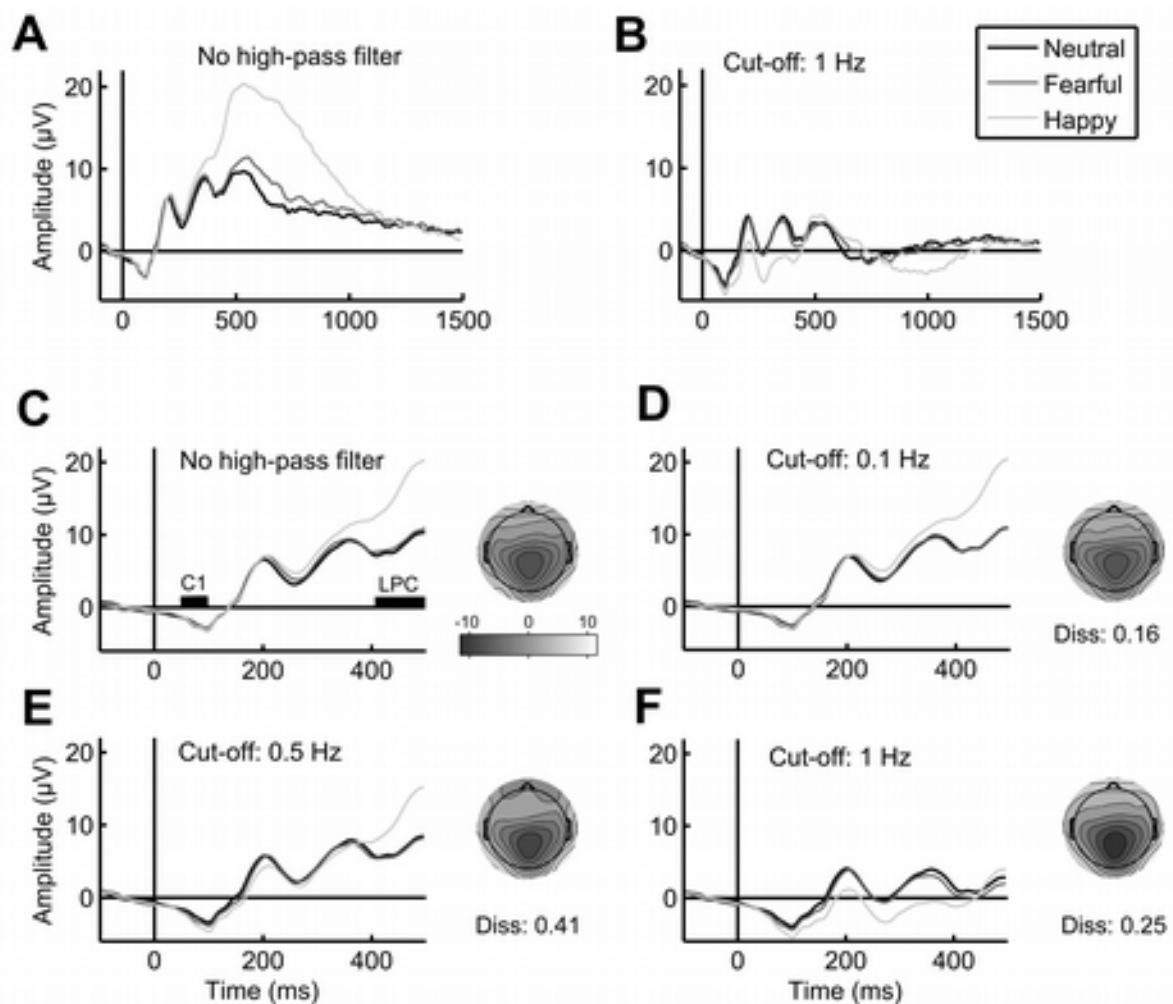


Figure 2. The effect of filtering on ERP data of an experiment on emotional face processing. ERPs are averaged over all 24 participants, and over electrodes CP1, CP2, P1 and P2, at various high-pass cut-off frequencies, and for each experimental condition. **(A, B)**: ERPs from -100 to 1500 ms, without a high-pass filter **(A)**, and with a zero phase-shift high-pass filter, a filter of cut-off of 1 Hz applied forward and backward **(B)**. The morphology of the waveform is drastically affected. The slow and later component disappears with filtering, and the earlier components are also pushed downwards. **(C-F)**: Zoom of the ERPs up to 500 ms post-stimulus onset, high-pass filtered with a zero phase-shift filter, with different cut-off frequencies. The time windows used to quantify C1 (50-100 ms) and LPC (400-500 ms) are indicated in **(C)**. Scalp topographies calculated from the C1 time window is shown for each filter condition. The potential was re-referenced to the average. Dissimilarity indicates how each map is different from the non-filtered map. A dissimilarity of 0 means that the two maps are identical, and a dissimilarity of 2 means the two maps are inverted (see Lehmann and Skrandies, 1980). Due to eye-blink artifacts occurring after 500 ms, the ERPs shown in **(A)** and **(B)** were calculated using a subset of the epochs used in **(C-F)**.

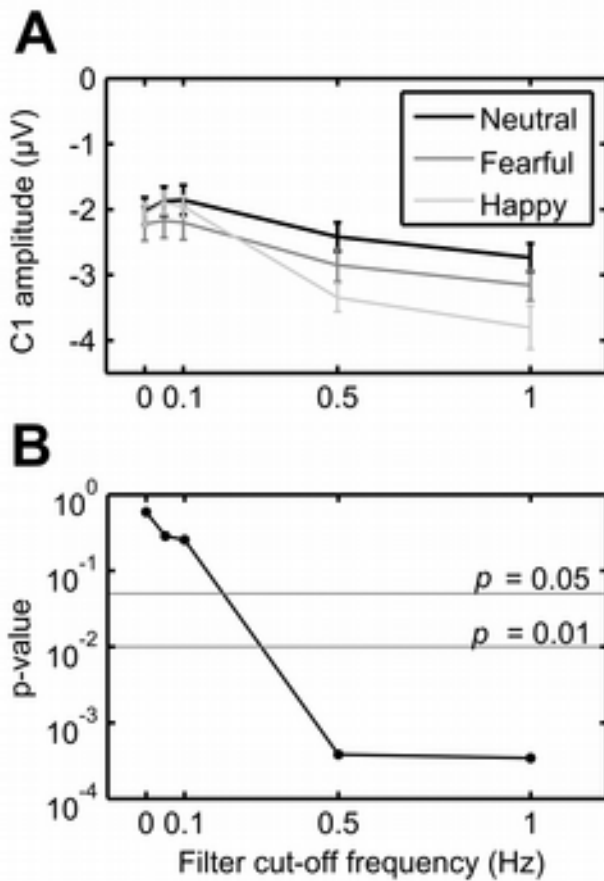


Figure 3. Effects of high-pass cut-off frequency on the C1 component across experimental conditions. **(A)**: Average amplitude at electrodes CP1, CP2, P1 and P2, between 50 and 100 ms, for each experimental condition (Neutral, Fearful, Happy), and for different cut-off values of the zero-phase shift FIR filter. **(B)**: P-value from one-way ANOVAs. High cut-offs lead to highly significant, but erroneous, results. A cut-off of 0 Hz indicates that no high-pass filter was applied.

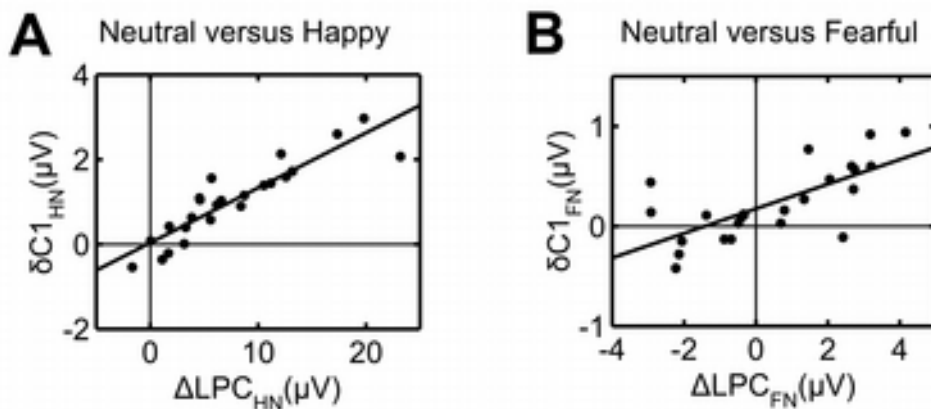


Figure 4. Correlation between the LPC effect in the unfiltered condition and the change of C1 effect, in the Neutral and Happy conditions **(A)** and the Neutral and Fearful conditions **(B)**. Each point corresponds to a participant. The significant correlations show that the observed C1 effect is mainly due to a systematic bias generated by a subsequent LPC effect, observable in the unfiltered waveforms (see Fig. 2A and 2C).

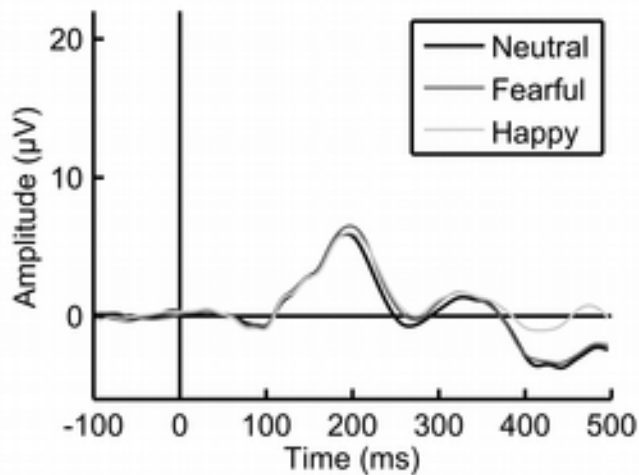


Figure 5. ERP waveforms generated with data subjected to a causal high-pass filter. The filter used was a 4th order Butterworth filter of cut-off 1 Hz. The shape of the waveform is strongly altered, but no early artifactual C1 effect is present in comparison to the ERP generated after applying a zero phase-shift high-pass filter (cf. Fig. 2F). This is consistent with the hypothesis that the C1 effect observed with a zero-phase filter at the same cut-off value was due to signal differences occurring subsequently.