

INNOVATIVE METHODOLOGY |

## Differentially Variable Component Analysis: Identifying Multiple Evoked Components Using Trial-to-Trial Variability

Kevin H. Knuth,<sup>1</sup> Ankoor S. Shah,<sup>2,3,5</sup> Wilson A. Truccolo,<sup>6</sup> Mingzhou Ding,<sup>7</sup> Steven L. Bressler,<sup>8</sup> and Charles E. Schroeder<sup>3,4</sup>

<sup>1</sup>Department of Physics, University at Albany, State University of New York, Albany; <sup>2</sup>Department of Neuroscience, Albert Einstein College of Medicine, Bronx; <sup>3</sup>Cognitive Neuroscience and Schizophrenia Program, Nathan Kline Institute, Orangeburg; <sup>4</sup>Department of Psychiatry, Columbia College of Physicians and Surgeons, New York, New York; <sup>5</sup>Department of Internal Medicine, Yale University School of Medicine, New Haven, Connecticut; <sup>6</sup>Neuroscience Department, Brown University, Providence, Rhode Island; <sup>7</sup>The J. Crayton Pruitt Family Department of Biomedical Engineering, University of Florida, Gainesville; and <sup>8</sup>Center for Complex Systems and Brain Sciences, Florida Atlantic University, Boca Raton, Florida

Submitted 24 June 2005; accepted in final form 2 February 2006

**Knuth, Kevin H., Ankoor S. Shah, Wilson Truccolo, Mingzhou Ding, Steven L. Bressler, and Charles E. Schroeder.** Differentially variable component analysis: identifying multiple evoked components using trial-to-trial variability. *J Neurophysiol* 95: 3257–3276, 2006. First published February 8, 2006; doi:10.1152/jn.00663.2005. Electric potentials and magnetic fields generated by ensembles of synchronously active neurons, either spontaneously or in response to external stimuli, provide information essential to understanding the processes underlying cognitive and sensorimotor activity. Interpreting recordings of these potentials and fields is difficult because detectors record signals simultaneously generated by various regions throughout the brain. We introduce a novel approach to this problem, the differentially variable component analysis (dVCA) algorithm, which relies on trial-to-trial variability in response amplitude and latency to identify multiple components. Using simulations we demonstrate the importance of response variability to component identification, the robustness of dVCA to noise, and its ability to characterize single-trial data. We then compare the source-separation capabilities of dVCA with those of principal component analysis and independent component analysis. Finally, we apply dVCA to neural ensemble activity recorded from an awake, behaving macaque—demonstrating that dVCA is an important tool for identifying and characterizing multiple components in the single trial.

### INTRODUCTION

Neurophysiology relies on the analysis of electric potentials or magnetic fields produced by the brain in response to sensory stimulation or in association with its cognitive and/or motor operations. These fields arise from transmembrane current flow produced by multiple ensembles of synchronously firing neurons. The underlying neural ensembles, also called generators or sources, are often dynamically coupled in unknown ways that are of interest to the experimenter. As a result of the property of linear superposition of electric currents and magnetic fields, both invasive and noninvasive electroencephalographic (EEG) recordings and magnetoencephalographic

(MEG) recordings reflect linear mixtures of the activity from these sources in addition to ongoing background activity and sensor noise. Thus even in single-trial recordings, the individual responses of each of the sources are mixed within the recorded signal, making it difficult to identify them and to study their dynamical interactions. Furthermore, it is standard practice to enhance the signal-to-noise ratio by averaging event-related potentials (ERPs) or fields (ERFs) over experimental trials. However, implicit in this construction is the assumption that the evoked waveform is constant over trials and that any variability represents noise. In this practice, the possibility of assessing trial-dependent effects in the data is unobtainable.

The last decade has seen great developments in blind source separation (BSS) and independent component analysis (ICA) techniques, such as Infomax ICA (Bell and Sejnowski 1995), FastICA (Hyvärinen and Oja 1997), and second-order blind identification (SOBI) (Belouchrani et al. 1993). These algorithms have been useful in identifying sources in EEG and MEG signals using both ensemble-averaged data (Makeig et al. 1997; Särelä et al. 1998; Vigário et al. 1999) and single trials (Cao et al. 2000; Jung et al. 1999; Makeig et al. 2002; Tang et al. 2002). However, along with its strengths, each technique has limitations, and often these limitations can be addressed. ICA, for example, allows reliable source (component) separation with minimal a priori assumptions and constraints. Its limitation is that although trial-to-trial variability can assist in separation, these effects are not explicitly considered and quantified, and these are substantial opportunities missed. Also, like many other techniques, ICA solves for maximal independence of components, despite the fact that components are often dynamically coupled. Thus although ICA may be reasonable for source separation per se, it is not explicitly designed to quantify the dynamical interactions between the neuronal ensembles that generate the components. This is one

Address for reprint requests and other correspondence: K. Knuth, Department of Physics, University at Albany (SUNY), Albany, NY 12222 (E-mail: kknuth@albany.edu).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked "advertisement" in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

of the most interesting and important aspects of the behavior of brain function and a major focus of our efforts.

Here we describe a model of the sensory-evoked neural response that is more realistic than previous models in that it explicitly models trial-to-trial amplitude as well as latency variability in single-trial responses. Using this model, we derive the differentially variable component analysis (dVCA) algorithm, and demonstrate how different variability patterns in neural ensemble activity can be used to separate and identify their component signals. Using simulations, we evaluate not only the ability of dVCA to characterize single-trial responses, but also its robustness to noise. We also compare the specific capabilities of dVCA with two other popular decomposition techniques: principal component analysis (PCA) and ICA. Finally, we demonstrate how to apply dVCA by analyzing neuronal ensemble responses recorded intracortically in macaque V1. Throughout this paper we emphasize the ability of the dVCA algorithm to 1) more accurately account for observed trial-to-trial variability, 2) use differential variation in the amplitudes and latencies to separate and identify sources, 3) avoid enforcing statistical independence of the sources, and 4) more accurately estimate the ongoing background activity. It also merits emphasis that, in addition to its uses in ERP and ERF source separation, the dVCA algorithm is a generally useful instrument for defining and analyzing single-trial neural responses.

## METHODS

### Modeling evoked responses

Trial-to-trial variability of evoked responses can conceivably take many forms: latency shifts, amplitude variation, and even waveshape changes. Our earlier findings have shown that the observed variability can be well described by amplitude and latency variations of stereotypic response waveshape components (Truccolo et al. 2002). Of course, waveshape variability also exists, but robust single-trial amplitude and latency estimates are nonetheless obtainable with the assumption of fixed component waveshapes. For this reason, we model the response evoked from a single neural ensemble as a stationary stereotypic waveshape that can vary in amplitude and onset latency from trial to trial. We write the response evoked in a given trial mathematically as  $\alpha s(t - \tau)$ , where the function  $s(\cdot)$  represents the waveshape of the response as a function of elapsed time  $t$ ,  $\alpha$  represents the trial-specific amplitude scaling factor of the response, and  $\tau$  represents the trial-specific onset latency shift. Furthermore, when multiple neural ensembles are engaged in the response to a stimulus, the activity of each ensemble is represented in the model as a separate waveform with a distinct trial-specific amplitude and latency shift. It is important to note that by modeling both the waveshape and its amplitude and latency, there is degeneracy in the model because an overall change in amplitude scale or latency shift can be described either by the amplitude and latency parameters or by the overall scale and shift of the waveshape. To eliminate this degeneracy, we take as a convention that the ensemble average amplitude scaling factor of the response over the recorded trials is unity,  $\langle \alpha \rangle = 1$ , and the ensemble average latency shift is zero,  $\langle \tau \rangle = 0$ .

In many experimental paradigms, investigators use several detectors positioned at different locations to measure the evoked responses of neural ensembles. The degree to which a detector records a signal evoked by a particular source depends on many factors including the position and orientation of the source relative to the detector. To describe this source–detector coupling, we introduce a coupling

matrix  $\mathbf{C}$ , where the matrix element  $C_{mn}$  describes the degree to which the  $m$ th detector detects the  $n$ th source. This coupling matrix is known as the mixing matrix in the source separation literature and as the lead-field matrix in electrophysiology.

During the course of an experiment the investigator records responses to multiple presentations of a stimulus. Each presentation is typically called a trial. For the  $r$ th recorded trial, we model the data  $x(t)$  recorded by the  $m$ th detector in component form as

$$x_{mr}(t) = \sum_{n=1}^N C_{mn} \alpha_{nr} s_n(t - \tau_{nr}) + \eta_{mr}(t) \quad (1)$$

where  $n$  indexes the  $N$  neural sources activated by stimulus presentation,  $C_{mn}$  is the coupling between the  $m$ th detector and the  $n$ th source,  $\alpha_{nr}$  is the amplitude scale of the  $n$ th source during the  $r$ th trial,  $\tau_{nr}$  is the latency shift of the  $n$ th source during the  $r$ th trial,  $s_n(\cdot)$  is the waveshape of the  $n$ th source, and  $\eta_{mr}(t)$  is the unmodeled part of the data recorded in the  $m$ th detector during the  $r$ th trial. The unmodeled part of the data record is typically a combination of the recorded background activity along with any noise in the detector. For simplicity, we assume that  $\eta_{mr}(t)$  has zero mean. Thus Eq. 1 describes the data recorded in a given detector during a given trial as a sum of the signals generated by each of the  $N$  neural sources appropriately scaled in amplitude and shifted in latency for that trial and also scaled according to the coupling between each source and that detector plus the signals that we do not yet understand or care to understand. We call Eq. 1 the *multiple component event-related potential* (mcERP) model of evoked activity.

By adopting the mcERP model of the evoked responses, we implicitly adopt a well-defined set of characteristics capable of describing a neural source. The term “component” refers to the waveshape describing the temporal activation pattern of a particular neural source. Because no information regarding the spatial locations or distributions of the neural sources has gone into the model, this model does not distinguish between two spatially distinct groups of neurons that produce the same waveshape varying identically with latency and proportionally in amplitude. However, in such a situation, examination of the estimated coupling matrix would reveal two spatially distinct sources if there are detectors positioned within the proximity of each source. The major advantage obtained by estimating the coupling matrix is that practical experience in conjunction with previously obtained physiological or anatomical data suggesting source distributions can be used to independently evaluate the solutions obtained with this technique. The disadvantage, aside from withholding information from the algorithm, is that there remain two degeneracies in the model. First, there is no specified order to the sources. Second, there is no specified scale for the coupling matrix; one could halve the coupling matrix elements while doubling the magnitude of the source waveshapes and obtain an equally valid solution. These degeneracies, which are present in every other linear source separation algorithm, including PCA and ICA, pose no difficulties to the interpretation of the results. In our implementation we have chosen to normalize the columns of the coupling matrix so that the maximum value is equal to one. However, it should be noted that this scaling degeneracy could be remedied by adopting a physical model of the source currents (Knuth and Vaughan 1998). With these conventions defined, it is a straightforward matter to map the parameter values estimated by the dVCA algorithm back to a model of the single-trial response. A single-trial component is found by applying the single-trial estimates to the relevant portion of the mcERP model above, which is simply  $\alpha_{nr} s_n(t - \tau_{nr})$ . The single-trial ERP measured by the  $m$ th detector can be found by summing the contributions from each component,  $\sum_{n=1}^N C_{mn} \alpha_{nr} s_n(t - \tau_{nr})$ , which gives a noise-free estimate of the recorded evoked response.

One last notable strength of the mcERP model is that no component waveform is required to be present in every trial. In other words, a

single-trial amplitude of zero is allowable for any component. This is important because it acknowledges the empirical fact that in a given site in the brain, the trial-to-trial response is *not* stereotypic. The general resemblance of most single-trial responses to the averaged waveform arises from the fact that in most experimental trials, similar configurations of neural elements are activated with trial-varying latency and amplitude; by the same token, the occasional extreme deviation from the “stereotypic” waveform shows that essentially different ensemble configurations can substitute on a subset of the trials. By not assuming a priori that identical neural processes are active during every trial of the experiment, one is able to investigate the possibility that different processing strategies are used by the brain during the course of the recordings.

### The dVCA algorithm

Bayesian probability theory is used to derive equations that use the recorded data to identify the maximum probability estimates of the mcERP model parameters in Eq. 1. Like PCA and ICA, each source wavelshape is modeled as a pointwise time series where each point of the time series is considered to be a model parameter. The entire set of model parameters consists of these time-series points along with the mcERP parameters described above. Equations are obtained for each model parameter and an iterative fixed-point method is used to solve them jointly. The implemented algorithm is hierarchical in the sense that components are added one by one as the data analysis

progresses. Figure 1 provides a flowchart describing the analysis stages used by the algorithm. In the text that follows, we present a basic outline of the operation of the algorithm. A more detailed description, including derivations of the equations, can be found in APPENDIX A.

The key idea is that the algorithm starts with the ensemble average ERP, which is widely accepted in the community as a meaningful description of the ERP. The average ERP can be considered to be a one-component model of the data. Our algorithm improves on this result by estimating the amplitude and latency of this component waveform in each of the single trials. With this new information in hand, an amplitude-weighted and latency-shifted average of the single-trial responses is then computed. This result is iterated an appropriate number of times, resulting in a refined estimate of a one-component model of the data. This one-component model can be subtracted from the data, resulting in residual signals—called the “residuals” for short—that represent aspects of the recordings that were not captured by the model. If the residuals contain signals with significant temporal or spatial structure, another component can be added to the model. Both components are then jointly refined, allowing us to resolve details in the wavelshapes that are not visually apparent either in the raw data or in the average ERP. This process of adding components is repeated until the data are well described. This optimization heuristic is commonly referred to as a hierarchical iterative fixed-point algorithm.

The dVCA algorithm improves on the average ERP in two ways: by extending our ability to address single-trial responses and trial-to-trial

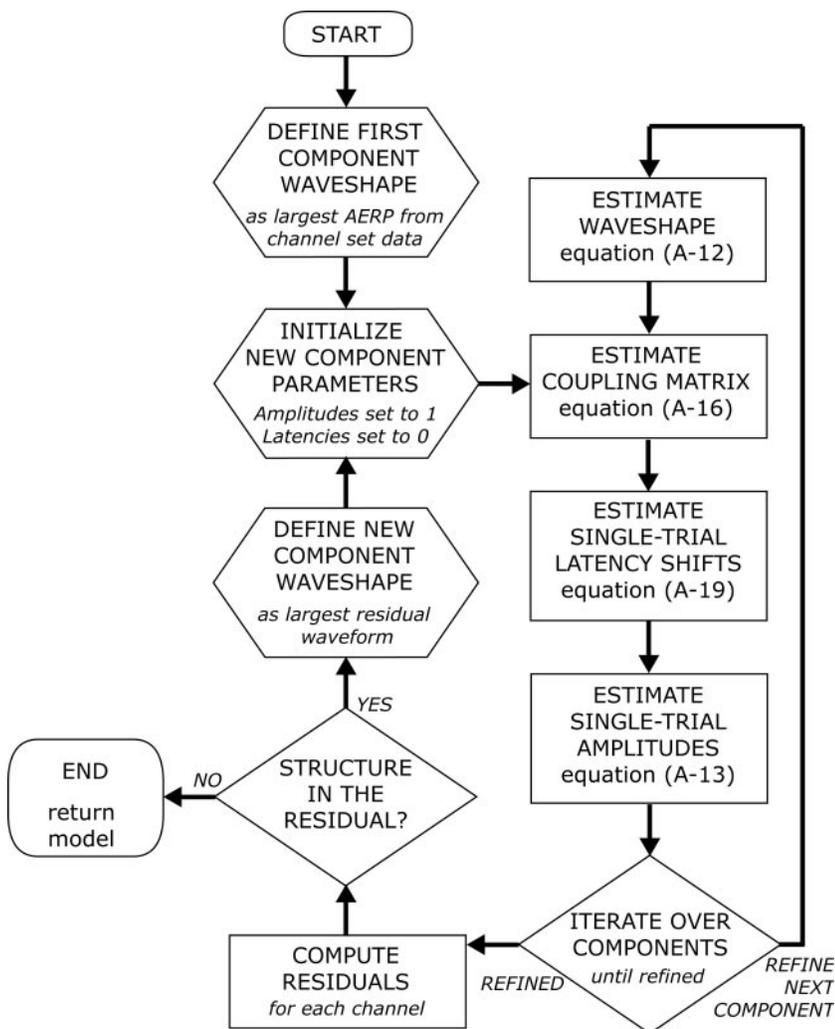


FIG. 1. Flowchart describing the differentially variable component analysis (dVCA) algorithm. This flowchart describes the implementation of the dVCA algorithm, which relies on Eqs. A12, A16, A19, and A13 described in APPENDIX A. It is based on a hierarchical model, which begins with a single-component model and adds components until the investigator chooses to stop.

variability, and by generalizing the one-component stereotypic waveshape model of the average ERP to a multiple-component stereotypic waveshape model. The trial-to-trial variability addressed by dVCA includes only single-trial amplitude and latency variation. Although retaining the concept of a stereotypic component waveshape from the average ERP may be seen as a liability, the state-of-the-art in signal processing does not allow us to consider radical multiple-component waveshape variations because that approach could easily lead to an algorithm that is overly sensitive to noise. Below, we will apply the dVCA algorithm to field potential data recorded from cortical area V1 of the macaque and demonstrate how dVCA can be used to address significant waveshape variations.

The fact that this algorithm is derived from a signal model that describes the recordings as a mixture of components implies that dVCA has something in common with techniques such as factor analysis (FA), the more familiar principal component analysis (PCA), and the more recent independent component analysis (ICA). This is in fact the case, and by reverting to a more simplistic signal model, assigning the appropriate prior probabilities and marginalizing over the source waveshape parameters one can derive, using the same methodology described in APPENDIX A, these more familiar techniques (Knuth 1997, 1999, 2005). PCA assumes that the source amplitudes are Gaussian-distributed and chooses a solution such that a single source will contribute as much of the signal variance as possible. The result is an unphysical separation, which is constrained only by this unnatural variance assumption. ICA allows for more general assumptions regarding the distribution of the source amplitudes. ICA basically assigns non-Gaussian priors to the source amplitude distributions. By assuming that the sources are independent (or equivalently that these priors are factorable) and that there are as many sources as detectors, an integration over all possible source waveshapes can be performed, resulting in a straightforward learning rule for the coupling (mixing) matrix. Mixed signals will tend to have more Gaussian amplitude distributions, so ICA strives to find a separation matrix that minimizes the Gaussianity of the results, thus optimally separating the signals.

dVCA relies on the prior information about the phenomenon of trial-to-trial variability. This information is incorporated by explicitly introducing new relevant model parameters into the signal model (Eq. 1). In the derivation, we assume that we are ignorant about the values of the model parameters and therefore assign uniform priors to the source waveshapes, single-trial amplitudes, and single-trial latencies. One can also view this as an assignment of a multidimensional uniform density, which is maximally ignorant, but also conveniently factorable, because the probability of each variable is independent of one another. However, because it is a multidimensional uniform density, any combination of values of amplitude and latency from any source is as equally probable as any other. So the algorithm is not biased toward any particular values of amplitude or latency given the amplitude or latency of the same source in another trial or another source altogether. In this sense, dVCA does not impose any condition of independence in the basic assumptions underlying the algorithm. All possibilities are equally probable.

Because dVCA makes no explicit assumptions regarding independence, this algorithm has the ability to identify distinct sources with similar waveshapes or similar amplitude and latency variability patterns. We have performed some early experiments demonstrating the ability of dVCA to identify coupled sources and sources with similar waveshapes (Shah et al. 2003). Not surprisingly, if the sources have similar waveshapes, but different trial-to-trial variability patterns, they can be separated. However, sources with identical waveshapes and identical trial-to-trial variability are inseparable. Later, we demonstrate using field potential data recorded from cortical area V1 of the macaque that coupled sources can be identified. Specifically, we identify amplitude–amplitude coupling across components, latency–

latency coupling across components, and amplitude–latency both within and across components.

## RESULTS

### Simulations

The dVCA algorithm was evaluated using synthetic data to demonstrate the utility of amplitude and latency variability in the identification of multiple evoked components and also to assess the robustness of the algorithm in the presence of noise. We simulated electric field potentials recorded from a linear-array multielectrode with 15 channels spanning the cortical laminae in V1. Specifically, we designed three synthetic ERP components (Fig. 2A) sampled at 2 kHz to approximate the neural ensemble response to diffuse red-light stimulation in macaque V1 (Givre et al. 1995; Mehta et al. 2000). Figure 2B shows the field potentials derived from the noise-free synthetic data as the detectors in the multielectrode would record them. Superimposing the field potentials from the three sources and approximating the second spatial derivative of this summed activity yields the current source density (CSD) profile shown in Fig. 2C. The value of the CSD technique (Freeman and Nicholson 1975; Nicholson and Freeman 1975; Mitzdorf 1985; Tenke et al. 1993) is evident as it both localizes the neural activity at the current sources and sinks, and eliminates volume-conducted, or far-field, activity (i.e., see component c3 below). Many of our results will be displayed using the CSD profile. We emphasize that dVCA works with the recorded field potentials only, and that displaying the CSD profile is used only to emphasize the quality of our results by retrieving accurate (in simulations) or physiologically reasonable (in real data) localizations of neural activity.

The first component, c1, represents the initial biphasic activation in lamina 4, followed by the second component, c2, representing activation in the supragranular layers. The third component, c3, represents a far-field source that volume conducts to the electrode channels and is observable in the field potential waveforms (Fig. 2B), but it is absent from the CSD plot (Fig. 2C) since the coupling between this component and the channels is nearly constant (linear with a small slope). The spatial distribution of component amplitudes across the electrode array is defined by the coupling matrix, which was chosen to simulate the expected spatial distributions of the neural ensembles (current source–sink pairs) located in lamina 4, in the supragranular layers, and at a distant site. Although this is a simplistic model, it captures features we expect to see in actual field potential recordings.

In the majority of simulations, uncorrelated Gaussian-distributed, additive noise was introduced, as in Eq. 1. The signal-to-noise ratio (SNR) is different for each component from trial to trial because the three simulated components have different single-trial amplitudes. For this reason, we specify the trial-average SNR for each component. Since the mean amplitude scale of the components is set to one, this is easily computed from the standard deviation (SD) of the detected component waveshape divided by the SD of the additive noise

$$\text{SNR}_{\text{component}} = 20 \log_{10} \frac{\sigma_{\text{component}}}{\sigma_{\text{noise}}} \quad (2)$$

where  $\sigma_{\text{component}} = 0.876, 0.174, \text{ and } 0.935$  for c1, c2, and c3, respectively. Because we are using multiple detectors, the

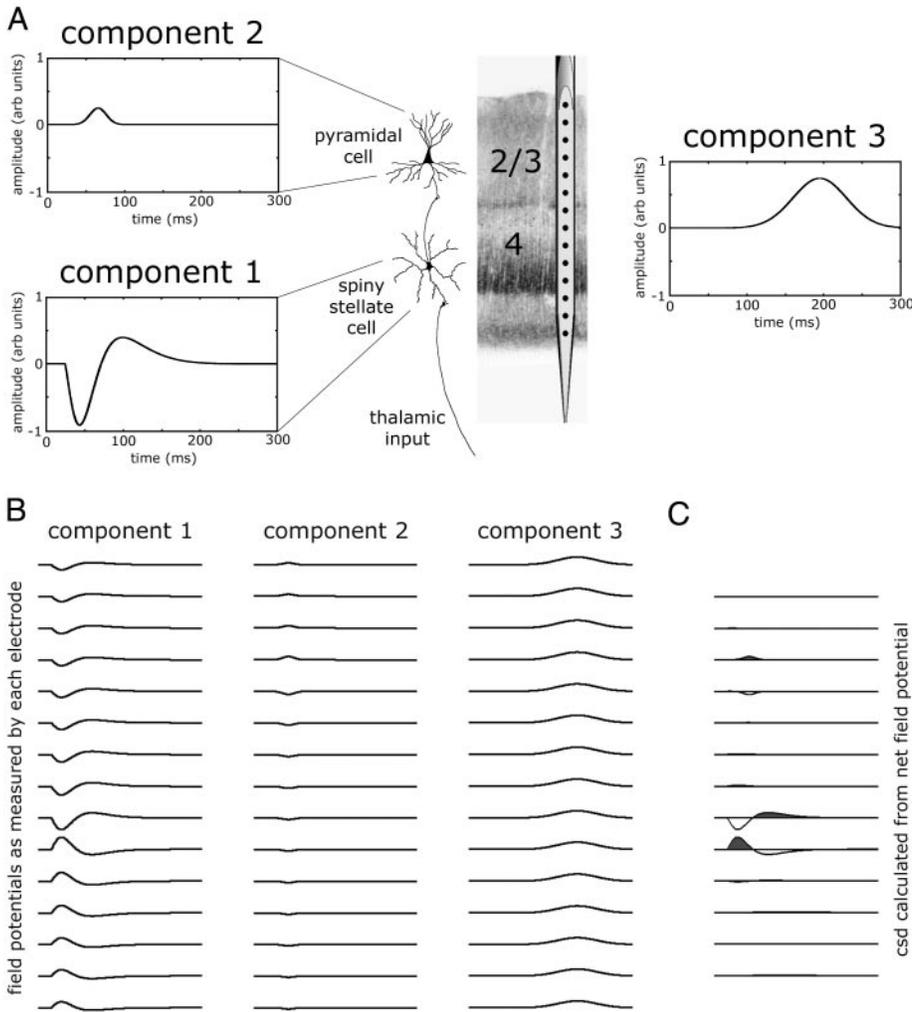


FIG. 2. Synthetic data used in the simulations. Synthetic data used herein are derived from a model of expected responses in macaque V1 when stimulated by a red-light flash. *A*: these data represent electric field potentials recorded from a 15-channel linear-array multielectrode spanning the cortical laminae in V1. Thalamic input activates the spiny stellate cells in layer IV, generating a biphasic field potential [component 1 (c1)]. Feedforward connections from the stellate cells activate the pyramidal cells in the supragranular layers [component 2 (c2)]. Component 3 (c3) models a signal generated by a far-field source that volume conducts to the multielectrode array. *B*: noise-free synthetic field potentials generated by these 3 components are recorded differently by each electrode in the multielectrode array. Notice that the polarity and amplitude of the recordings depend on the physical positions of the current sources and sinks in the cortical laminae (see *C*). *C*: current source density (CSD) of the summed field potentials in *B* is computed using an approximation of their second spatial derivative. CSD focuses the activity at the location of the current sinks (negative) and sources (positive) relative to the detector positions. This technique clarifies the laminae in which the field potentials are generated. Notice that far-field sources do not appear in the CSD.

values of  $\sigma_{component}$  were found by multiplying the SD of the waveshape  $\sigma_{wave}$  by the L2 norm of the coupling coefficients for that component. For the  $n$ th component, this is

$$\sigma_{component_n} = \sigma_{wave_n} \sqrt{\sum_{m=1}^M C_n^2} \quad (3)$$

We will often also provide the SNR ratio, which is a ratio of the component and noise SDs.

The performance of the dVCA algorithm was evaluated in several ways. First, the ability of dVCA to separate the three mixed signals was measured using a unit-normalized Amari error (Amari et al. 1996), which describes the degree of component mixing (details can be found in APPENDIX B, Eq. B5). A completely separated solution gives an Amari error of zero, whereas the worst case mixture gives an Amari error of one. For three signals of equal variance, an Amari error of 0.065 corresponds to 10% mixing.

Second, the ability of dVCA to estimate each component waveshape was evaluated by calculating the fractional root-mean-square (RMS) error of the estimated waveshape as compared with the correct waveshape. Note that before this comparison could be made the estimated components  $\hat{s}(t)$  needed to be rescaled and permuted because of the indeterminacy described by Eq. B1. For the  $j$ th component

$$E_{wave}^j = \frac{\sqrt{\sum_{t=1}^T [s_j(t) - \hat{s}_j(t)]^2}}{\sqrt{\sum_{t=1}^T [s_j(t)]^2}} \quad (4)$$

Finally, the accuracies of the amplitude and onset latency estimates for the  $j$ th component were evaluated by computing the deviation from the correct values for the entire set of single-trial estimates using

$$E_{amp}^j = \alpha_{jr} - \hat{\alpha}_{jr} \quad (5)$$

and

$$E_{lat}^j = \tau_{jr} - \hat{\tau}_{jr} \quad (6)$$

and then computing the range of values contained within the 68th percentile. This not only provides a measure of accuracy, but also allows us to compare the advantage of employing dVCA over the standard technique of averaging. The key idea is to note that if the trial-to-trial variability can be summarized by the SD ( $\sigma$ ), then the error one obtains by working with the average ERP necessarily must be greater than or equal to  $\sigma$  because it does not account for trial-to-trial variability. Thus errors in the dVCA estimates below the SDs of the variability

represent a significant improvement over the standard assumption that amplitudes and latencies do not vary.

### Effects of amplitude variability

We first demonstrate that amplitude variability aids in the separation process. Eleven experiments using synthetic data, each consisting of 50 simulated trials from the source component configuration described above, were performed where the degree of variability of the single-trial amplitudes was controlled. To generate the synthetic data, 50 single-trial amplitudes for each component were randomly sampled from a log-normal distribution with a sample mean  $\mu_{amp} = 1.0$  and sample SDs of  $\sigma_{amp} = \{0, 0.0625, 0.125, 0.1875, 0.219, 0.25, 0.375, 0.5, 0.625, 0.75, 1.0\}$  for each of the 11 sets of synthetic data. Because latency variability was not simulated in this set of experiments, the single-trial latency parameters were set to zero ( $\tau_{jr} = 0$ ). Given the component waveshapes, the coupling matrix, and the single-trial amplitudes and latencies, Eq. 1 was used to generate the synthetic data. The SD of  $\sigma_{noise} = 0.217$  was used for the noise, resulting in SNRs of 12.1, -1.9, and 12.7 dB (4.04, 0.80, and 4.31 in terms of SD ratios) for c1, c2, and c3, respectively. The dVCA algorithm was used to estimate the coupling matrix, component waveshapes, and single-trial amplitudes and latencies from the synthetic data. Because the true parameters were known, the performance of the algorithm was evaluated as previously described.

Without amplitude variability, dVCA was unable to separate the components as demonstrated by the high Amari error of 0.219 (see Fig. 3A). A small degree of amplitude variability,  $\sigma_{amp} = 0.25$ , renders the problem solvable as demonstrated in Fig. 3A, where the Amari error drops below 0.05 corresponding to a fractional RMS waveshape error of <10% (SNR  $\approx$  23 dB) for three equal variance components, and remains relatively constant with increasing variability hovering about an average of 0.028. Figure 3B shows the dramatic improvement in the waveshape estimation as quantified by the RMS errors, which rapidly drop to levels commensurate with the SNRs of the individual components. The effect of amplitude variability on the single-trial amplitude estimates (not shown) remained relatively constant for  $\sigma_{amp} \geq 0.25$  with mean SDs of the errors in the single-trial amplitude estimates of 0.014, 0.076, and 0.010 and for c1, c2, and c3 respectively. This is well below the SD of the amplitude variability. It should be noted that even though the true onset latencies were set to zero in these

simulations, dVCA still estimates these quantities. The errors of the onset latency estimates (not shown) also remained relatively constant with respective mean SDs of 0.417, 2.059, and 1.000 ms.

Finally, experiments were repeated with normally distributed single-trial amplitudes rather than a log-normal distribution. Negative amplitude values were discarded and the sample set was controlled to maintain a sample mean of  $\mu_{amp} = 1.0$  and the specified sample variance. The two analyses gave similar quantitative and qualitative results.

### Effects of latency variability

Next we demonstrate that latency variability also aids in the separation process. Eight experiments using synthetic data, each consisting of 50 synthetic trials from the source component configuration described above, were performed where the variability of the single-trial latencies was controlled. In these experiments there was no amplitude variability ( $\alpha_{jr} = 1$ ). To generate the simulated data, 50 single-trial latencies were randomly drawn from a Gaussian distribution with sample mean  $\mu_{lat} = 0$  and sample SDs of  $\sigma_{lat} = \{0, 1.25, 2.5, 3.75, 5.0, 6.25, 7.5, 10.0\}$  ms for each of the eight simulations. The same noise variance was used as in the amplitude variability case.

Amari error was found to decrease with increasing latency variability (Fig. 3C), dropping to <0.05 with  $\sigma_{lat} \geq 7.5$  ms. Component waveshape estimation was also found to improve with increasing onset latency variability, although the effect is not nearly so dramatic as in the amplitude variability case. Moreover, with the onset latency variability ranges considered, the accuracy of the algorithm's estimates never attained the levels found with amplitude variability. This difference between the effects of amplitude and latency variability was not attributable to the sampling distributions because we controlled for both the sample mean and variance. Instead, the most likely reason for this effect is the fact that in terms of signal amplitude, amplitude variability is a first-order effect, whereas latency variability, when written as a Taylor expansion, is a second-order effect dependent on the first derivative of the waveshape. Although the amplitudes were all set to one in the simulated data, dVCA still estimates these quantities. Again, these amplitude estimate errors were low for  $\sigma_{lat} \geq 7.5$  ms with the mean SD of the errors in the estimates equal to 0.017, 0.077, and 0.011 for c1, c2, and c3, respectively, which is on the order of the errors seen in the amplitude variability trials.

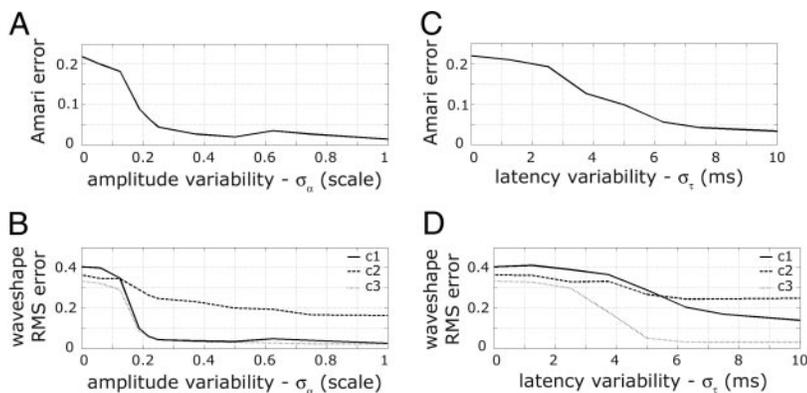


FIG. 3. Amplitude and latency variability aid component identification. A: Amari error, which measures the degree of signal separation, decreases with increasing amplitude variability where  $\sigma_{amp} \geq 0.25$  is sufficient for signal separation. B: waveshape root-mean-square (RMS) errors also indicate the importance of amplitude variability. C: Amari error decreases with increasing latency variability, where  $\sigma_{lat} \geq 7.5$  ms is necessary to achieve effective separation. D: increasing latency variability also improves the estimate of the component waveshapes, but not as dramatically as amplitude variability.

Onset latency estimate errors also remained relatively constant with SDs of 0.250, 2.250, and 1.142 ms, respectively.

#### Sensitivity to additive white Gaussian noise

In this first noise study we examined the robustness of dVCA in the face of additive white Gaussian noise by simulating 12 different noise levels. Each simulation relied on 50 trials of synthetic data where variability in both the amplitudes and latencies of the components was modeled. Variability ranges were in accordance with those observed in our previous investigations (Truccolo et al. 2001). Single-trial amplitudes were drawn from a log-normal distribution with sample mean  $\mu_{amp} = 1.0$  and sample SD  $\sigma_{amp} = 1.0$ . Single-trial latencies were drawn from a normal distribution with sample mean  $\mu_{lat} = 0$  and sample SD  $\sigma_{lat} = 10.0$  ms. The synthetic signal from each electrode (detector) was contaminated with a unique white Gaussian noise waveform, which is specified by  $\eta_{mr}(t)$  in Eq. 1.

The component-specific SNRs and resulting Amari errors, listed in Table 1 and shown in Fig. 4A, indicate a relatively smooth increase in Amari error with decrease in c1 SNR. However, a significant jump in error occurs as the SNR of c2 passes below  $-17$  dB (note that this corresponds to a c1 SNR =  $-3$  dB; see figure caption), suggesting that an SNR level of about  $-17$  dB may denote a transition from a useful data set to a prohibitively noisy one. These results can be compared with the waveshape fractional RMS errors, which grow exponentially with decreasing SNR (Fig. 4B). The errors reach 1.0, signifying that the deviations in the estimates are on the scale of the waveshape itself, at about  $-15$  to  $-17$  dB for the localized components and about  $-17$  dB for the far-field component. Figure 4E shows waveshape results for four different noise levels, providing a better idea of the quality of the separation under these conditions. Most noteworthy is the fact that the majority of the waveshape error arises from the high-frequency contamination of the Gaussian noise rather than mixing of the components. Much of this could easily be improved by incorporating a prior probability describing the expectation that components should be slowly varying with respect to typical sampling rates, which would effectively low-pass filter the results. As we will demonstrate, explicitly

filtering a real data set can remove important signals, and alter others (Bogacz et al. 2002; Mocks et al. 1986). Regardless, some distortion in the positive phase of c1 can be seen at  $SNR_1 = -3.0$  dB, which is much more easily noticeable at  $SNR_1 = -15$  dB, where only c1 and c3 remain detectable. The improved accuracy in the estimation of the far-field component c3 (illustrated in Fig. 4B) is most likely explained by the fact that each electrode in the array provides information about the far-field source, whereas for local sources some detectors have small signals.

Next we examine the quality of the single-trial amplitude estimates. First, it is important to realize that, if we simply take the trial average as an estimate of the resulting evoked response, we would be assuming that during each trial the response had the same average amplitude. Because this is not the case, the errors in our simplistic amplitude estimate would be on the order of the physiologic variability of response amplitudes. If the performance of dVCA is such that its amplitude estimation errors are less than the trial-to-trial amplitude variability, then dVCA is necessarily performing better than the trial average. Put another way, the performance of dVCA exceeds that of signal averaging when the estimation errors are less than the degree of variability in the original single-trial amplitudes. Figure 4C shows the relationship between the errors of the estimates and the range of variability of the amplitudes. The horizontal black lines in Fig. 4C represent the degree of variability of the component amplitude as 1 SD of the single-trial amplitudes about their mean. The solid blue curve represents 1 SD of the dVCA single-trial amplitude estimate errors. When the blue curves are contained within the black lines, dVCA is outperforming standard averaging. Again, the far-field estimates (c3) were more accurate than those of the local sources (c1 and c2). Ninety-five percent of the amplitude estimates were within the degree of variability down to SNRs between  $-10$  and  $-15$  dB with 68% of the estimates within the range well down to  $-15$  to  $-18$  dB, which is consistent with the performance indicated by both the Amari error and the waveshape RMS error. To demonstrate the quality of the amplitude estimates, Fig. 4F shows scatterplots of the true amplitude scales versus the estimates for c1 at the same four SNR levels as in Fig. 4E. Note that different values for the single-trial amplitudes were used in each simulation. Correct estimates will lie on the diagonal, whereas incorrect estimates are off diagonal. Although the errors in the values of the amplitude estimates become unacceptable around  $-15$  dB the pattern still exhibits a strong correlation ( $r = 0.948$ ).

Figure 4D shows the behavior of the onset latency estimates, which are less well estimated than the amplitudes. In addition, the degradation of the quality of the far-field estimates was not noticeably different from those of the local sources with 95% of the estimates being within the range of variability down to an SNR of 3 to  $-1$  dB, and 68% of the estimates within the range of variability down to  $-15$  to  $-18$  dB. Figure 4G shows scatterplots of the true onset latencies versus the estimates for c1. The diagonal pattern, which indicates a high level of predictability, is almost lost by  $-15$  dB ( $r = 0.392$ ). Note that as a result of the difference in variability levels of the amplitudes and onset latencies, the distribution of points in Fig. 4F should not be directly compared with those in Fig. 4G because they are effectively at different magnifications.

TABLE 1. Effect of additive white Gaussian noise on separation

Case	Noise SD	Component SNR, dB			Amari Error
		1	2	3	
1	0.155	15.0	1.0	15.6	0.004
2	0.310	9.0	-5.0	9.6	0.005
3	0.437	6.0	-8.0	6.6	0.013
4	0.618	3.0	-11.0	3.6	0.011
5	0.734	1.5	-12.5	2.1	0.014
6	0.873	0	-14.0	0.6	0.026
7	1.037	-1.5	-15.5	-0.9	0.017
8	1.233	-3.0	-17.0	-2.4	0.050
9	1.742	-6.0	-20.0	-5.4	0.108
10	2.460	-9.0	-23.0	-8.4	0.100
11	4.909	-15.0	-29.0	-14.4	0.198
12	9.794	-21.0	-35.0	-20.4	0.421

Component SNRs are calculated using the component SDs of 0.290, 0.059, and 0.262 and coupling L2 norms of 3.02, 2.98, and 3.57 for components c1, c2, and c3, respectively.

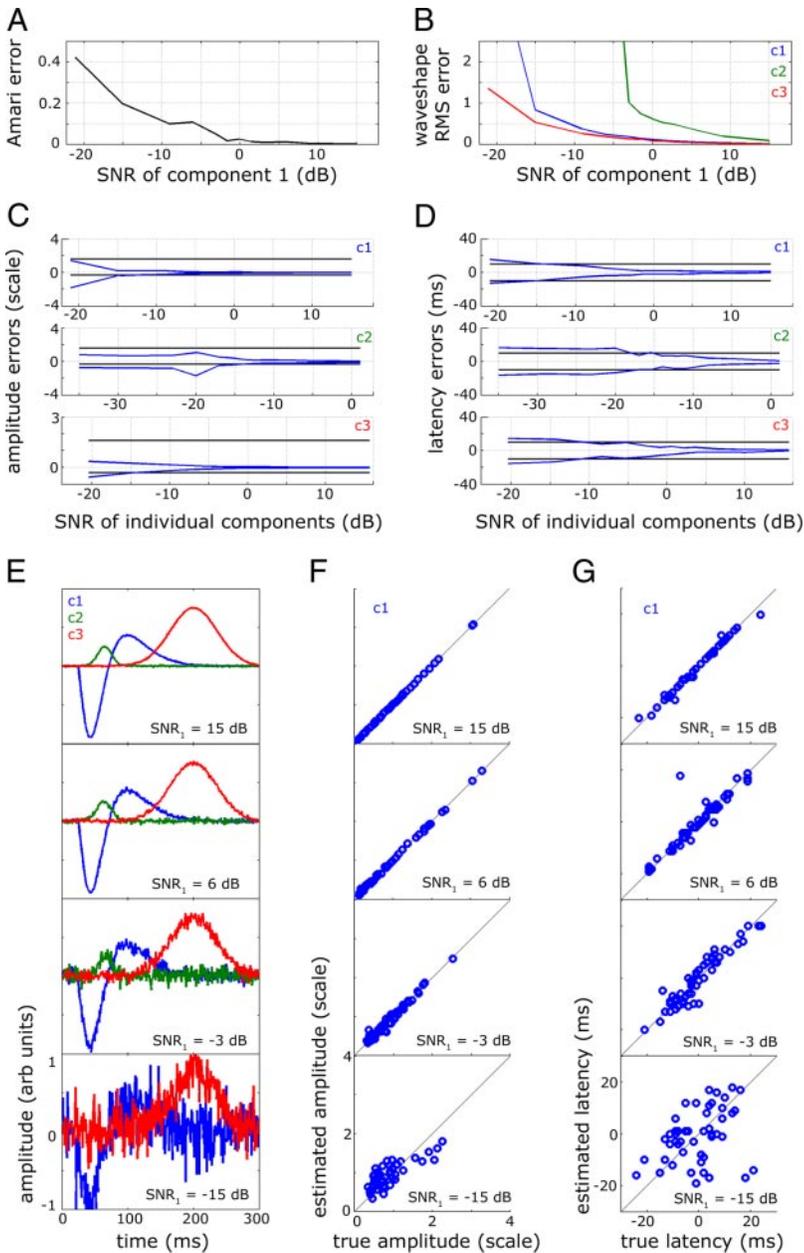


FIG. 4. dVCA algorithm is robust to noise. In this figure we study the behavior of dVCA with varying signal-to-noise ratio (SNR). Because of their different magnitudes each component has a different SNR, and for simplicity the plots are made with respect to the SNR of component 1, labeled  $\text{SNR}_1$ . Recall that  $\text{SNR}_2 \approx \text{SNR}_1 - 14$  dB and that  $\text{SNR}_3 \approx \text{SNR}_1 - 0.6$  dB. **A:** Amari error decreases with increasing SNR (see text). **B:** quality of the waveshape estimates improves with increasing SNR. Note that the graph is drawn with respect to c1 SNR. Waveshape RMS error reaches 1.0, signifying that the deviations in the estimates are on the order of the waveshape itself, at about  $-15$  to  $-17$  dB for the localized components and about  $-17$  dB for the far-field component. **C:** single-trial amplitude estimates are also robustly recovered. Horizontal black lines denote the level of variability in the single-trial quantities. Blue curves indicate 1 SD of the single-trial amplitude estimate errors. When the blue curves are within the black lines, dVCA is performing better than the standard practice of averaging (see text). High-quality estimates are achieved down to  $-15$  to  $-18$  dB. **D:** single-trial latency estimates are slightly less well estimated than amplitudes, although high-quality estimates are again possible down to  $-15$  to  $-18$  dB. **E:** these plots demonstrate the quality of the waveshape estimates for 4 SNR levels. At a c1 SNR of  $-15$  dB, the smallest component, c2, was unable to be extracted from the data. **F:** scatterplots of the true c1 single-trial amplitudes vs. the estimated c1 single-trial amplitudes demonstrate the quality of amplitude estimates, and show that useful information can be obtained down to an SNR of  $-15$  dB. **G:** scatterplots of the true c1 single-trial latencies vs. the estimated c1 single-trial latencies demonstrate that latencies are less well estimated than amplitudes. Latencies can be estimated down to SNR levels well below  $-3$  dB, but become inaccurate by  $-15$  dB.

#### Sensitivity to correlated far-field noise

In an effort to more accurately simulate the conditions that may be experienced during a real experiment, we designed a second noise study to determine the effect of ongoing correlated far-field activity on dVCA performance. The ongoing far-field activity (noise) was modeled as a time series with a  $1/f$  spectrum so that correlations would exist at all timescales. To simulate its far-field nature the ongoing activity was coupled identically to each detector, which can also be viewed as a spatial correlation across channels.

Again, 12 levels of the ongoing noise amplitude were tested; however, in this study Gaussian noise was not added to the individual electrodes, only the ongoing correlated far-field noise was used. Noise amplitude was computed from the sample variance of the time-series noise signal. The specific noise levels, SNRs of the three components and resulting Amari errors are listed in Table 2. Figure 5A shows the Amari

error smoothly decreasing with increasing SNR. The Amari error reaches the same level of error as in the Gaussian noise case with 20 dB less noise, amounting to large errors around a c1 SNR of 3 to 8 dB. It is apparent that this type of noise more severely limits the ability of dVCA to separate signals. Figure 5B shows the component waveshape errors blowing up around 1 dB, with the effect on the far-field component c3 (red) being understandably catastrophic because both c3 and the noise are correlated across detectors. The amplitude errors reach unacceptable levels around 0 to 3 dB (Fig. 5C). However, most interesting are errors of the onset latencies (Fig. 5D). The errors for the local sources reach high levels (68% of estimates being within the range of variability) between 0 and 1 dB, whereas the errors for the far-field component c3 are large across the entire SNR range considered with 95% of the estimates never being within the range of variability. This is because the correlated far-field noise is severely interfering

TABLE 2. Effect of additive l/f far-field noise on separation

Case	Noise SD	Component SNR, dB			Amari Error
		1	2	3	
1	0.036	27.8	13.8	28.4	0.015
2	0.048	25.3	11.3	25.9	0.013
3	0.063	22.8	8.8	23.4	0.017
4	0.113	17.8	3.8	18.4	0.008
5	0.150	15.3	1.3	15.9	0.035
6	0.201	12.8	-1.2	13.4	0.068
7	0.268	10.3	-3.7	10.9	0.113
8	0.357	7.8	-6.2	8.4	0.143
9	0.476	5.3	-8.7	5.9	0.159
10	0.634	2.8	-11.2	3.4	0.191
11	0.846	0.3	-13.7	0.9	0.359
12	1.128	-2.2	-16.2	-1.6	0.365

Component SNRs are calculated using the component SDs of 0.290, 0.059, and 0.262 and coupling L2 norms of 3.02, 2.98, and 3.57 for components c1, c2, and c3, respectively.

with the ability to identify the far-field component in any given trial.

#### PCA, ICA, and dVCA comparison

Neuroscientists have several available techniques to analyze evoked responses, although each of these techniques is based on different assumptions and signal models. Here we briefly evaluate two popular techniques, PCA and ICA, and compare them to dVCA using our synthetic data, which as described above, is based on experimental findings regarding both the componentry and their laminar distribution (Givre et al. 1995; Mehta et al. 2000) and observed trial-to-trial variability (Trucolo et al. 2002). ICA was performed by applying the Infomax ICA algorithm in the EEGLAB toolbox (Delorme and Makeig 2004). The runica.m algorithm (Makeig et al. 1996) was used both with and without the “extend” option (extended ICA), which allows ICA to model sub-Gaussian sources (Lee et al. 1999). Because extended ICA consistently gave better results in the single-trial data sets exhibiting trial-to-trial variability, we will present those results here and in Fig. 6. PCA was performed using the EEGLAB routine runpca.m (Colin Humphries, CNL/Salk Institute, August 1997), which com-

putes the singular value decomposition (SVD) of the data and identifies the resulting eigenvectors as the estimates of the source waveshapes. Because each data set has 15 channels, both PCA and ICA estimate 15 source waveshapes. Because it is known that the synthetic data are composed of three sources, the three estimated source waveshapes with highest correlation to the three original sources were chosen for analysis, whereas the remaining estimated sources were ignored. This avoids the difficult issues of estimating the number of sources in both PCA and ICA, and thus gives them a slight advantage.

Two synthetic data sets were used to perform the comparison. The first data set exhibited neither amplitude nor latency variability and was contaminated by low-amplitude additive Gaussian noise (component SNRs of 15, 1.0, and 15.6 dB). The second data set, examined earlier in the section on *Sensitivity to additive white Gaussian noise*, exhibited log-normally distributed amplitude variability with sample mean  $\mu_{amp} = 1.0$  and sample SD  $\sigma_{amp} = 1.0$ , normally distributed latency variability with sample mean  $\mu_{lat} = 0$  and sample SD  $\sigma_{lat} = 10.0$  ms, and possessed additive Gaussian noise identical in SNR level to that of the no variability case.

These two data sets were analyzed using PCA and ICA in two different ways: first by averaging the epochs and analyzing the average response (average analysis), and second by treating the data as a long string of concatenated single-trial responses (single-trial analysis). Figure 6A shows the source waveshape results for PCA, extended ICA, and dVCA in the single-trial analysis cases. Note that only the dVCA analysis of the variable responses results in accurate identification of the underlying source waveshapes. The Amari errors in Fig. 6B quantify the degree to which the components were correctly identified in the four cases. In this example, the performance of the original ICA algorithm was comparable to PCA. Because dVCA can analyze only single-trial responses, the dVCA analysis consisted of analyzing the single-trial cases, resulting in only two bars in the bar graph. Note that dVCA performs poorly when there is no trial-to-trial variability. However, the presence of trial-to-trial variability dramatically improves the performance of dVCA, enabling it to surpass both PCA and extended ICA in separation quality.

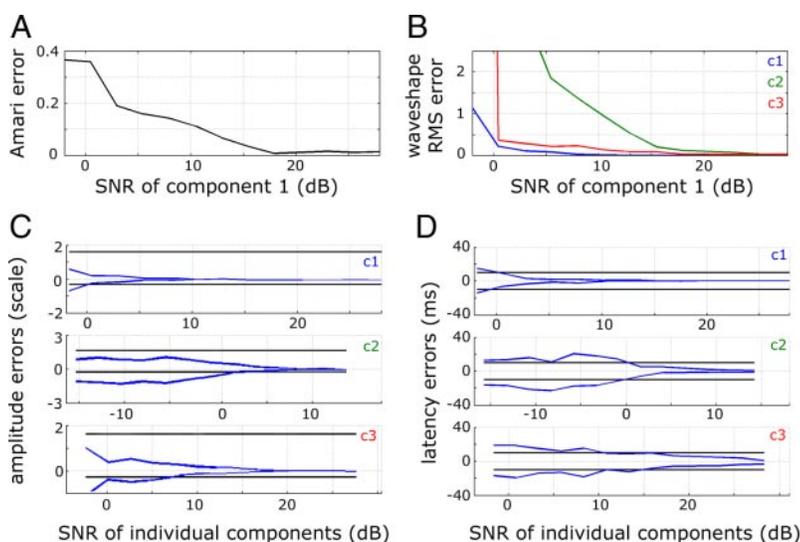


FIG. 5. dVCA algorithm is also robust to correlated far-field noise, although to a lesser degree than independent Gaussian noise. *A*: Amari error decreases with increasing SNR and reaches the same level of error as in the Gaussian case with 20 dB less noise. *B*: component waveshapes blow up around 0 dB, with the far-field component, c3, being more dramatically affected. *C*: amplitude errors reach unacceptable levels from 0 to 3 dB, whereas the far-field component amplitudes can be estimated only down to around 8 dB. *D*: latency estimates become unacceptable around 0 dB for the local sources, whereas the far-field latencies can be estimated only at levels above 10 dB.

TABLE 3. Symbols used throughout the text

Numbers and Indices	
$N$	Number of components
$M$	Number of recording channels
$R$	Number of trials
$T$	Number of time points recorded
$n, j$	Index for number of components $N$
$m, i$	Index for number of channels $M$
$r, p$	Index for number of trials $R$
$t, q$	Index for number of time points $T$
Model parameters (Boldface denotes matrix form)	
$\alpha_{nr}$ , $\alpha$	Amplitude of $n$ th component during $r$ th trial
$C_{mv}$ , $\mathbf{C}$	Coupling (mixing matrix) between channel $m$ and source $n$
$s_n(t)$ , $\mathbf{s}(t)$	Magnitude of source $n$ at time $t$
$\tau_{nr}$ , $\tau$	Onset latency of $n$ th component during $r$ th trial
$\eta_{mr}(t)$ , $\eta(t)$	Noise recorded in $m$ th channel at time $t$ during $r$ th trial
$x_{mr}(t)$ , $\mathbf{x}(t)$	Data recorded in $m$ th channel at time $t$ during $r$ th trial
Error analysis and variability	
$\mu$	Mean value (quantity indicated by subscript)
$\sigma$	Standard deviation (quantity indicated by subscript)
SNR	Signal-to-noise ratio in units of decibels (dB)
$E_{Amari}$	Amari error between mixing matrix and estimated mixing matrix
$E_{wave}$	Fractional root mean square error of the $j$ th source waveshape
$E_{amp}$	Difference between estimated and correct single-trial amplitude
$E_{lat}$	Difference between estimated and correct single-trial latency
Miscellaneous	
$I$	Prior information
$Q$	Sum of squares difference between model and data
$\mathbf{M}$	Matrix transformation (such as a mixing matrix)
$\Sigma$	Diagonal scaling matrix
$\pi$	Permutation matrix

Note that estimated parameters are indicated with "hats" (carets), as in  $\hat{s}_n(t)$ .

Finally, we compared the ability of extended ICA and dVCA to identify source waveshapes in the face of additive white Gaussian noise. These data sets are identical to those analyzed in the earlier section on *Sensitivity to additive white Gaussian noise*. Note also that the data set with the smallest noise level (SNR<sub>1</sub> = 15 dB) is the same data set that was used in the variable single-trial case considered above. Figure 6C demonstrates that dVCA consistently outperforms extended ICA with respect to robustness to noise. It should be noted that the extended ICA algorithm outperformed the original ICA algorithm in all but two cases. At SNR<sub>1</sub> = -9 dB, the Amari error of extended ICA was 0.240, whereas the Amari error of ICA was 0.177 compared with the Amari error of dVCA, which was 0.100. In the second case, the Amari error of extended ICA and ICA differed by only 0.001.

#### Application to cortical field potential data

To further demonstrate the utility of dVCA, we applied the algorithm to field potentials recorded from primary visual cortex of a male macaque monkey during the cuing period of an earlier selective attention study (Mehta et al. 2000) for which all animal care and use procedures were approved by the Institutional Animal Care and Use Committee, and were in accordance with National Institutes of Health publication no. 86-23 (rev. 1087). During data collection, the subject was required to discriminate between standard and target visual stimuli for a juice reward. The standard visual stimulus was a 10- $\mu$ s red-light flash presented through a diffusing plate subtending a 20° visual angle, centered on the point of visual fixation, while the target was presented similarly, although

differing slightly in intensity. Stimuli were presented at irregular interstimulus intervals (ISIs) (minimum of 350 ms, average of 2 stimuli/s). Intracortical field potential profiles were recorded using a linear-array multielectrode with 14 contacts (or channels), equally spaced at 150  $\mu$ m, inserted into V1 and positioned so that the channels spanned all six laminae (see Fig. 2A, *middle*). The continuous field potential record from all channels, incorporating the stimulus tags, was sampled at 2 kHz and recorded on a PC-based data acquisition system (Neuroscan, El Paso, TX).

The signals examined with dVCA were recorded during 171 trial presentations of the standard visual stimulus and were epoched from 0 to 300 ms (0 ms indicates stimulus presentation). No other preprocessing of the data was performed. The average field potential signals in each electrode contact were calculated and used to determine the current source density (CSD) profile of these data (Fig. 7A). The CSD is an estimate of the second spatial derivative of the electric potential, which is proportional to the vector divergence of the current field. This profile was approximated using a three-point, second-order difference of the field potentials (Nicholson and Freeman 1975; Schroeder et al. 1995). Because the CSD represents the divergence of the electric current, it indexes the local regions of transmembrane current flow and eliminates volume-conducted activity generated at a distant site, which often contaminates field potential recordings. This makes the CSD a useful tool for studying the electric signals detected by an electrode array. Consistent with earlier studies reviewed by Schroeder et al. (1995, 2001), the visually evoked CSD profile sampled from V1 during this experimental session shows that the earliest activation occurs in the granular subdivisions of Layer

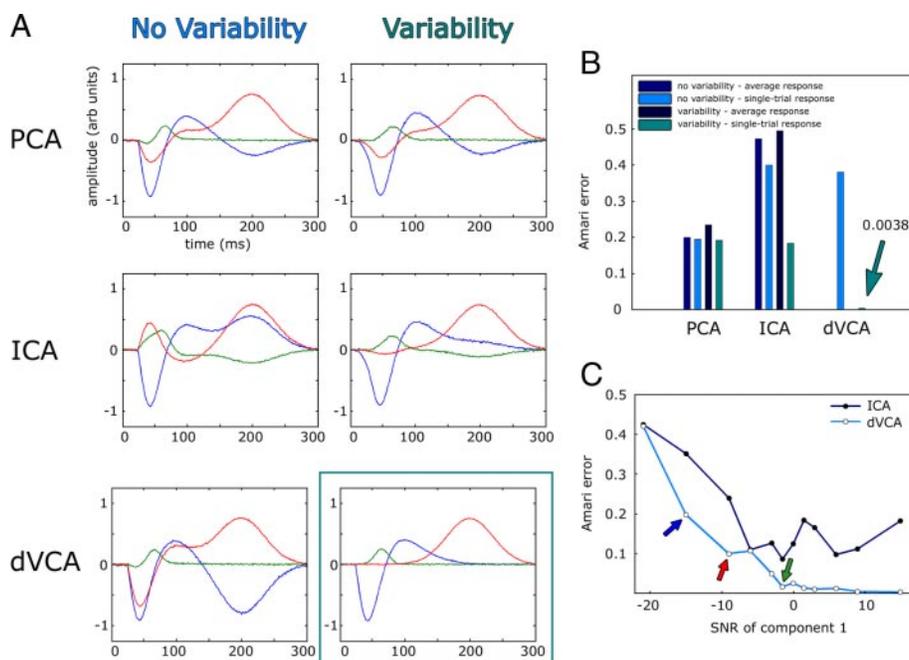


FIG. 6. Comparison between principal component analysis (PCA), independent component analysis (ICA), and dVCA. Two synthetic data sets were generated, one with trial-to-trial variability ( $\sigma_{amp} = 1.0$ ,  $\sigma_{lat} = 10.0$  ms) and one without (see text for more details). Both data sets have the same SNR level ( $SNR_1 = 15$  dB). **A:** PCA, ICA, and dVCA were used to analyze these data as one long time series of multiple single-trial epochs. Note that PCA and ICA were not able to accurately separate these waveshapes, whereas dVCA, which fails in the case where there is no variability, succeeds in the variable case (*boxed graph*). **B:** accuracy of these analyses is quantified by considering the Amari error. For PCA and ICA, the analyses were also performed for averaged epochs. Both PCA and ICA results are improved when single-trial data are used over analysis of averaged responses. dVCA, which can be used only on single-trial data, dramatically outperforms both PCA and ICA when the responses exhibit trial-to-trial variability. Resulting waveshapes corresponding to the single-trial results are shown in **A**. **C:** ICA and dVCA are compared using noisy data. Data set with the smallest noise level ( $SNR_1 = 15$  dB) is the same data set that was used in the variable, single-trial case considered above (compare with Amari errors in **B**). dVCA consistently outperforms ICA with respect to robustness to noise. Each colored arrow indicates the SNR levels at which dVCA can no longer detect its associated component. For example, the blue arrow indicates that below an SNR of  $-15$  dB, component 1 (blue) can no longer be detected. Likewise, the red arrow indicates the SNR below which component 2 was no longer detectable and the green arrow indicates the point below which component 3 was no longer detectable. To the *left* of each of these points our Amari error necessarily increases; this is because we can no longer detect one of the components and have an incomplete picture of the componentry. With high noise (*far to the left*) we have lost all ability to detect components, and dVCA's performance is now similar to that of ICA.

4, which constitute the major target of thalamic afferents. A second focus of activity localizes to the supragranular (laminae 2 and 3) layers and is thought to index a feedforward activation of pyramidal cells by interneurons in the granular layers (Schroeder et al. 1990, 1991). Note that because we are using a three-point, second-order difference to approximate the CSD, the top waveform in Fig. 7A is associated with electrode channel 2 rather than channel 1.

Although it is an attractive idea to apply dVCA as an automated computer program that analyzes the data and produces an answer (as we did with simulated data), it is often more productive to apply dVCA in a more facultative approach (Shah et al. 2004). To illustrate and stress the importance of this point, instead of following the flowchart exactly as in Fig. 1, we first applied dVCA to the single-trial field potential signals to extract a single component. As shown in Fig. 7B, the CSD profile of this component depicts a biphasic response localized in the granular layer, suggesting that it represents the initial response to the thalamic inputs. Not surprisingly, this waveshape and its distribution are very similar to the CSD profile of the ensemble averaged ERP of Fig. 7A because the initial approximation to the first component derives from the ensemble average. However, note that it excludes much of the minor variations not associated with this main activation pattern. In addition to providing information about the component waveshape and its spatial distribution, dVCA provides infor-

mation about the amplitude and latency shift of the component in each trial. Figure 7C shows the distribution of single-trial amplitude scales for this component. In accordance with the dVCA normalization condition, the mean of the sample equals one ( $\mu_1 = 1$ ). In this case, the distribution shows that there is very little single-trial amplitude variability ( $\sigma_1 = 0.1044$ ), which is on the order of 10% variation. In contrast, the distribution of single-trial latency shifts (Fig. 7D) shows something quite different and unexpected—it is bimodal with an early latency peak at  $-4.625$  ms and a late latency peak at  $2.125$  ms with a difference of  $6.75$  ms. The ratio of late to early responses is 2.29:1 (119:52), as defined by a  $-1.25$ -ms cutoff between the two modes.

We performed several different analyses to confirm the existence of the bimodal latency distribution displayed in Fig. 7D. First, Fig. 7E displays a colorized plot of the actual, single-trial field potentials (FPs) recorded from electrode channel 10, which is located in the granular layer. Time runs along the horizontal axis, chronological trial number runs along the vertical axis, and color indicates the amplitude of the FP. The red dashes on the *left side* of the plot indicate trials for which dVCA determined the latency to be early. It is easily seen that these highlighted trials have a response onset before the late trials, confirming that the bimodal latency distribution of the dVCA estimation concurs with observations from the actual data. It is also apparent that early and late responses are

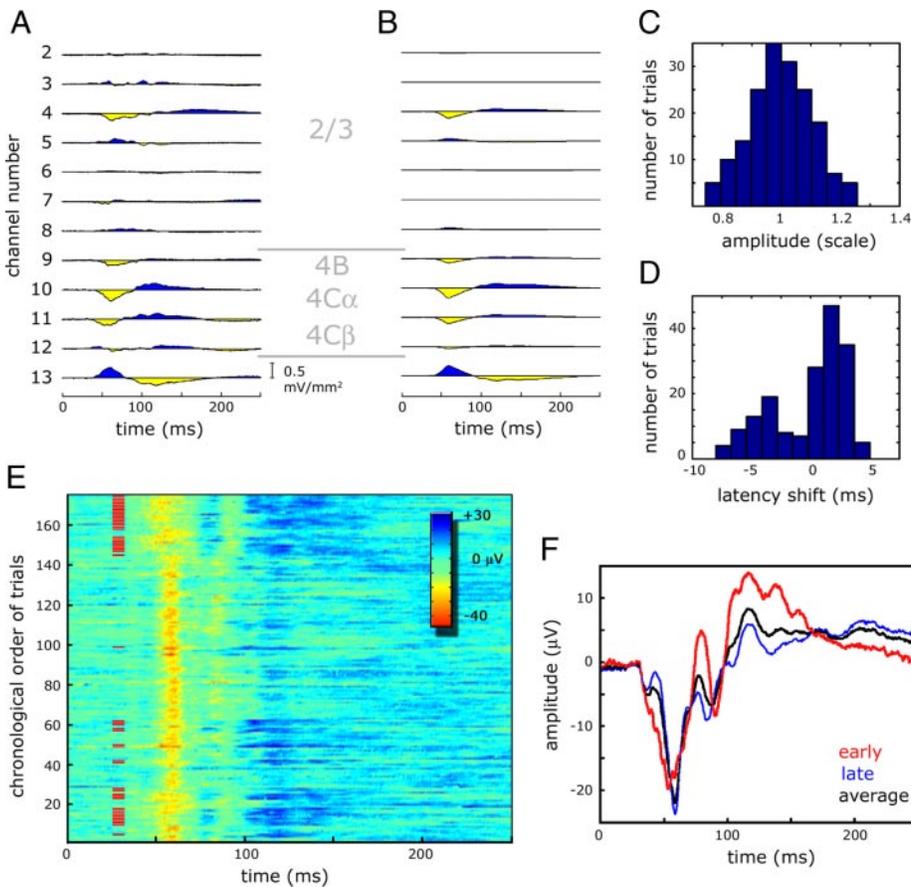


FIG. 7. Single-component analysis of V1 responses. *A*: CSD profile of the trial-average event-related potential shows granular and supragranular activation in macaque V1 in response to a red-light flash (number of trials = 171). *B*: estimating a single component with dVCA results in a waveshape with a CSD profile that captures the most prominent responses in the data. *C*: a histogram of the single-trial amplitudes of this component shows that the amplitude rarely varies more than  $\pm 20\%$ . *D*: histogram of the single-trial latencies reveals that there are 2 response modes: an early response that happens in 1/3 of the trials and a late response that happens in about 2/3 of the trials. Peak latency difference between these 2 modes of activation is 6.75 ms. *E*: to verify the existence of these 2 activation modes, this figure shows all 171 trials of the field potential recorded in channel 10. Each trial is represented as a horizontal line with its time-varying color representing the time-varying amplitude of the field potential. Trials designated as belonging to the “early” subset are indicated by the red dashes on the left side of the plot. Note that the “early” trials are characterized by a negative (yellow) field potential onset occurring before the onset seen in the “late” trials. *F*: further verification of this finding is provided by comparing the average FP recorded in channel 10 with the average obtained from the “early” subset and the average obtained from the “late” subset. Although both the early and late responses show initial activation occurring at the same time, the early response’s activation continues to grow, whereas the late response’s activation decreases before growing again. Difference in latency between the minima of the 2 subaverages is 5.5 ms.

grouped over trials, which suggests a slow cyclical shift between different response states with one state dominated by faster and the other by slower responses. Second, we performed a cross-correlation analysis between each single-trial waveshape and the trial-averaged, FP waveform between 25 and 95 ms (this time period captures the initial negative deflection of the averaged FP waveform). This analysis also yielded a bimodal distribution (not shown). Finally, we selectively averaged the early and late FP recordings in channel 10 and showed them to be significantly different in both latency and waveshape (Fig. 7*F*). The trial-averaged FP waveshape (black) lies between the early (red) and late (blue) waveshapes. As expected, the minimum of the early averaged responses occurs before the minimum of the late averaged responses (5.5-ms difference). Although this time difference reflects the fact that these early responses precede the late responses, it is however not as accurate as the dVCA estimate because it is derived from a single time point rather than from the entire waveshape. It is also not surprising that the late responses more closely resemble the ensemble-averaged FP because they outnumber the early responses by more than 2 to 1.

Because these two types of granular responses exhibit distinct waveshapes and latencies this suggests that different physiological mechanisms are at work in these two processes. These mechanisms will be much easier to distinguish by separating the original data set into two subsets—an early subset and a late subset—and performing separate analyses on each and comparing. We then used dVCA to reestimate the first component in each subset. Figure 8, *A* and *B* shows the CSD of the first component from the early and late subsets,

respectively. The onset latency of the prominent granular sink over source configuration was determined by descending down the left side of the peak and identifying the point in time where five consecutive previous data points were monotonically increasing. With this measure the response onset latency of the early subset (Fig. 8*A*) was 28 ms, which is clearly earlier than its counterpart in the late subset (Fig. 8*B*), which was 42.5 ms, as indicated by the drop lines. The waveshapes of the two types of Layer 4 responses differ significantly, indicating a difference in activation patterns between subsets. This is also confirmed by the different degrees to which this activation appears in the supragranular layers because the later responses seem to be more strongly coupled to the supragranular activation than the early responses.

After any analysis, the residual signals must be examined because they represent activity that was not modeled by the algorithm. This residual activity is often a combination of unstructured “noise” and structured unmodeled signal. In Fig. 8*C* we show the trial-averaged FP residuals for the early (red) and late (blue) subsets for odd-numbered electrode channels. First, the black arrows indicate activation that is almost identical in both subsets (mean peak time at 37.75 ms for the early subset and 37.68 ms for the late subset). Because of the timing and the laminar distribution of this field potential, we believe that this negativity reflects the initial signal from the thalamocortical afferents. The early responses onset just after onset of this thalamic signal, which occurs at about 25 ms, whereas the late responses are delayed. After 50 ms the residual signals from the two subsets diverge significantly. This difference is striking in channel 1 where the oscillations are 180° out of

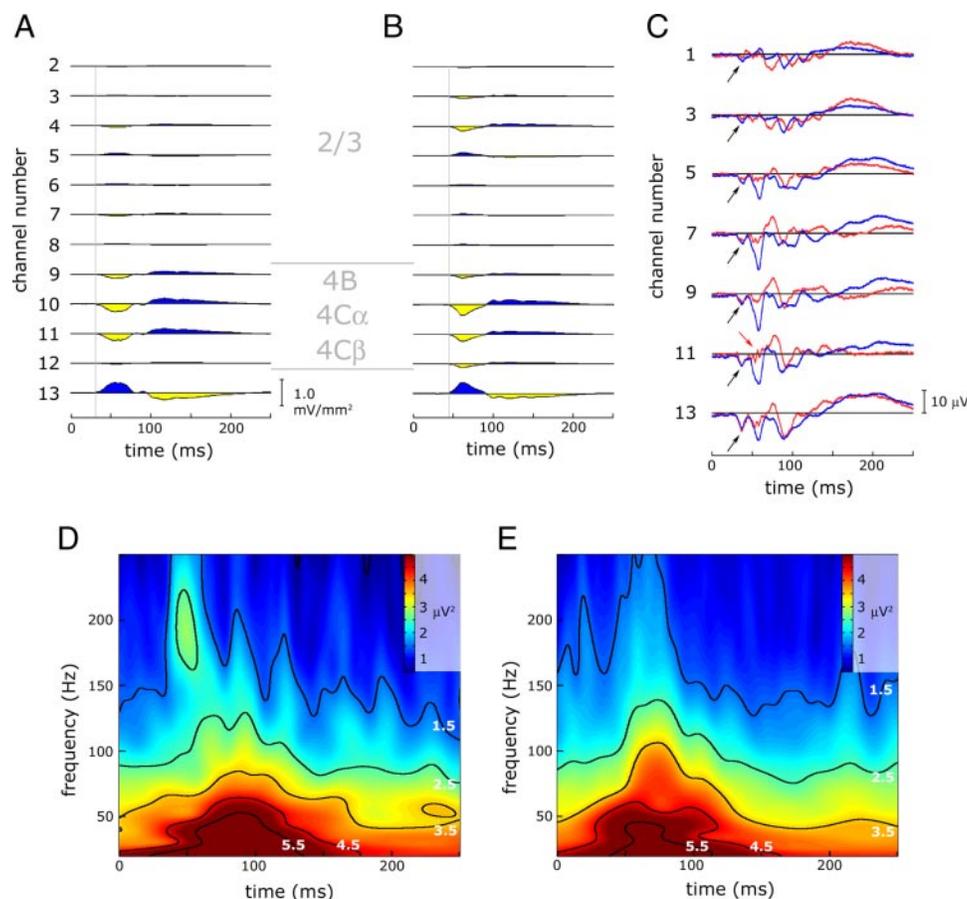


FIG. 8. Examination of the early and late responses. Data set has been split into 2 subsets: the early subset and the late subset, and a single component has been reestimated for each subset. *A*: CSD profile of the early component with a drop line showing onset of the major response at 28 ms. *B*: CSD profile of the late component. Drop line shows its major response onset at 42.5 ms. Wave-shapes of the 2 responses are noticeably different, as are their laminar profiles. *C*: average residual field potentials, computed by subtracting the single-trial model from the single-trial data and averaging over all trials, further demonstrates the differences between these 2 modes of activation (early: red; late: blue). Only the odd channels are shown. These residuals represent responses not modeled by the single component in each subset. Black arrows indicate what is most likely the thalamic input, which is time locked to the stimulus because it is visible in the average. Subsequent activation of the 2 response types is very different. Late oscillations in channel 1 after 100 ms are 180° out of phase. Red arrow in channel 11 at about 50 ms shows time-locked ultra-high-frequency (UHF) oscillations in the early responses that are not present in the late subset. *D*: wavelet analysis was performed on the residuals to characterize these oscillations and verify that the categorization indicated by the dVCA results is justified. Residuals in the early subset show a burst of UHF oscillation between 42 and 57 ms with frequency ranging from 160 to 220 Hz. *E*: these oscillations are completely absent in the late subset.

phase. Of significant interest are the ultra-high-frequency (UHF; 160–220 Hz) oscillations in the granular and the supra-granular layers of the early subset (red arrow in channel 11). These UHF oscillations, which onset with the initial granular response, are absent in the averaged residuals of late responses.

To verify that these UHF oscillations were not present in the late responses, we applied a Morlet-based wavelet transformation to the two residual subsets of channel 11. To preserve information about time–frequency behavior of oscillations not precisely time-locked to the stimulus, we applied the wavelet transform to each single trial separately and the wavelet results were then averaged. These results show that there is a burst of UHF activity ranging from about 160 to 220 Hz and occurring between 42 and 57 ms in the early subset, which is represented as an “island of activity” in the time–frequency plot (Fig. 8*D*). Such an island of activity is completely absent in the late subset (Fig. 8*E*). This finding demonstrates the ability of dVCA to identify distinct physiological modes of activation using the single-trial characteristics of the responses, even in situations where the data are unfiltered.

To demonstrate how dVCA can be used to extract multiple components, we focus on the late subset. A complete analysis of the entire data set will be published elsewhere. In this example we model three components (Fig. 9, *A–C*). Components 2 and 3 (*c2* and *c3*) have been magnified by a factor of 2 to make them easily visible, and are thus not on the same scale as component 1 (*c1*), which is the largest response. To ensure that the SNR is within the acceptable range for dVCA, we used Eq. B4 to compute the SNR for each component in

each channel. For a given component this is accomplished by estimating the SD of the component in that channel and the SD of the residual signal in that channel. The average SNR for each component was then computed by averaging its SNRs across channels. We found that all three components are well within the range of acceptable performance with average SNRs of 1.59, 1.68, and 0.98 dB for *c1*, *c2*, and *c3*, respectively.

First, *c1* is in all practical respects identical to the single component we previously estimated from the late subset. The prominent granular sink over source configuration is activated at about 35.5 ms in this more detailed model. Although the onset time is easy to quantify in the large biphasic response, components *c2* and *c3*, which onset with low-amplitude oscillatory behavior, required a more sophisticated onset measure. By fitting the preresponse interval with a line, we computed the SD of the preresponse signal deviations from that line. Component onset was then defined as the first of five consecutive time points where the component amplitude was greater than this line by 2 SDs. This was useful because we did not filter the data, and low-frequency oscillations could produce confounding baseline deviations. With this measure, we found small, but significant, activations in *c1* as early as 26 ms (see drop line in Fig. 9*A*) occurring after the thalamic signal at 25 ms and before the massive biphasic response at 35.5 ms.

A low-frequency biphasic response is described by *c2* (Fig. 9*B*). Significant onset of this component occurs at 37 ms (see drop line), which is after *c1* begins the biphasic activation in the granular layers. Note that the source and sink distribution of *c2* are in superficial layers compared with the small supra-

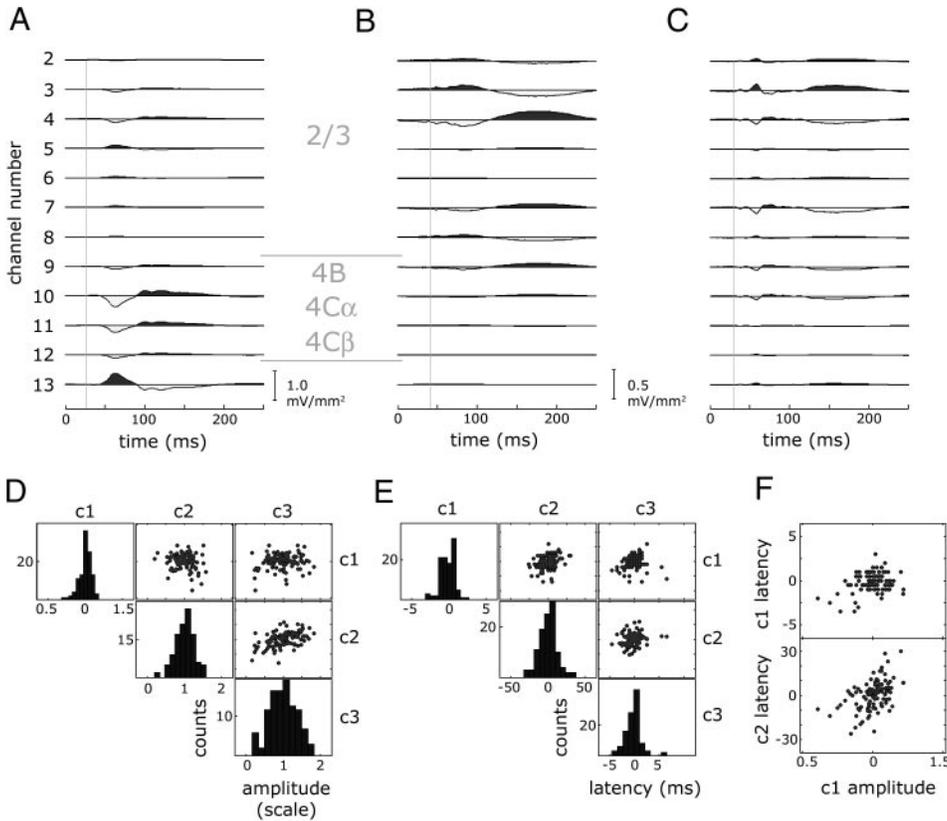


FIG. 9. Three components were estimated from the late subset. *A*: CSD profile of component 1 shows that it represents the granular response with some activation in the supragranular layers. Drop line marks the onset of the response triggered by the thalamic input at 26 ms. *B*: component 2 represents slow-wave activity in the supragranular layers. Its onset is considerably later at 37 ms. *C*: component 3 also describes supragranular activation with a pulslike activation followed by some slow wave activity. Similarity in the laminar profiles of c2 and c3 suggests that these responses are probably taking place in the same population of cells. Note, however, that these laminar profiles are not identical as c3 displays some granular activation. *D*: single-trial amplitude histograms and amplitude scatterplots. C1 shows very little amplitude variability, whereas that of c2 and c3, although comparable, are not equal. Scatterplot of c2 and c3 amplitudes are correlated with  $r = 0.495$  ( $P < 10^{-7}$ ). *E*: single-trial latency histograms and scatterplots. Note that the slow-wave activity in c2 shows great latency variability. Scatterplot of c1 and c2 latencies shows some correlation at  $r = 0.327$  ( $P < 0.001$ ). *F*: these scatterplots show the relationship between c1 amplitudes and c1 and c2 latencies with  $r = 0.297$  ( $P < 0.002$ ) and  $r = 0.483$  ( $P < 10^{-7}$ ), respectively.

granular response of c1. At first glance it appears that the spatial distribution of c3 is the same as that of c2. However, whereas these components may involve the same populations of supragranular neurons, c3 also involves others in the granular layers. The initial activation of c3 is at 30 ms, but it is not until 45 ms when the component begins to display a biphasic activation followed later by a slow wave.

Figure 9*D* shows the single-trial amplitude histograms along the diagonal, along with scatterplots showing the relationships between the single-trial amplitudes of the three components. The initial granular response, c1, has the lowest amplitude variability ( $\sigma_{amp1} = 0.102$ ), whereas the later supragranular responses display greater variability ( $\sigma_{amp2} = 0.248$ ,  $\sigma_{amp3} = 0.362$ ). Figure 9*E* shows the single-trial latency results along with scatterplots depicting the trial-by-trial relationships between component latencies. Component c1 again displays the least variability ( $\sigma_{lat1} = 1.08$  ms), but in this case, c2 displays greater variability than that of c3 ( $\sigma_{lat2} = 9.85$  ms,  $\sigma_{lat3} = 1.58$  ms).

Correlations in the scatterplots between component parameters indicate dynamic interactions among these components. Of the amplitude–amplitude interactions (Fig. 9*D*), we see that the amplitudes of c2 are correlated with c3 ( $r = 0.495$ ,  $P < 10^{-7}$ ), so that when c2 is larger than average, c3 is also larger than average. Such correlation might suggest that these components have not been separated. However, note that the latency variability of each component is very different (Fig. 9*E*), and that there is little correlation between the latency of c2 and the latency of c3 ( $r = 0.171$ ,  $P = 0.075$ ).

Of the latency–latency interactions, the largest correlation is between the c1 and c2 latencies ( $r = 0.327$ ,  $P < 0.001$ ), so that when c1 is earlier than average, c2 is also earlier than average.

More interesting are the amplitude–latency interactions. Figure 9*F* shows the two most probable relationships. First the amplitude and latency of c1 are correlated ( $r = 0.297$ ,  $P < 0.002$ ), so that when c1 is early its response is smaller, and when it is late its response is larger. Self-interactions, such as this one illustrated by c1, contain much information about the underlying dynamics of the neural response. More apparent is the relationship between the amplitude of c1 and the latency of c2 ( $r = 0.483$ ,  $P < 10^{-7}$ ). In this case, when the c1 granular response is larger than average, c2 onsets later than average. This result is somewhat counterintuitive because generally we expect that if the c1 is driving the c2, a larger c1 might produce an earlier-onset c2. The difference between c1 and c2 is further highlighted by noting that the amplitude of c1 is anticorrelated to the latency of c3 ( $r = -0.190$ ,  $P = 0.048$ ) (not shown).

## DISCUSSION

Detailed studies of single-trial interactions among neural components in a variety of experimental situations have the potential of providing novel insight into the information-processing strategies used by the brain. In this paper, our specific goal was to describe the dVCA algorithm and to show how it can tease apart evoked responses in different neuronal ensembles and facilitate the study of their dynamic interactions.

Maximum-likelihood techniques have previously been used to approach the problem of trial-to-trial variability of evoked responses. These past works have relied on single-component signal models that incorporate latency variability (Pham et al. 1987; Woody 1967), and more recently both amplitude and latency variability (Jaskowski and Verleger 1999) to characterize evoked responses. Multiple-component models were

introduced by Lange et al. (1997) by adopting a template model of the ERP waveshapes, and were used to characterize both the amplitude and latency of multiple components in single-channel EEG recordings. Early versions of the dVCA algorithm also dealt with multiple-component, single-channel estimation (Knuth et al. 2001; Truccolo et al. 2001, 2002, 2003), but instead used a completely general waveshape model that did not rely on a given functional form. The dVCA algorithm presented here allows one to use multiple-channel recordings to identify and characterize multiple ERP components in the single trial by taking advantage of the trial-to-trial variability.

The dVCA algorithm was derived by approaching the problem as an exercise in Bayesian parameter estimation. There are several advantages to a Bayesian approach. First, the strategy is model based in the sense that given a quantitative model of the phenomena of interest, probability theory can be used to estimate the values of the model parameters from the data. Any failure in the algorithm can be traced back to either inadequacies of the model in representing the physical situation, assumptions, or approximations made in its implementation, or a situation produced by an insufficiently informative data set. Second, inadequacies in the model, once identified, are readily remedied—leading to improvements in the algorithm. Third, once the model components have been adequately estimated, the residual data can be examined to identify previously hidden phenomena. For example, herein we have discovered UHF oscillations associated with early initial granular layer responses and absent during the more frequent late responses. By accurately identifying the evoked components in the single-trial epochs, one can more accurately study the ongoing activity, which has been purported to contain signals important for communication among brain regions (Bressler et al. 1993; Singer 1993; Truccolo et al. 2003), in perception, working memory, and sensorimotor integration (Engel et al. 2001; Lee et al. 2003). Finally, the Bayesian methodology allows one to incorporate additional prior information into the problem to improve one's inferences, which is a major advantage that we plan to capitalize on in future work.

dVCA has a number of technical strengths. The first is that it capitalizes on both amplitude and latency variability to aid in component separation. Second, neither the components nor their underlying neural sources are assumed to be independent of one another. This avoids the adoption of physiologically implausible assumptions and enables one to study the dynamical interactions among neural sources. A third strength is that by accounting for trial-to-trial variability in amplitude and latency, one is able not only to quantify the interactions among sources but also to study their time-dependent properties over the duration of the entire experiment. Single-trial analysis also allows dVCA to detect components, which are not present in every experimental trial, thus allowing an experimenter to study cognitive or sensorimotor processes, which may use different processing strategies at different times. Additionally, one can use dVCA to study how systemic factors such as attention, arousal, or varying disease states may modify neural responses on both a single-trial and a single-component basis. Finally, although no explicit prior information about the values of the model parameters was included in the development of the algorithm, much prior information went into the choice of the model. This is in contrast to other approaches such as ICA

and other general BSS techniques, which strive to make very general assumptions about the model and the distributions of the parameter values (Knuth 1997, 1999).

dVCA is robust in the presence of noise, allowing accurate estimates of all parameters down to SNRs on the order of  $-15$  dB for white Gaussian noise and down to around  $0-1$  dB for correlated far-field noise. The estimation of far-field signals in the presence of correlated far-field noise was difficult. This is undoubtedly explained by the fact that this noise possesses the same spatial distribution as the source, making the two difficult to distinguish. Although this is a problem for local field potential measurements, it will not be a problem for whole-head recordings because, in that case, a far-field source for one detector will be local for another. Also in most cases, the behavior of the background noise will fall somewhere between the pure white Gaussian noise and pure correlated far-field noise extremes. When using dVCA, we recommend that one calculate the SNR of the estimated components to ensure that the algorithm is operating in a regime where the quality of the parameter estimates is guaranteed. In addition, the spatial distribution of the sources across the array should be examined to ensure that they are physiologically reasonable. We have been able to construct cases where the sources remain inseparable by the algorithm. This can typically be detected by examining the CSD map of the estimated components across the array. In cases where two sources remain mixed, each is characterized by multiple sets of neural sources in identical locations. However, such a criterion is not necessarily conclusive in intracortical laminar recordings because the cortical architecture allows for only a small number of current sources and sinks, and tight couplings between sources is extremely likely (Shah et al. 2003). One possible solution to this problem is to restart the algorithm using starting conditions consistent with the source locations indicated by the CSD. The value of the posterior probability of these solutions obtained from different starting points can also be computed and compared (when using the same model order) to ensure that one has the most probable solution and is not merely stuck at a local maximum. Searching such enormous solution spaces is a difficult problem faced by all source separation algorithms.

We are currently working to improve the dVCA algorithm along several lines.

First, there are situations where the experimenter has knowledge about the forward problem, which describes the propagation of the signals to the detectors. Such knowledge can be incorporated by adopting a more specific source model (e.g., current dipole model) and expressing the coupling coefficients in terms of the new ERP source parameters, detector coordinates, and head geometry (Knuth and Vaughan 1998; Schmidt et al. 1999), or by using information about the propagation law to derive prior probabilities for the coupling matrix elements (Knuth 1998, 1999, 2005). Similarly, more detailed information about the correlation structure of the background noise can be used to derive more accurate likelihood functions (Sivia 2003). After much experimentation in a specific cortical region, one can assign Gaussian distributions to the priors for the amplitude and latency variability of specific sources based on previously observed means and variances.

Second, by using a discrete model of the component waveshapes, we are restricted to estimating discrete values of onset latency shift. By adopting a waveshape model that relies on a

set of continuous basis functions, continuous values of latency shift could be investigated. An example of a continuous time model is the frequency domain model of the ERP used by Pham et al. (1987) and Jaskowski and Verleger (1999). They model the waveshape as a sum of a discrete set of sinusoids, which is continuous in the time domain, but discrete in the frequency domain. In principle, such a model allows one to describe latency shifts with arbitrary precision. It should also be noted that dVCA does not accommodate trial-dependent waveshape changes or spatiotemporal wave activity across cortex. Incorporating such activity into the signal model will require great care not to erroneously overfit the data.

Third, our algorithm represents a maximum a posteriori (MAP) estimate based on an iterative, or fixed-point, solution. Although each step in our algorithm has intuitive appeal, it is perhaps not the most effective means for obtaining a solution. Markov chain Monte Carlo (MCMC) with simulated annealing is well suited to performing simultaneous model selection calculations and parameter estimations. Several years ago, we briefly investigated this approach only to find it to be extremely time consuming given the large number of parameters in the model. It is expected that given the advances in computing power and in MCMC technology, such an approach would most likely outperform the algorithm presented here by automatically identifying the number of components warranted by the data, avoiding local optimal solutions, and readily providing error bars for the results.

Fourth, oscillatory responses of sufficiently low frequency pose great difficulties when it comes to accurate estimation of their onset latency. Furthermore, ongoing oscillations that are weakly time locked to stimuli are not well described by the multiple-component event-related model and suggest that modifications in the model to allow for ongoing oscillations or oscillatory bursts would lead to new results.

Finally, in principle this algorithm is equally applicable to human scalp EEG and MEG data. In fact, the Woody filter, which was the first method to estimate single-trial evoked responses, was used on human scalp EEG with success (Woody 1967). The signal-to-noise levels are certainly well within the applicable range (0 to  $-2$  dB, or 1.0 to 0.1 in terms of SD ratios) for many of the evoked responses currently studied. However, in practice, several challenges remain. The number of detectors used in these paradigms is typically an order of magnitude greater than those we have demonstrated here. It is certain that the algorithm in its present implementation will run more slowly. Also, whole-head studies expose the experimenter to an order-of-magnitude more sources than we work with in our intracortical recordings. This increases the possibility that the dVCA algorithm can become trapped in nonoptimal local solutions. This of course is a potential problem for all source separation and localization techniques, and dVCA is no exception. We are beginning to examine the application of dVCA in human scalp studies, and expect that if such pathological solutions are encountered, they can be avoided by using more sophisticated search algorithms.

#### APPENDIX A: ALGORITHM DEVELOPMENT

Bayes' theorem is the natural starting point for deriving the dVCA algorithm because it allows one to describe the probability of the model in terms of the likelihood of the data and the prior probability of the model parameters

$$p(\text{model}|\text{data}, I) = \frac{p(\text{data}|\text{model}, I)p(\text{model}|I)}{p(\text{data}|I)} \quad (\text{A1})$$

where  $I$  represents any prior information one may have about the physical situation. The probability on the left-hand side of Eq. A1 is referred to as the *posterior probability*. It represents the probability that a given set of hypothesized values of the model parameters accurately describes the physical situation in which the data were collected. The first term in the numerator on the right-hand side—the likelihood of the data given the model—describes the degree of accuracy with which we believe the model can predict the data. The second term in the numerator is the prior probability of the model, or the *prior*. This prior describes the degree to which we believe the hypothesized values of the model parameters to be correct based only on our prior information about the problem. The term in the denominator is called the *evidence*, which in this problem acts only as a normalization factor. It is through the assignment of the likelihood and the priors that we express all of our knowledge about the particular source separation problem. Bayes' theorem can be viewed as describing how one's prior probability,  $p(\text{model}|I)$ , is modified by the acquisition of some new information in the form of data.

To apply this theorem to our problem, we consider the change in our knowledge about the model with the acquisition of new data, which consists of the set of recorded data  $\mathbf{x}(t)$  from  $M$  detectors over the course of  $R$  trials. In this case, Bayes' theorem can be written as

$$p[\mathbf{C}, \mathbf{s}(t), \boldsymbol{\alpha}, \boldsymbol{\tau}|\mathbf{x}(t), I] = \frac{p[\mathbf{x}(t)|\mathbf{C}, \mathbf{s}(t), \boldsymbol{\alpha}, \boldsymbol{\tau}, I]p[\mathbf{C}, \mathbf{s}(t), \boldsymbol{\alpha}, \boldsymbol{\tau}|I]}{p[\mathbf{x}(t)|I]} \quad (\text{A2})$$

where boldface symbols represent the entire set of parameters of each type in the mcERP model, e.g.,  $\boldsymbol{\alpha} = \{\alpha_1, \alpha_2, \dots, \alpha_R\}$ . The most probable set of model parameters maximizes the probability in Eq. A2, and thus in practice the equation can be expressed as a proportionality with the inverse of the evidence  $p[\mathbf{x}(t)|I]$  as the implicit proportionality constant. Equation A2 then becomes

$$p[\mathbf{C}, \mathbf{s}(t), \boldsymbol{\alpha}, \boldsymbol{\tau}|\mathbf{x}(t), I] \propto p[\mathbf{x}(t)|\mathbf{C}, \mathbf{s}(t), \boldsymbol{\alpha}, \boldsymbol{\tau}, I]p[\mathbf{C}, \mathbf{s}(t), \boldsymbol{\alpha}, \boldsymbol{\tau}|I] \quad (\text{A3})$$

For simplicity, the joint prior can be factored into four terms, each describing what we know about the source-detector coupling, the source waveshapes, the single-trial amplitudes, and the single-trial latency offsets

$$p[\mathbf{C}, \mathbf{s}(t), \boldsymbol{\alpha}, \boldsymbol{\tau}|\mathbf{x}(t), I] \propto p[\mathbf{x}(t)|\mathbf{C}, \mathbf{s}(t), \boldsymbol{\alpha}, \boldsymbol{\tau}, I]p(\mathbf{C}|I)p(\mathbf{s}(t)|I)p(\boldsymbol{\alpha}|I)p(\boldsymbol{\tau}|I) \quad (\text{A4})$$

For the amplitude and latency priors,  $p(\boldsymbol{\alpha}|I)$  and  $p(\boldsymbol{\tau}|I)$ , respectively, we assign uniform densities with appropriate cutoffs denoting a range of physiologically realizable values. Note that a joint uniform density  $p(\boldsymbol{\alpha}, \boldsymbol{\tau}|I)$  can always be factored in this way, and is not necessarily a statement about independence. Because the amplitude and latency priors are each represented by a uniform density, we can absorb those two factors into the implicit proportionality constant.

Our derivation continues by using the principle of maximum entropy to assign a Gaussian likelihood (e.g., Jaynes 2003; Sivia 1996) by introducing a parameter  $\sigma$  reflecting the expected square deviation between our predictions and the mean

$$p[\mathbf{C}, \mathbf{s}(t), \boldsymbol{\alpha}, \boldsymbol{\tau}, \sigma|\mathbf{x}(t), I] \propto (2\pi\sigma^2)^{-\text{MRT}/2} \exp\left[-\frac{1}{2\sigma^2} Q\right] p(\sigma|I) p(\mathbf{C}|I) p(\mathbf{s}|I) \quad (\text{A5})$$

where  $p(\sigma|I)$  is the prior probability for  $\sigma$ .  $Q$  represents the sum of the square of the residuals between the data and our model in Eq. 1, written

$$Q = \sum_{m=1}^M \sum_{r=1}^R \sum_{t=1}^T [x_{mr}(t) - \sum_{n=1}^N c_{mn} \alpha_{nr} s_n(t - \tau_{nr})]^2 \quad (\text{A6})$$

where  $M$  represents the number of detectors,  $R$  is the number of

experimental trials, and  $T$  is the number of recorded time points per trial. Some consider this a fancy way to state that the noise is Gaussian distributed. However, as Jaynes (2003) demonstrates, this is only a statement that the variance of the noise is equal to some value  $\sigma^2$ . Thus higher-order statistical structure in the noise is tolerated and the noise does not have to be Gaussian distributed.

The fact that we do not actually know the value of  $\sigma$  is not a problem because we can obtain a conservative result by considering all possible noise levels by integrating the joint posterior over all possible values of  $\sigma$ . Symmetry considerations require that we assign what is called a Jeffreys prior for  $\sigma$ ,  $p(\sigma|I) = \sigma^{-1}$  (Sivia 1996). Performing the integral of the joint posterior over all possible values of  $\sigma$ , we obtain a marginal posterior probability for our original set of model parameters. To do this, it helps to make the change of variables by defining  $t = \sigma^{-1}$  and then integrating over  $t$  by iteratively integrating by parts. The proportionality constant one obtains depends on whether the product  $MRT$  is even or odd, so there are two cases to consider. However, the functional form of the result is independent of this fact

$$p[\mathbf{C}, \mathbf{s}(t), \boldsymbol{\alpha}, \boldsymbol{\tau} | \mathbf{x}(t), I] \propto Q^{-MRT/2} p(\mathbf{C}|I) p(\mathbf{s}|I) \quad (\text{A7})$$

Some may recognize this as being related to the Student's  $t$ -distribution (Student 1908).

If we were more knowledgeable, prior information regarding the source waveforms could be used to improve our inferences. In addition, knowledge of the source-detector coupling, which is found by solving the electromagnetic forward problem, could be used to create an algorithm that simultaneously performs source separation and localization (Knuth 1998; Knuth and Vaughan 1998). For simplicity, in this development we choose to assign uniform priors to  $p(\mathbf{C} | I)$  and  $p(\mathbf{s} | I)$ , and absorb the terms into the implicit proportionality constant. It is more convenient to work with the logarithm of the posterior probability (Eq. A5), which can be compactly written as

$$\ln P = -\frac{MRT}{2} \ln Q + \text{const} \quad (\text{A8})$$

where  $P$  is the posterior probability  $p[\mathbf{C}, \mathbf{s}(t), \boldsymbol{\alpha}, \boldsymbol{\tau} | \mathbf{x}(t), I]$ .

An iterative algorithm to identify mERPs, which we call *differentially variable component analysis* (dVCA), is implemented by solving Eq. A8 for the most probable set of model parameters. Such a solution is called the *maximum a posteriori* (MAP) estimate. This most probable set of model parameters is represented as a peak in the space of posterior probabilities. At this maximum, the partial derivative of the posterior probability with respect to any one of the model parameter values is zero. For this reason, our first step in deriving an optimal estimate of the component waveshape by equating each of the partial derivatives of Eq. A8 with zero. In the current dVCA implementation, each source waveshape is modeled as a pointwise time series where each point of the time series is considered as a model parameter. To solve this we must look at the partial derivative of the log posterior with respect to the waveshape amplitude at each of the times over which the waveshape is defined. The partial derivative with respect to the  $j$ th component waveshape at time  $q$  gives

$$\frac{\partial \ln P}{\partial s_j(q)} = -\frac{MRT}{2} Q^{-1} \frac{\partial Q}{\partial s_j(q)} \quad (\text{A9})$$

with

$$\frac{\partial Q}{\partial s_j(q)} = -2 \sum_{m=1}^M \sum_{r=1}^R [WC_{mj}\alpha_{jr} - (C_{mj}\alpha_{jr})^2 s_j(q)] \quad (\text{A10})$$

where

$$W = x_{mr}(q + \tau_{jr}) - \sum_{\substack{n=1 \\ n \neq j}}^N C_{mn}\alpha_{nr}s_n(q - \tau_{nr} + \tau_{jr}) \quad (\text{A11})$$

The term  $W$  is important because it deals with the data, which has been time-shifted according to the latency shift of the component being estimated,  $x_{mr}(q + \tau_{jr})$ . From this, one subtracts all other components after each has been appropriately scaled and time-shifted,  $C_{mn}\alpha_{nr}s_n(q - \tau_{nr} + \tau_{jr})$ . The derivative of the log probability is zero when the scaled, estimated waveshape equals  $W$ . Thus one can obtain an expression for the optimal waveshape of the  $j$ th source at time  $q$  in terms of the other sources

$$\hat{s}_j(q) = \frac{\sum_{m=1}^M \sum_{r=1}^R WC_{mj}\alpha_{jr}}{\sum_{m=1}^M \sum_{r=1}^R (C_{mj}\alpha_{jr})^2} \quad (\text{A12})$$

Similarly for the source amplitudes, one obtains an optimal estimate by considering the derivative of the log probability with respect to the amplitude of the  $j$ th source during the  $p$ th trial. Setting this derivative equal to zero results in

$$\hat{\alpha}_{jp} = \frac{\sum_{m=1}^M \sum_{r=1}^T [U V]}{\sum_{m=1}^M \sum_{r=1}^T V^2} \quad (\text{A13})$$

where

$$U = \left[ x_{mp}(t) - \sum_{\substack{n=1 \\ n \neq j}}^N C_{mn}\alpha_{np}s_n(t - \tau_{np}) \right] \quad (\text{A14})$$

and

$$V = C_{mj}s_j(t - \tau_{jp}), \quad (\text{A15})$$

such that the solution is given by the projection of the detector-scaled component  $C_{mj}s_j(t - \tau_{jp})$  onto the data after removing the other scaled, time-shifted components. This can be viewed in terms of a dot product, which is related to a matching filter solution.

The optimal source-detector coupling coefficients are found similarly with

$$\hat{C}_{ij} = \frac{\sum_{r=1}^R \sum_{t=1}^T [X Y]}{\sum_{r=1}^R \sum_{t=1}^T Y^2} \quad (\text{A16})$$

where

$$X = \left[ x_{ir}(t) - \sum_{\substack{n=1 \\ n \neq j}}^N C_{in}\alpha_{nr}s_n(t - \tau_{nr}) \right] \quad (\text{A17})$$

and

$$Y = \alpha_{jr}s_j(t - \tau_{jr}) \quad (\text{A18})$$

Estimating the latency shift of the  $j$ th source during the  $p$ th trial using the approach taken for the other parameters leads to a complex solution as the latency appears implicitly as the argument of the waveshape function. Instead we examine the necessary conditions for maximizing the quadratic form  $Q$ . Expanding the square in Eq. A6,

one can see that as the latency shift  $\tau_{jp}$  is varied, only the cross-terms corresponding to the  $j$ th source change. The optimal estimate of the latency shift  $\hat{\tau}_{jp}$  can be found by maximizing the cross-correlation between the estimated source and the data after the contributions from the other sources have been subtracted such that

$$\hat{\tau}_{jp} = \arg \max Z(\tau_{jp}) \quad (\text{A19})$$

where

$$Z(\tau_{jp}) = \sum_{m=1}^M \sum_{t=1}^T \left[ C_{mj} \alpha_{jp} s_j(t - \tau_{jp}) \left( x_{mp}(t) - \sum_{\substack{n=1 \\ n \neq j}}^N C_{mn} \alpha_{np} s_n(t - \tau_{np}) \right) \right] \quad (\text{A20})$$

In practice, as a discrete model is being used for the component waveshapes  $\mathbf{s}(t)$ , we use a discrete set of latency shifts with resolution equal to the sampling rate.

Perhaps the greatest challenge is determining the number of sources warranted by the data. In our investigations to date we have been focused on understanding the major sources responsible for the recorded data. This typically entails first modeling a single component and examining the un-modeled residual signal in detail. We then attempt to model the data with two components and so on to greater numbers of sources. As we demonstrate, the responses we are examining are sufficiently complex that this approach rapidly reveals previously unknown characteristics of the neurophysiology.

There are many ways that one can use the equations above to implement the dVCA algorithm. A useful iterative method (Fig. 1) begins by modeling a single neural source with the single-trial amplitudes set to unity and the latency shifts set to zero.

1) Event-related potentials (ERPs) are computed for each detector by averaging the data over all trials. The ERP for each detector is full-wave rectified, and its integral (area under the curve) is approximated. The ERP having the largest area, or total signal content, is chosen as the initial approximation of the first component waveshape.

2) The single-trial amplitude scales are all initialized to one and the single-trial latency shifts are initialized to zero. This is consistent with the implicit assumptions of the ERP.

3) The coupling matrix is estimated using Eq. A16.

4) Single-trial latency shifts are estimated using Eq. A19.

5) Single-trial amplitudes are estimated using Eq. A13.

6) Equations A12, A16, A19, and A13 are then iterated until the average change in the waveshape from the previous iteration is  $<1\%$  or until a maximum number of iterations has been performed.

7) At this point the residual signal for each detector is computed by subtracting the model from the data, as in the argument of Eq. A6. The residual signals are then averaged across trials to obtain a residual ERP for each channel.

8) The investigator can explore the benefit of introducing a new source to be modeled.

9) The initial approximation of the next component waveshape is chosen to be the residual ERP of the detector having the largest total signal.

10) The single-trial amplitudes and latency shifts of the new component are set to one and zero, respectively.

11) Equations A12, A16, A19, and A13 are used to obtain the parameters accompanying the new component in addition to refining the estimate of the first component. The set of equations is further iterated to refine the solutions until the average changes in each component waveshape is again  $<1\%$  or until the maximum number of iterations has been reached.

12) Additional components are added until the investigator chooses to stop.

This implementation represents only one possible approach to using these equations to obtain useful solutions.

## APPENDIX B: AMARI ERROR

The Amari error (Amari et al. 1996) derives from the fact that the estimated coupling matrix  $\hat{\mathbf{C}}$  can at best be determined to within a scaled permutation of the true coupling matrix

$$\hat{\mathbf{C}} = \mathbf{C}\mathbf{\Sigma}\mathbf{\Pi} \quad (\text{B1})$$

where  $\mathbf{\Sigma}$  is a diagonal scaling matrix and  $\mathbf{\Pi}$  is a permutation matrix. Here the quantities decorated with a caret (also known as a "hat") denote estimated quantities and those without a caret denote the unknown true values. Ideally, the source estimates are therefore related to the original sources by a simple matrix transformation

$$\hat{\mathbf{s}} = \mathbf{M}\mathbf{s} \quad (\text{B2})$$

where  $\mathbf{s}$  is the  $N \times T$  matrix of the original source waveshapes,  $\hat{\mathbf{s}}$  is the  $N \times T$  matrix of the estimated source waveshapes, and  $\mathbf{M}$  is a transformation matrix found by

$$\mathbf{M} = \mathbf{\Pi}^{-1}\mathbf{\Sigma}^{-1}. \quad (\text{B3})$$

The deviation of  $\mathbf{M}$  from the ideal form suggested in Eq. B3 provides a measure of the quality of separation. By postmultiplying both sides of Eq. B2 by  $\mathbf{s}^T$ , we form a square matrix  $(\mathbf{s}\mathbf{s}^T)$  postmultiplying  $\mathbf{M}$  on the right-hand side. Assuming that this square matrix is invertible, this allows us to estimate  $\mathbf{M}$  by multiplying both sides by its inverse  $(\mathbf{s}\mathbf{s}^T)^{-1}$ , giving

$$\hat{\mathbf{M}} = (\hat{\mathbf{s}}\mathbf{s}^T)^{-1} \quad (\text{B4})$$

The Amari error (Amari et al. 1996) can be computed by summing the rescaled cross-terms, which describes the degree of component mixing

$$E_{\text{Amari}} = \frac{\sum_{i=1}^N \left( \sum_{j=1}^N \frac{|\hat{M}_{ij}|}{\max_k |\hat{M}_{ik}|} - 1 \right) + \sum_{j=1}^N \left( \sum_{i=1}^N \frac{|\hat{M}_{ij}|}{\max_k |\hat{M}_{kj}|} - 1 \right)}{2(N^2 - N)} \quad (\text{B5})$$

Note that we have normalized the Amari error to have a maximum value of one rather than a number dependent on the number of sources.

## ACKNOWLEDGMENTS

We thank Drs. Ashesh Mehta and Istvan Ulbert for data collection, Dr. Peter Lakatos for helpful comments and discussion, and Dr. Leonard Trejo and S. Clanton of National Aeronautics and Space Administration Ames Research Center for assistance in optimizing the coding of portions of the algorithm. We also thank the anonymous reviewers for extensive input and valuable comments, which greatly helped to improve the quality of this work.

## GRANTS

K. H. Knuth was supported in part by a National Alliance for Research on Schizophrenia and Depression Young Investigator Award, the National Aeronautics and Space Administration (NASA) Intelligent Data Understanding/Intelligent Systems/Computing, Information, and Communications Technology Program, the NASA Software, Intelligent Systems, and Modeling Human and Robotic Technology Embedded Real-Time Advisory System for Crew-Automation Reliability, and the University at Albany (SUNY). A. S. Shah was supported by Medical Scientist Training Program Grant T32M07288 and by National Institute of Mental Health (NIMH) Grant MH-060358. C. E. Schroeder was supported by NIMH Grant MH-060358. Macaque data came from experiments supported by NIMH Grant MH-060358. S. L. Bressler and M. Ding were supported by National Science Foundation Grant IBN0090717 and NIMH Grants MH-64204 and MH-42900, and M. Ding was also supported by NIMH Grant MH-70498.

## REFERENCES

Amari S, Cichocki A, and Yang HH. A new learning algorithm for blind signal separation. In: *Advances in Neural Information Processing Systems*,

- edited by Touretzky D, Mozer M, and Hasselmo M. Cambridge, MA: MIT Press, 1996, vol. 8, p. 752–763.
- Bell AJ and Sejnowski TJ.** An information-maximization approach to blind source separation and deconvolution. *Neural Comp* 7: 1129–1159, 1995.
- Belouchrani A, Abed-Meraim K, Cardoso J-F, and Moulines E.** Second-order blind separation of correlated sources. *Proc Int Conf Digital Sig Proc Nicosia, Cyprus*, 1993, p. 346–351.
- Bogacz R, Yeung N, and Holroyd CB.** Detection of phase resetting in the electroencephalogram: an evaluation of methods. *Soc Neurosci Abstr* 28: 506.9, 2002.
- Bressler SL, Coppola R, and Nakamura R.** Episodic multiregional cortical coherence at multiple frequencies during visual task performance. *Nature* 366: 153–156, 1993.
- Cao J, Murata N, Amari S, Cichocki A, Takeda T, Endo H, and Harada N.** Single-trial magnetoencephalographic data decomposition and localization based on independent component analysis approach. *IEICE T Fund Electr* 9: 1757–1766, 2000.
- Delorme A and Makeig S.** EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J Neurosci Methods* 134: 9–21, 2004.
- Engel AK, Fries P, and Singer W.** Dynamic predictions: oscillations and synchrony in top-down processing. *Nat Rev Neurosci* 2: 704–716, 2001.
- Freeman JA and Nicholson C.** Experimental optimization of current source-density technique for anuran cerebellum. *J Neurophysiol* 38: 369–382, 1975.
- Givre SJ, Arezzo JC, and Schroeder CE.** Effects of wavelength on the timing and laminar distribution of illumination-evoked activity in macaque V1. *Vis Neurosci* 12: 229–239, 1995.
- Hyvärinen A and Oja E.** A fast fixed-point algorithm for independent component analysis. *Neural Comput* 9: 1483–1492, 1997.
- Jaskowski P and Verleger R.** Amplitude and latencies of single-trial ERP's estimated by a maximum-likelihood method. *IEEE Trans Biomed Eng* 46: 987–993, 1999.
- Jaynes ET.** *Probability Theory—The Logic of Science*. Cambridge, UK: Cambridge Univ. Press, 2003.
- Jung T-P, Makeig S, Westerfeld M, Townsend J, Courchesne E, and Sejnowski TJ.** Independent component analysis of single-trial event-related potentials. In: *Proceedings of the First International Workshop on Independent Component Analysis and Signal Separation: ICA'99*, edited by Cardoso J-F, Jutten Ch, and Loubaton Ph. Aussois, France: ICA'99, 1999, p. 173–178.
- Knuth KH.** Difficulties applying recent blind source separation techniques to EEG and MEG. In: *Maximum Entropy and Bayesian Methods, Boise 1997*, edited by Erickson GJ, Rychert JT, and Smith CR. Dordrecht, The Netherlands: Kluwer Academic, 1997, p. 209–222.
- Knuth KH.** Bayesian source separation and localization. In: *Proceedings of SPIE: Bayesian Inference for Inverse Problems*, edited by Mohammad-Djafari A. Bellingham, WA: SPIE, 1998, vol. 3459, p. 147–158.
- Knuth KH.** A Bayesian approach to source separation. In: *Proceedings of the First International Workshop on Independent Component Analysis and Signal Separation: ICA'99*, edited by Cardoso J-F, Jutten Ch, Loubaton Ph, Aussois France: ICA'99, 1999, p. 283–288.
- Knuth KH.** Informed source separation: a Bayesian tutorial (Invited paper). In: *Proceedings of the 13th European Signal Processing Conference (EU-SIPCO 2005)*, edited by Sankur B, Cetin E, Tekalp M, and Kuruoğlu E. Antalya, Turkey, 2005.
- Knuth KH, Truccolo WA, Bressler SL, and Ding M.** Separation of multiple evoked responses using differential amplitude and latency variability. In: *Proceedings of the Third International Workshop on Independent Component Analysis and Blind Signal Separation: ICA2001*, edited by Lee T-W, Jung T-P, Makeig S, and Sejnowski TJ. San Diego, CA: ICA2001, 2001, p. 463–468.
- Knuth KH and Vaughan HG Jr.** Convergent Bayesian formulations of blind source separation and electromagnetic source estimation. In: *Maximum Entropy and Bayesian Methods, Garching, Germany 1998*, edited by von der Linden W, Dose V, Fischer R, and Preuss R. Dordrecht, The Netherlands: Kluwer Academic, 1998, p. 217–226.
- Lange DH, Pratt H, and Inghar GF.** Modeling and estimation of single evoked brain potential components. *IEEE Trans Biomed Eng* 44: 791–799, 1997.
- Lee KH, Williams LM, Breakspear M, and Gordon E.** Synchronous gamma activity: a review and contribution to an integrative neuroscience model of schizophrenia. *Brain Res Brain Res Rev* 41: 57–78, 2003.
- Lee T-W, Girolami M, and Sejnowski TJ.** Independent component analysis using an extended infomax algorithm for mixed sub-Gaussian and super-Gaussian sources. *Neural Comput* 11: 609–633, 1999.
- Makeig S, Bell AJ, Jung T-P, and Sejnowski TJ.** Independent component analysis of electroencephalographic data. In: *Advances in Neural Information Processing Systems*, edited by Touretzky D, Mozer M, and Hasselmo M. Cambridge, MA: MIT Press, 1996, vol. 8, p. 145–151.
- Makeig S, Jung T-P, Bell A, Ghahremani D, and Sejnowski TJ.** Blind separation of auditory event-related brain responses in independent components. *Proc Natl Acad Sci USA* 94: 10979–10984, 1997.
- Makeig S, Westerfeld M, Jung T-P, Enghoff S, Townsend J, Courchesne E, and Sejnowski TJ.** Dynamic brain sources of visual evoked responses. *Science* 295: 690–694, 2002.
- Mehta AD, Ulbert I, and Schroeder CE.** Intermodal selective attention in monkeys. I. Distribution and timing of effects across visual areas. *Cereb Cortex* 10: 343–358, 2000.
- Mitzdorf U.** Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. *Physiol Rev* 65: 37–100, 1985.
- Mocks J, Gasser T, Kohler W, and De Weerd JP.** Does filtering and smoothing of average evoked potentials really pay? A statistical comparison. *Electroencephalogr Clin Neurophysiol* 64: 469–480, 1986.
- Nicholson C and Freeman JA.** Theory of current source-density analysis and determination of conductivity tensor for anuran cerebellum. *J Neurophysiol* 38: 356–368, 1975.
- Pham DT, Mocks J, Kohler W, and Gasser T.** Variable latencies of noisy signals: estimation and testing in brain potential data. *Biometrika* 74: 525–533, 1987.
- Särelä J, Vigário R, Jousmäki V, Hari R, and Oja E.** ICA for the extraction of auditory evoked fields. In: *4th International Conference on Functional Mapping of the Human Brain*. Montreal, Canada: HBM'98, 1998.
- Schmidt DM, George JS, and Wood CC.** Bayesian inference applied to the electromagnetic inverse problem. *Hum Brain Mapp* 7: 195–212, 1999.
- Schroeder CE.** Defining the neural bases of visual selective attention: conceptual and empirical issues. *Int J Neurosci* 80: 65–78, 1995.
- Schroeder CE, Mehta AD, and Foxe JJ.** Determinants and mechanisms of attentional modulation of neural processing. *Front Biosci* 6: D672–D684, 2001.
- Schroeder CE, Mehta AD, and Givre SJ.** A spatiotemporal profile of visual system activation revealed by current source density analysis in the awake macaque. *Cereb Cortex* 8: 575–592, 1998.
- Schroeder CE, Steinschneider M, Javitt DC, Tenke CE, Givre SJ, Mehta AD, Simpson GV, Arezzo JC, and Vaughan HG Jr.** Localization of ERP generators and identification of underlying neural processes. *Electroencephalogr Clin Neurophysiol Suppl* 44: 55–75, 1995.
- Schroeder CE, Tenke CE, Givre SJ, Arezzo JC, and Vaughan HG Jr.** Laminar analysis of bicuculline-induced epileptiform activity in area 17 of the awake macaque. *Brain Res* 515: 326–330, 1990.
- Schroeder CE, Tenke CE, Givre SJ, Arezzo JC, and Vaughan HG Jr.** Striate cortical contribution to the surface-recorded pattern-reversal VEP in the alert monkey. *Vision Res* 31: 1143–1157, 1991.
- Shah AS, Knuth KH, Lakatos P, and Schroeder CE.** Lessons from applying differentially variable component analysis (dVCA) to electroencephalographic activity. In: *Bayesian Inference and Maximum Entropy Methods in Science and Engineering, Jackson Hole WY 2003, AIP Conference Proceedings 707*, edited by Zhai Y and Erickson GJ. Melville, NY: American Institute of Physics, 2004, p. 167–181.
- Shah AS, Knuth KH, Truccolo WA, Ding M, Bressler SL, and CE.** A Bayesian approach to estimating coupling between neural components: evaluation of the multiple component event related potential (mcERP) algorithm. In: *Bayesian Inference and Maximum Entropy Methods in Science and Engineering, Moscow ID 2002, AIP Conference Proceedings 659*, edited by Williams C. Melville, NY: American Institute of Physics, 2003, p. 23–38.
- Singer W.** Synchronization of cortical activity and its putative role in information processing and learning. *Annu Rev Physiol* 55: 349–374, 1993.
- Sivia DS.** *Data Analysis. A Bayesian Tutorial*. Oxford, UK: Clarendon Press, 1996.
- Sivia DS.** Some thoughts on correlated noise. In: *Bayesian Inference and Maximum Entropy Methods in Science and Engineering, Jackson Hole WY 2003, AIP Conference Proceedings 707*, edited by Erickson GJ and Smith CR. Melville, NY: American Institute of Physics, 2003, p. 303–313.

- Student.** The probable error of a mean. *Biometrika* 6: 1–24, 1908.
- Tang AC, Pearlmuter BA, Malaszenko NA, and Phung DB.** Independent components of magnetoencephalography: single-trial response onset times. *Neuroimage* 17: 1773–1789, 2002.
- Tenke CE, Schroeder CE, Arezzo JC, and Vaughan HG Jr.** Interpretation of high-resolution current source density profiles: a simulation of sublaminal contributions to the visual evoked potential. *Exp Brain Res* 94: 183–192, 1993.
- Truccolo WA, Ding M, Knuth KH, Nakamura R, and Bressler SL.** Trial-to-trial variability of cortical evoked responses: implications for the analysis of functional connectivity. *Clin Neurophysiol* 113: 206–226, 2002.
- Truccolo WA, Knuth KH, Ding M, and Bressler SL.** Bayesian estimation of amplitude, latency and waveform of single trial cortical evoked components. In: *Bayesian Inference and Maximum Entropy Methods in Science and Engineering, Baltimore MD 2001, AIP Conference Proceedings 617*, edited by Fry RL and Bierbaum M. Melville, NY: American Institute of Physics, 2001, p. 64–73.
- Truccolo WA, Knuth KH, Shah AS, Bressler SL, Schroeder C, and Ding M.** Estimation of single-trial multicomponent ERPs: differentially variable component analysis (dVCA). *Biol Cybern* 89: 426–438, 2003.
- Vigário R, Särelä J, Jousmäki V, and Oja E.** Independent component analysis in decomposition of auditory and somatosensory evoked fields. In: *Proceedings of the First International Workshop on Independent Component Analysis and Signal Separation: ICA'99*, edited by Cardoso J-F, Jutten Ch, and Loubaton Ph. Aussois, France: ICA'99, 1999, p. 167–172.
- Woody CD.** Characterization of an adaptive filter for the analysis of variable latency neuroelectric signals. *Med Biol Eng* 5: 539–553, 1967.