Applying Differentially Variable Component Analysis (dVCA) to Event-related Potentials

Ankoor S. Shah[*†], Kevin H. Knuth[**], Peter Lakatos[†] and Charles E. Schroeder[†*]

[*]*Department of Neuroscience, Albert Einstein College of Medicine, Bronx, NY 10461*
[†]*Cognitive Neuroscience & Schizophrenia, Nathan Kline Institute, Orangeburg, NY 10962*
[**]*Computational Sciences Division, NASA Ames Research Center, Moffett Field, CA 94035*

Abstract. Event-related potentials (ERPs) generated in response to multiple presentations of the same sensory stimulus vary from trial to trial. Accumulating evidence suggests that this variability relates to a similar trial-to-trial variation in the perception of the stimulus. In order to understand this variability, we previously developed differentially Variable Component Analysis (dVCA) as a method for defining dynamical components that contribute to the ERP. The underlying model asserted that: (i) multiple components comprise the ERP; (ii) these components vary in amplitude and latency from trial to trial; and (iii) these components may co-vary. A Bayesian framework was used to derive *maximum a posteriori* solutions to estimate these components and their latency and amplitude variability. Our original goal in developing dVCA was to produce a method for automated estimation of components in ERPs. However, we discovered that it is better to apply the algorithm in stages because of the complexity of the ERP and to use the results to define interesting subsets of the data, which are further analyzed independently. This paper describes this method and illustrates its application to actual neural signals recorded in response to a visual stimulus. Interestingly, dVCA of these data suggests two distinct response modes (or states) with differing components and variability. Furthermore, analyses of residual signals obtained by subtracting the estimated components from the actual data illustrate gamma-frequency (circa 40 Hz) oscillations, which may underlie communication between various brain regions. These findings demonstrate the power of dVCA and underscore the necessity to apply this algorithm in a guided rather than a ballistic fashion. Furthermore, they highlight the need to examine the residual signals for those features of the signals that were not anticipated and not modeled in the derivation of the algorithm.

INTRODUCTION

Imagine seeing something move in the distance while on safari in Africa. As you approach the object you notice that it is yellow, has a tail, four legs, and a mane. Instantly, you piece together these different facts and "perceive" a lion. Our brain performs this operation repeatedly as our senses continuously sample the external environment looking for stimuli that are relevant to the task. However, the mechanisms underlying perception are still hypothetical. Understanding these methods of perception has been an interesting question for hundreds of years, and it is becoming more important in understanding the dysfunction of various neurological and psychiatric illnesses such as schizophrenia [1].

Electroencephalography (EEG) is often used to study ensembles of neurons responding to sensory stimuli with the assumption that it will yield insight into neural encoding and perception. Typically, experimentalists present the same sensory stimulus to a subject multiple times while continuously recording EEG. The electric potential generated

in response to a single presentation of the sensory stimulus is a single-trial, event-related potential (ERP), where the term trial is synonymous with stimulus presentation. Single-trial ERPs are highly variable [2], and they are typically averaged across trials to reduce noise. However, if trial-by-trial variability represents dynamical changes in the state of the system instead of noise, averaging would obscure information on such processes. Moreover, single-trial ERPs capture electrical activity from multiple generators affected by stimulus presentation, and averaging yields no information about these generators or their interaction. Characterizing the variability and the multiple generators of neural responses may further our understanding of the mechanisms of perception.

We previously introduced differentially Variable Component Analysis (dVCA) [3, 4, 5, 6, 7] to analyze the variability of single-trial ERPs and their underlying generators. This technique assumes that multiple components or stereotypical waveshapes compose the single-trial ERP and that each component may vary differentially from trial to trial. The dVCA algorithm estimates the components and their single-trial variability, and this information can be utilized to examine how different aspects of the single-trial ERP may be functioning during the neural response to stimulus presentation. When we first introduced this algorithm, our goal was to create an automated tool. Although this worked well in simulations [5, 7], we discovered it was far better to apply a guided approach when using dVCA on actual data. In this paper, we outline the guided approach by demonstrating it on data recorded in the primary visual cortex V1 of a macaque monkey. We will highlight some of the major capabilities of this technique, but we will not provide an in-depth analysis of these data.

THE MODEL

The multiple component ERP (mcERP) model, which underlies the dVCA algorithm, describes the neural activity associated with the presentation of the stimulus as the sum of all components activated in response to that stimulus. Components are defined as stereotypical waveshapes that describe the temporal activation pattern of neural ensembles. Since neural responses display trial-to-trial variability [2], the mcERP model states that each component may vary in time and amplitude across trials. Mathematically, we write:

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

where $x_{mr}(t)$ represents the recorded signal in electrode m during stimulus trial r over time t and N denotes the total number of components $s_n(t)$. The first term of the summation on the right-hand side of Equation 1 can be thought of as the stimulus-evoked activity, and the second term is the unpredictable portion. The evoked portion refers to activity that is relatively time-locked and phase-locked to stimulus presentation. Since evoked activity is thought to be repetitive across trials, each component n has a stereotypical waveshape denoted as $s_n(t)$. To account for trial-to-trial variability in the evoked activity of any particular trial r, each component waveshape is modified by a specific latency shift τ_{nr} and a specific amplitude scale α_{nr}. The latency shift allows a component to vary in time across trials, while the amplitude scale allows the

component to vary in amplitude across trials. The coupling coefficient C_{mn} describes how a particular electrode m relates to each component n. This term is most important in cases where multiple electrodes record neural activity simultaneously. In these instances, more than one electrode may record a single component differentially because the electric potential will generate a measurable, but varying, field across space. In contrast to the evoked activity, the unpredictable signal in Equation 1 contains activity that is (i) neither time- nor phase-locked to the stimulus presentation, (ii) biological noise, and (iii) recording apparatus noise.

Several assumptions underlie the mcERP model. First, each presentation of the stimulus generates electrical activity from multiple neural ensembles that sum together to form the ERP. Second, each component displays a stereotypic waveshape that varies only in amplitude and latency from presentation to presentation. Though this assumption allows for variability in the waveshape of summed activity through differential variability of each component, it does not take into account any potential changes in the waveshape of a single component. It does allow for a particular component to disappear in trial r because the single-trial amplitude scale α can be set to zero. In this way, two different waveshapes that are alternately expressed by the neural populations would be considered as two distinct components.

In order to find the most probable values for the parameters of the evoked activity in the mcERP model, we utilized Bayes' Theorem to express the posterior probability of our model:

$$p(model|data,I) = \frac{p(data|model,I)p(model|I)}{p(data|I)}, \tag{2}$$

where I represents prior information about the physical situation. For simplicity, we assigned uniform densities for the prior probabilities of each model parameter, and we assigned a Gaussian likelihood based on the mcERP model in Equation 1. This gave

$$p(\mathbf{C},\mathbf{s}(t),\alpha,\tau,\sigma|\mathbf{x}(t),I) \propto (2\pi\sigma^2)^{\frac{-MRT}{2}} Exp[-\frac{1}{2\sigma^2}Q]p(\sigma|I), \tag{3}$$

where

$$Q = \sum_{m=1}^{M}\sum_{r=1}^{R}\sum_{t=1}^{T}(x_{mr}(t) - \sum_{n=1}^{N} C_{mn}\alpha_{nr}s_n(t-\tau_{nr}))^2 \tag{4}$$

and where M denotes the total number of recording electrodes, R is the total number of stimulus presentations, and T denotes the total number of time points. Assigning a Jeffreys prior for $p(\sigma|I)$ and marginalizing over all possible values of the uncertainty σ, we obtained an expression for the posterior probability as

$$p(\mathbf{C},\mathbf{s}(t),\alpha,\tau|\mathbf{x}(t),I) \propto Q^{\frac{-MRT}{2}}, \tag{5}$$

and taking the logarithm of both sides, we got

$$\ln P = -\frac{MRT}{2}\ln Q + const, \tag{6}$$

where P is the posterior probability. From this expression, we calculated *maximum a posteriori* (MAP) equations for each of the model parameters by taking the partial

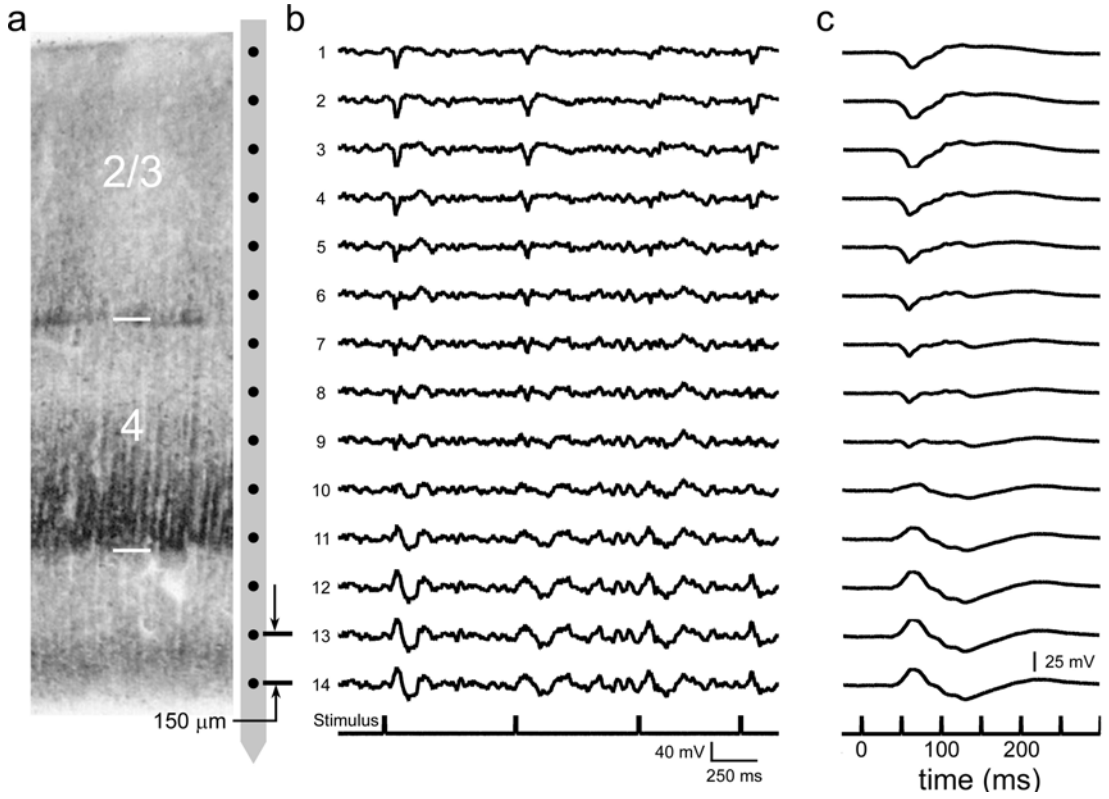

FIGURE 1. Example of linear-array multielectrode readings taken from V1 of a macaque monkey. (a) Histological section of primary visual cortex of a macaque monkey is shown on the left with the layers clearly marked. A schematic of the linear-array multielectrode is shown to the right with each "black dot" representing a separate recording channel. (b) Field potential recordings from each channel of the multielectrode. The last trace shows ticks denoting individual presentations of the red-light stimulus. (c) Averaged field potential profile across trials for the standard visual stimulus.

derivative of Equation 6 with respect to each of the parameters and setting them equal to zero. A more detailed description of the model can be found in Shah *et al.* (2003), while a complete derivation is found in Knuth *et al.* (Submitted) [5, 7].

DATA COLLECTION

Before discussing the algorithm that we utilized to solve the MAP equations, it is worth describing the data collection methods that we used in sampling neural field potentials from multiple electrodes simultaneously. Data were collected from the primary visual cortex (V1) of an awake macaque monkey by acutely inserting a linear-array electrode into the brain (Figure 1a). These data were collected in conjunction with a previous study, but they were never analyzed [8, 9]. Prior to recording the EEG data, the subject had been prepared surgically for these types of chronic recordings (see [10] for details of the surgical preparation). The linear-array multielectrode was inserted such that the individual electrodes or data channels spanned the different layers of cortex. The mul-

tielectrode contained 14 independent recording channels that were spaced at 150 μm increments. Field potential activity was sampled continuously at 2000 Hz (Figure 1b) and stored on a PC-based data acquisition system (Neuroscan Inc., El Paso, Texas). During the sampling of electrical activity, the subject was presented with randomly interspersed standard and target visual stimuli at an average rate of two stimuli per second. The standard visual stimulus consisted of a 10-μs, red-light flash, while the target stimulus varied slightly in intensity. Stimulus presentation began when the monkey depressed a lever and looked within a central fixation window thereby ensuring that he was attentive to the task. The subject's task was to release a lever when the target stimulus was presented. Correct responses were rewarded with a drop of juice, while incorrect responses generated a "timeout period" during which no stimuli were presented and thus no reward could be earned.

The data were prepared offline for dVCA by cutting the continuous data record into 325-millisecond epochs. The epochs spanned -25 to 300 ms with 0 denoting onset of the visual stimulus. One-hundred-and-thirty-seven presentations of the stimulus were found within the data record. Typically, these 137 epochs are averaged across trials for each channel, and the result in each channel is called the average ERP (Figure 1c). Since we are most often interested in the local activity generated by stimulus presentation and since field potential recordings often contain activity that has volume conducted

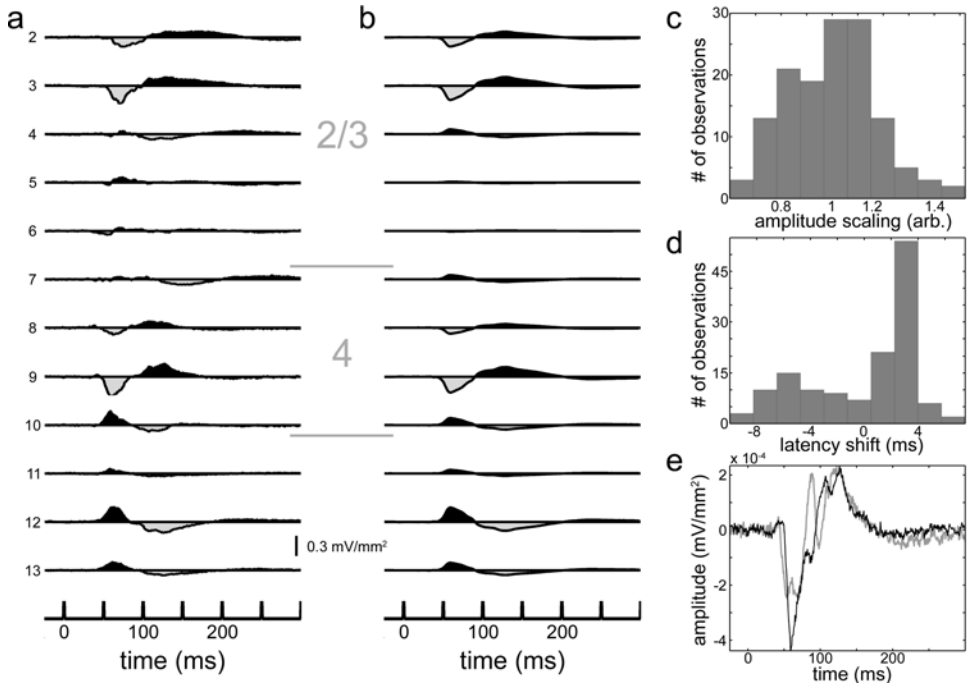

FIGURE 2. One-component model of data recorded from primary visual cortex of a macaque monkey. (a) Averaged current source density (CSD) profile of the actual data with the cortical layers marked on the right. (b) CSD profile of the 1-component model after dVCA has been applied to the full data set. The scale bar applies to panels (a) and (b). (c) and (d) Distributions of the single-trial amplitude scales and latency shifts for the 1-component model of the full data set. (e) CSD signal in channel 9 of the multielectrode after selectively averaging the actual data. The full data set was split into "early" (gray) and "late" (black) subsets based on cutoff of -2 ms in the latency shifts shown in (d).

from distant sites, we visualize the results by applying a current source density (CSD) transformation to the data (see Figure 2a). The laminar CSD profile is computed by taking the second spatial derivative of the field potential recordings [11, 12, 13]. This sacrifices the first and last recording channels, but the resultant indexes transmembrane current flow in the local neuronal ensemble. Current sinks (shaded gray in Figure 2a) represent current flow patterns from the extracellular medium into the intracellular compartment of neurons while current sources (shaded black in Figure 2a) denote the opposite. As with any generator of electrical current, the sink-source patterns must balance each other if the linear-array multielectrode has "bracketed" the generators properly (for example, see channels 8-11 in Figure 2a). This structure of balanced sink-source combinations indicates local neural activity.

In applying dVCA, we will always apply the algorithm to field potential recordings because computing a second derivative (approximated as a double finite difference) adds significant noise to the signal. However, since we are interested in determining the generators of the field potential activity, we show the current source density profiles of each estimated component by performing the CSD calculation on the product of the coupling matrix and the component waveshape.

THE ALGORITHM

The MAP equations allow us to estimate the most probable set of parameters for the mcERP model given our data; however, there are several methods to solve this system of equations. The method that we apply is a guided, fixed-point algorithm where each parameter is estimated sequentially and updated through successive iterations. It is similar to our previous automated algorithm [5, 7], but there are several decision points where the user must direct the algorithm. The algorithm begins by modeling a single component in the data (Figure 3). The average ERP for each multielectrode channel is calculated, and the total power for each average waveform is determined. The waveform with the largest signal power is the initial guess at the component waveshape $s_1(t)$, where the 1 denotes the first component. The single-trial amplitude scales α and the single-trial latency shifts τ are set to 1 and 0, respectively. We will use the boldfaced font to denote the entire set of values (the values across all 137 trials) for that parameter. The next step is to estimate the coupling coefficients using the MAP equation for $\mathbf{C}$ and our initial guesses for $s_1(t)$, α, and τ. Then, the single-trial latencies τ are updated given the newly estimated coupling coefficients $\mathbf{C}$ and the initial guesses for $s_1(t)$ and α. Similarly, the single-trial amplitude scales α are updated using the newly acquired values for τ, $\mathbf{C}$, and the initial guess for $s_1(t)$. Finally, the waveshape of $s_1(t)$ is updated with the newly estimated parameters $\mathbf{C}$, τ, and α. At this stage, we can begin the process again by re-estimating the coupling coefficients $\mathbf{C}$. Through successive iterations, the parameters are updated and likely come closer to the most probable set of parameters. This iterative process is halted if (i) the change in the waveshape of the component between successive iterations is less than one percent or if (ii) the maximum number of iterations is reached. The maximum number of iterations is set by the user, and this is the first decision point in the algorithm. We chose 50 iterations for this study because the posterior probability

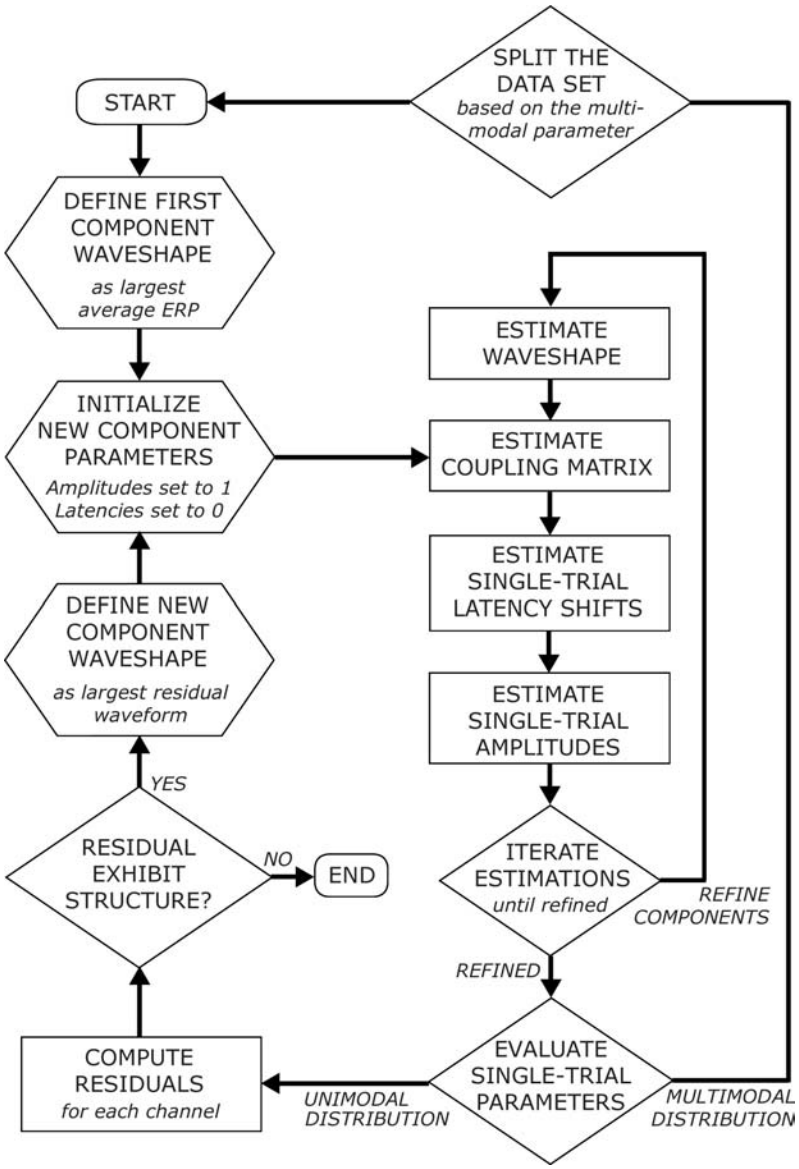

FIGURE 3. Flow chart of the dVCA algorithm.

of the model typically reaches a maximal plateau after approximately 18 to 20 iterations in our experience.

After refining the 1-component model, the estimated parameters must be examined. Specifically, the single-trial amplitude scales α and the single-trial latency shifts τ must be evaluated to determine whether the data contain multiple activation states. This determination is the second decision point in the guided application of dVCA. As will be demonstrated below, the distribution of these parameters may be multimodal thereby suggesting more than one activation state. If either α or τ shows multiple modes, the data set is split according to these modes, and the application of dVCA is started over again on each subset independently. By splitting the data set, we limit the number of

different components in the data making it less likely that the algorithm will get stuck at local maxima in probability space.

The third decision point determines whether a second component will be added to the model of the ERP, which comes once the first component has been estimated and the single-trial parameters have been evaluated. This decision is based on an evaluation of the residuals. Single-trial estimates of the ERP (based on the current model) are calculated by substituting the parameters of dVCA into the first term of the summation on the right-hand side of Equation 1. These estimates are then subtracted from each trial of the actual field potential data to yield single-trial, residual signals that contain activity not isolated by the 1-component model. These residual signals are averaged across trials to enhance activity that is relatively time- and phase-locked to stimulus presentation at the expense of the unpredictable signal. The CSD profile of this average residual signal is calculated to determine whether the residual contains current sources or sinks that resemble physiological generators. When the average residual CSD profile meets this subjective criterion, a second component is estimated.

Additional components in the mcERP model are estimated in a similar fashion to the first component. The first guess at the waveshape of any additional component is the average residual signal with the largest power. As the additional component is estimated, the parameters of all previous components are also updated. In this way, the components of the model change dynamically as each individual component better describes the physical situation. Refinement of the components, evaluation of potential subsets, and examination of the residuals continues until the CSD profile resembles unstructured noise. In this way, the dVCA algorithm allows us to model the most salient features of the average ERP. Activity such as that contained within the unpredictable signal is left behind in the residual, which can be studied after the modeled components are subtracted out of the actual data.

APPLYING DVCA TO ACTUAL DATA

The dVCA algorithm was applied to the 137 single-trial ERPs recorded in the 14 channels of the multielectrode. A single-component model was estimated for these data, and its CSD profile is illustrated (Figure 2b). This profile was calculated by multiplying the coupling coefficients **C** with the estimated component waveshape $s_1(t)$ and then calculating the CSD profile. Not surprisingly, the 1-component model captures much of the activity shown in the average profile of the actual data, which is displayed in Figure 2a. As outlined in Figure 3, the next step in evaluating the dVCA results is to examine the estimated amplitude scales and latency shifts for this component. The histogram of amplitude scales (Figure 2c) illustrates that the variability in this component's amplitude is quite small ($\sigma_{amplitude} = 0.173$). The histogram of the latency shifts (Figure 2d) shows a bimodal distribution with peaks at -5.63 and +3.13 ms. Based on the guidelines presented in Figure 3, we split the data set in two and generated an "early" and "late" subset. The early subset consisted of 49 trials in which the latency shift was less than or equal to -2 ms, while the late subset consisted of the remaining 88 trials. The cutoff of -2 ms was chosen because it falls between the two peaks of the estimated latency shifts (see

Figure 2d). Although the dVCA results suggest two types of neural responses to the same stimulus, we examined the actual data to verify these findings by selectively averaging them according to the early and late designations. This confirmed that the onset latency of the CSD signal in the early subset preceded that in the late subset (Figure 2e).

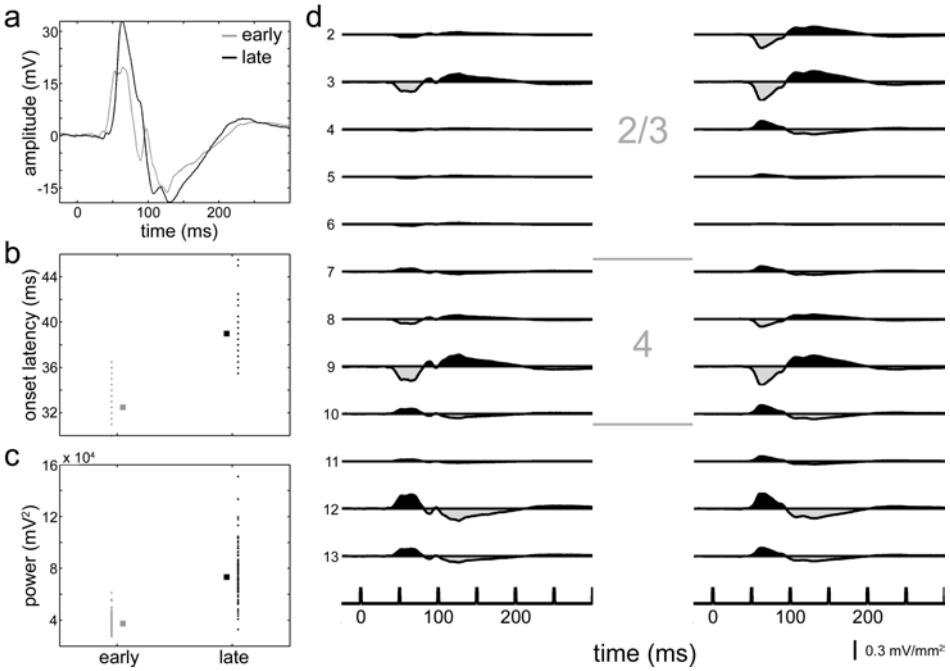

FIGURE 4. Applying dVCA to the early and late subsets individually. (a) Component waveshapes $s_1(t)$ for each subset. (b) Distribution of the onset latencies across trials in the model of the early and late subsets. The small dots mark individual trials, and the large squares mark the medians of each population. The medians of the two populations are statistically different from each other (Wilcoxon Rank-Sum test, $p < 0.001$). Since the onset latency is a discrete measure, individual measures are overlapping in this plot. (c) Distribution of the signal power across trials in the model of the early and late subsets. The small dots mark each trial, and the large squares mark the population median. The two populations are statistically different from each other (Wilcoxon Rank-Sum test, $p < 0.001$). (d) The CSD profiles of the 1-component model in the early (left) and late (right) subsets.

As specified by the dVCA algorithm (Figure 3), the analysis technique was applied to each subset separately. The modeling results on each subset show that the component waveshape for the early subset has an onset latency of 33 ms while that of the late subset has an onset of 39 ms (Figure 4a). These onset latencies were determined by finding the first time point where the amplitude of the signal exceeded the mean of the baseline (defined as -25 to 5 ms) by at least two standard deviations for 4 ms. By comparing the onset latency distributions of the model in each subset (Figure 4b), we verified a significant difference between the two populations using a non-parametric test comparing the medians to two samples (Wilcoxon Rank-Sum Test [14], $p < 0.001$). Second, the neural response in the early subset is smaller than that in late subset, and we verified this finding statistically (Wilcoxon Rank-Sum test, $p < 0.001$) by comparing the signal power of the model in each subset (Figure 4c). Although the 1-component model of each subset differs in onset latency and power, the component waveshape only represents the temporal activation pattern of the field potential generated by stimulus

presentation. It contains no information about the relative polarity and weighting with respect to any particular electrode channel. This information is contained in the coupling coefficients, which can help assess the spatial location of a component. The CSD profiles of the 1-component models for each subset show similar sinks and sources in and below layer 4 (Figure 4d). Activity in layers 2/3 differs especially in channel 4 where the early subset has little activity while the late subset has a large, initial source.

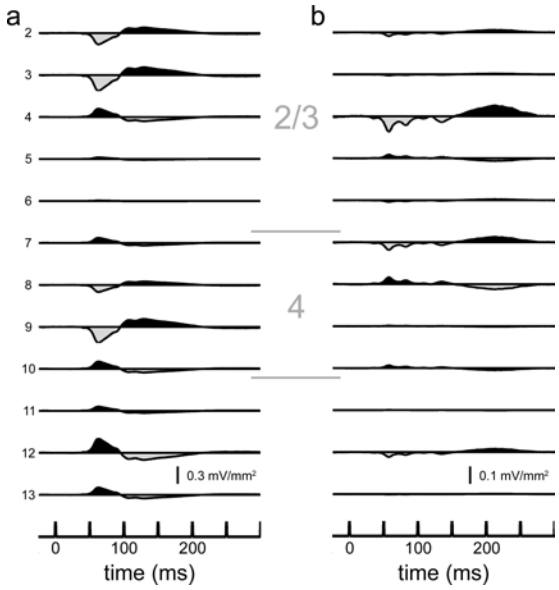

FIGURE 5. Two-component model of the late subset. The CSD profile of the first component $s_1(t)$ (a) and second component $s_2(t)$ (b) are illustrated. Note that the vertical scales in (a) and (b) are different.

For the purposes of this paper, we will only concentrate on the late subset from this point forward. Since the 1-component model of the late subset has been estimated and refined, we can now examine the residual signals. The residual activity (not shown) in channels 4-6 resembled activity in the actual data set (see Figure 2a, channels 4-6), and in accordance with the dVCA algorithm (Figure 3), we added a component and estimated the 2-component model for the late subset. The first component (Figure 5a) looks similar to the 1-component model and its onset latency is 38 ms, and the second component (Figure 5b) is much smaller, localizes to layers 2/3 and upper layer 4, and has an onset latency of 44 ms. This component also captures some interesting 30-40 Hz oscillations between 44 and 150 ms, and since they appear within a component of the mcERP model, they are likely relatively time- and phase-locked to stimulus presentation. Figure 6 displays the single-trial parameters for each component, and it shows that the amplitude variability of the first component ($\sigma_{\alpha 1} = 0.136$) is smaller than that of the second component ($\sigma_{\alpha 2} = 0.392$) (Figure 6a and b). However, there is a significant positive correlation between the two values ($r = 0.38, p < 0.001$) (Figure 6c). The latency-shift variability of component 1 ($\sigma_{\tau 1} = 3.28ms$) is much less than that for component 2 ($\sigma_{\tau 2} = 18.11ms$) (Figure 6d and e), and there is no relation between the latency-shift estimates even if the extreme points of component 2 are excluded. Thus, the 2-component model of the late subset indicates that the first component is activated prior to the second component, that it is less variable than the second component across

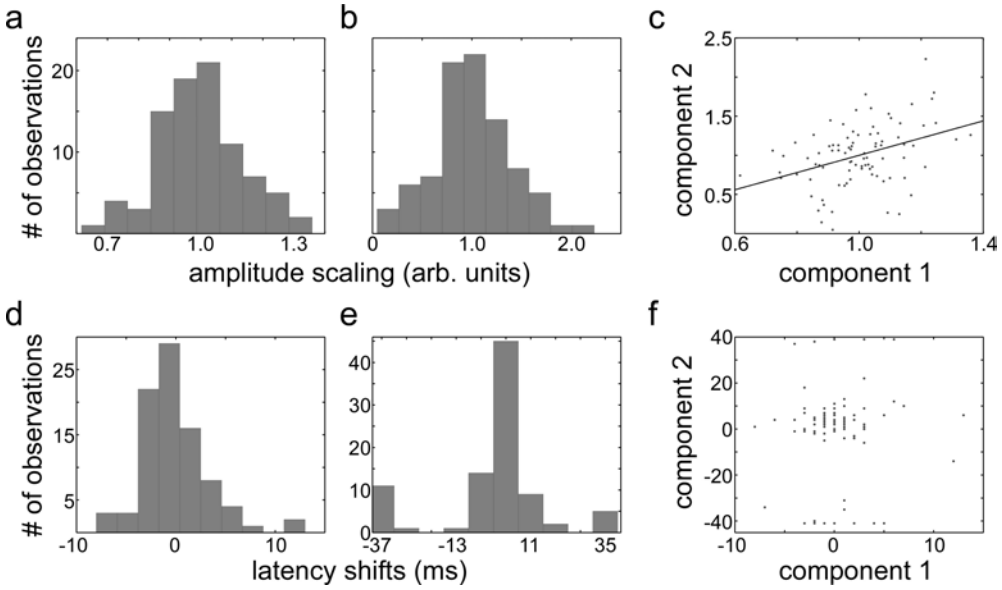

FIGURE 6. Single-trial amplitude and latency shifts of the 2-component model of late subset. Distributions of the single-trial amplitude scales for the first (a) and second (b) components co-vary across trials (c). Distributions of the single-trial latencies for the first (d) and second (e) components do not illustrate any correlation (f).

trials, and that the amplitude of the first and second component co-vary across trials.

As with any type of modeling, it is imperative to examine the residual activity to evaluate and understand the signals that were not modeled. By modeling the evoked activity, the residual signal should approach the unpredictable signal in Equation 1. This signal contains activity that is not phase-locked to stimulus presentation. This type of activity is of increasing importance in the field of neuroscience because mounting evidence suggests that gamma-frequency oscillations (> 30 Hz), which are non-phase-locked, may be involved in facilitating communication between different brain areas [15] and in amplifying behaviorally relevant stimuli [16]. To study the non-phase locked activity of the late subset, the residual signals were calculated by subtracting the 2-component model in each individual trial from the late subset, actual data in the same trial. Single-trial CSD profiles were calculated, and a Morlet-based wavelet was convolved with these signals to obtain a time-frequency (TF) map of the power. Single-trial maps were averaged across trials in each multielectrode channel, and the resultants were evaluated for gamma-frequency (> 30 Hz) oscillations. The TF maps for each of the layers in cortex showed differing time and frequency characteristics (Figure 7). Channel 2 of the multielectrode, located in layers 2/3, has activity in the traditional gamma-band range of 30-50 Hz and in the very high frequency (VHF) gamma-band range of 52-76 Hz (Figure 7a). The gamma-band activity peaks at 71.5 ms and 38 Hz, while the VHF activity peaks 89.5 ms and 65 Hz. Layer 4 activity from channel 9, on the other hand, shows a peak at 65 ms and 35 Hz (Figure 7b). Finally, the lower layers as represented by channel 12 have a peak at 80 ms and 36 Hz (Figure 7c). These results demonstrate that the residual signals in the late subset contain un-modeled gamma-frequency activity that is

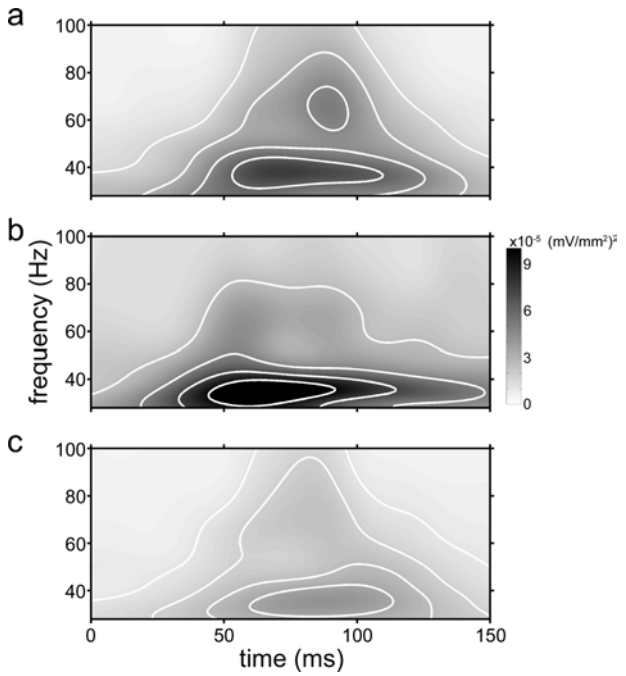

FIGURE 7. Time-frequency (TF) maps of the residual signal in the late subset. Examples from layers 2/3 (Channel 2, panel a), layer 4 (Channel 9, panel b), and below layer 4 (Channel 12, panel c) show that the residual signal varies in timing and frequency content across layers. All three TF maps are shown on the same gray scale, which is depicted in (b).

presumably not phase-locked to stimulus presentation. The characteristics of this activity may vary according to cortical layer.

DISCUSSION

Our two goals in this paper were to outline the guided approach of applying dVCA on real data and to illustrate some of the capabilities of this technique. Although an automated data analysis technique would have been convenient, applying dVCA in a guided manor is necessary considering the unexpected complexities of the brain responses. As mentioned above, there are three decision points in the guided algorithm. The first decision point occurs in choosing an appropriate value for the maximum number of iterations during the refinement of componentry. To avoid subjectivity in this assignment, the number must be assigned such that the posterior probability across iterations reaches a maximal plateau. In the cases where the component changes less than one percent in successive iterations, the maximal number of iterations does not come into play, but the posterior probability should still be evaluated to ensure that the parameter set is the most probable. The second decision point comes in evaluating the single-trial amplitude scales and latency shifts after refinement. In the case demonstrated above, it is quite easy to determine that there are multiple modes in the latency shift distribution. Other cases may not be so obvious, but if a data set were split when it should

not have been, one would expect very similar components and single-trial parameters for each of the subsets. The third decision point comes in determining whether to estimate additional components in modeling any particular data set. By comparing the residual signal with our expectations of physiologically realistic CSD profiles, we impose our ideas of what constitutes a valid neural generator. Certainly we risk missing interesting results due to our ignorance, but by examining the residual signals carefully, we provide ourselves the opportunity of recognizing new types of previously unimagined signals. The second goal of this work was to highlight some of the major capabilities of dVCA, and we show that it is capable of identifying the evoked activity and its variability in individual neural responses. By correlating these results with system level changes such as attention and arousal or with behavioral measures such as reaction time [6], one could study how dynamics in the evoked activity may contribute to higher cognitive functions such as perception. Moreover, this information may be utilized to evaluate how different components drive other components in the same brain area or across brain areas.

The modeling of the ERP responses in these data showed that the primary visual cortex displays at least two differing states of activation in response to a single stimulus. These two states differ markedly in waveshape, onset latency, overall power, and in their layers 2/3 activation (see Figure 5). Fries and colleagues [17] have suggested that shifts in the onset latency can be related to the phase of pre-existing gamma-band oscillations. Furthermore, Mehta *et al.* have shown that attention can alter the neural response profile in several cortical areas including the primary visual cortex [8], and either may explain the two states of activation. Additionally, minute variations in eye position within the fixation window may also account for the differing profiles. The phase of pre-existing gamma-band oscillations and the position of the eye are two hypotheses that could be tested in the present data. The attention hypothesis, however, would have to be tested with further experimentation by comparing the response to the visual stimulus when the subject attends to it versus when the subject ignores it (see [8] for a description of an attention experiment that would be amenable for this type of testing). Regardless of the mechanism, the identification of multiple response states by dVCA permits us to explore other variables that may affect sensory processing.

The residual activities characterized in three different locations within our multielectrode array reveal information about the un-modeled components of the visual response. There is structure to this activity, but it does not appear in the average of the residuals (results not shown) likely because it is neither time- nor phase-locked to stimulus onset. Applying a time-frequency power analysis allows us to eliminate phase information and concentrate on where and when gamma-frequency bursts occur. These non-phase-locked oscillations are of great interest to the neuroscience community because they are hypothesized to underlie our perception of objects by dynamically synchronizing different cortical areas during stimulation [18]. Interestingly, our results show that this activity in layer 4 of primary visual cortex precedes the same type of activity above and below layer 4. This temporal pattern is associated with information ascending through the system [19, 20]. Finally, layers 2/3 show a very-high-frequency, gamma-band oscillation that is not found in the other layers. The significance of this burst along with the other gamma-frequency bursts is yet to be determined.

The methods and data presented here show that the dVCA algorithm can characterize variability in ERPs and identify interesting subsets of data. The variability in neural

responses can be studied and correlating these changes with behavior may further our understanding of perception. Additionally, characterizing activity in patients suffering from neurological and psychiatric illnesses would enable a study of how neural processing varies in disease states. In combination with any study of the model parameters, a study of the residual activity must also be performed because these non-modeled signals may be critical in understanding the system. It is hoped that the new capabilities that dVCA provides will further our understanding of the neural mechanisms underlying perception.

ACKNOWLEDGMENTS

We thank Tammy McGinnis, Noelle O'Connell, Aimee Mills, and Chi-Ming Chen for their support and technical help. We also acknowledge the Medical Scientist Training Program at the Albert Einstein College of Medicine (T32M07288), the National Institute on Mental Health (MH60358), the NASA IDU/IS/CICT Program, and the NASA Aerospace Technology Enterprise for their support of this work. Data in this paper are from a thesis to be submitted in partial fulfilment of the requirements for the Degree of Doctor of Philosophy in the Graduate Division of Medical Sciences, Albert Einstein College of Medicine.

REFERENCES

1. Lee, K. H., Williams, L. M., Breakspear, M., and Gordon, E., *Brain Res Brain Res Rev*, **41**, 57–78 (2003).
2. Truccolo, W. A., Ding, M., Knuth, K. H., Nakamura, R., and Bressler, S. L., *Clin Neurophysiol*, **113**, 206–226 (2002).
3. Knuth, K. H., Truccolo, W. A., Bressler, S. L., 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, Dec. 9-12, 2001, San Diego CA*, edited by T.-W. Lee, T.-P. Jung, S. Makeig, and T. J. Sejnowski, 2001, pp. 463–468.
4. Truccolo, W. A., Knuth, K. H., Ding, M., and Bressler, S. L., "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*, edited by R. L. Fry and M. Bierbaum, AIP Conference Proceedings 617, American Institute of Physics, New York, 2002, pp. 64–73.
5. Shah, A. S., Knuth, K. H., Truccolo, W. A., Ding, M., Bressler, S. L., and Schroeder, C. E., "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*, edited by C. Williams, AIP Conference Proceedings 659, American Institute of Physics, New York, 2003, pp. 23–38.
6. Truccolo, W. A., Knuth, K. H., Shah, A. S., Bressler, S. L., Schroeder, C. E., and Ding, M., *Biol Cybern* (2003).
7. Knuth, K. H., Shah, A. S., Truccolo, W. A., Bressler, S. L., Ding, M., and Schroeder, C. E. (submitted).
8. Mehta, A. D., Ulbert, I., and Schroeder, C. E., *Cereb Cortex*, **10**, 343–358 (2000).
9. Mehta, A. D., Ulbert, I., and Schroeder, C. E., *Cereb Cortex*, **10**, 359–370 (2000).
10. Schroeder, C. E., Mehta, A. D., and Givre, S. J., *Cereb Cortex*, **8**, 575–592 (1998).
11. Freeman, J., and Nicholson, C., *J Neurophysiol*, **38**, 369–382 (1975).

12. Mitzdorf, U., *Physiol Reviews*, **65**, 37–100 (1985).
13. Schroeder, C. E., Steinschneider, M., Javitt, D. C., Tenke, C. E., Givre, S. J., Mehta, A. D., Simpson, G. V., Arezzo, J. C., and Vaughan, H. G., Jr., *Electroenceph Clin Neurophysiol*, **44**, 55–75 (1995).
14. Gibbons, J. D., *Nonparametric statistical inference*, M. Dekker, New York, 1985, ISBN 0824773276.
15. Gray, C. M., Konig, P., Engel, A. K., and Singer, W., *Nature*, **338**, 334–337 (1989).
16. Fries, P., Reynolds, J. H., Rorie, A. E., and Desimone, R., *Science*, **291**, 1560–1563 (2001).
17. Fries, P., Neuenschwander, S., Engel, A. K., Goebel, R., and Singer, W., *Nat Neurosci*, **4**, 194–200 (2001).
18. Engel, A. K., Fries, P., and W., S., *Nat Rev Neurosci*, **2**, 704–716 (2001).
19. Felleman, D. J., and Van Essen, D. C., *Cereb Cortex*, **1**, 1–47 (1991).
20. Schroeder, C. E., Shah, A. S., Mehta, A. D., Givre, S. J., and Ulbert, I. (in prep).