10

Linking Neuronal Variability to Perceptual Decision Making via Neuroimaging

Paul Sajda, Marios G. Philiastides, Hauke Heekeren, and Roger Ratcliff

Introduction

One of the fundamental questions posed by systems neuroscience is whether the variability observed in neuronal responses is largely reflective of neurons being noisy processing elements, is a result of unaccounted contextual effects of otherwise identical stimuli (e.g., a memory/hysteresis effect) and/or is reflective of latent processes in the underlying neuronal network. From a behavioral neuroscience perspective, decision making is also confronted with the issue of variability—namely, that even for very simple decisions, accuracy and response time (RT) can vary significantly for nominally identical stimuli. Over the last decade, there has been substantial work focused on linking neuronal variability to this behavioral variability. For the most part, much of the effort has been focused on animal studies, including nonhuman primates. Recent advances in neuroimaging, however, specifically methods for single-trial analysis of noninvasively measured neural activity, has enabled one to address the question with respect to variability and decision making in the human brain.

In this chapter, we review systems, methods, and models we have used to link neuronal variability to perceptual decision

making in the human brain. We begin by describing how we identify task-relevant EEG components and use signal detection theory to relate these components to the behavioral data. We then use the well-known diffusion model of two-choice decision making to show that the trial-to-trial variability of these EEG components can be used to improve model fits of the behavioral data, providing evidence that the trial-to-trial variability we see in the EEG contains meaningful information and is not purely noise. We then turn to our work on combining EEG and fMRI to infer the cortical networks underlying perceptual decision making. We briefly discuss how EEG can be used to inform an fMRI analysis to tease apart individual processes underlying perceptual decision making. We then show how the trial-to-trial fluctuations in the EEG can be used to construct regressors that yield fMRI activations that are unobservable, given only behavioral or stimulus derived regressors. These specific results suggest that the trial-to-trial fluctuations we identify in the EEG represent latent processes such as attentional "polling" of the sensory input. In general, our results demonstrate that analysis of trial-to-trial variability of neural activity yields new insights into the constituent brain processes of decision making in the human brain.

Single-Trial Analysis of EEG

Traditionally, the analysis of EEG has relied on averaging event-locked data across hundreds of trials as well as across subjects, to uncover the neural signatures of the neurocognitive process under investigation. The main assumption of this approach is that trial averaging increases signal-to-noise ratio (SNR) by minimizing the background EEG activity relative to the neural activity correlated with experimental events. While this assumption is generally valid, it inevitably conceals inter-trial and inter-subject response variability. This trial-by-trial variability may carry important information regarding the underlying neural processes, which in turn might have important behavioral consequences.

In contrast, single-trial methods are often designed to exploit the large number of sensor arrays by spatially integrating information across channels to generate an aggregate representation (i.e., component) of the data that optimally discriminates between experimental conditions of interest. Spatial integration enhances the signal quality without loss of temporal precision common to trial averaging while the resulting discriminating components are often a better estimator of the underlying neurophysiological activity.

Methods that have been developed to extract components of interest from the EEG include independent component analysis (ICA) (Jung et al., 2001;

Makeig et al., 2002; Onton et al., 2006), common spatial patterns (CSP) (Guger et al., 2000; Ramoser et al., 2000) support vector machines (SVM) (Lal et al., 2004; Thulasidas et al., 2006) and linear discrimination (LD) based on logistic regression (Parra et al., 2002; Parra et al., 2005). LD in particular can be used to compute a set of spatial weights which maximally discriminate between experimental conditions over several different temporal windows, thus allowing the monitoring of the temporal evolution of discriminating activity. Unlike CSP, which tries to identify orientations in sensor space that maximize power, LD tries to maximize discrimination between two classes. Also unlike ICA, which is designed to minimize the correlation between spatial components (i.e., make spatial components as independent as possible (Hyvarinen et al., 2001) LD is used to identify components that maximize the correlation with relevant experimental events. All of these techniques linearly transform the original EEG signal via the following transformation

$$\mathbf{Y} = \mathbf{WX} \tag{1}$$

where $\mathbf{X}$ is the original data matrix, $\mathbf{W}$ is the transformation matrix (or vector) estimated using the different techniques, and $\mathbf{Y}$ is the resulting component/source matrix (or vector). Figure 10.1 summarizes how the LD technique can be used for binary discrimination.

Using the single-trial LD approach highlighted here, we explored the temporal characteristics of perceptual decision making in humans in an attempt to quantify the relationship between neural activity and behavioral output (Philiastides and Sajda, 2006). Motivated from the early work by Newsome and colleagues in primates (Britten et al., 1996; Britten et al., 1992), we reported the first noninvasive neural measurements of perceptual decision making in humans, that lead to neurometric functions predictive of psychophysical performance on a face versus car categorization task (see Figure 10.2A for examples of stimuli that were used). Specifically, we manipulated the difficulty of the task by changing the spatial phase coherence of the stimuli in a range that spanned psychophysical threshold. Carrying out the LD approach at different time windows and coherence levels revealed two EEG components that discriminated maximally between faces and cars as seen in Figure 10.2B for one subject. The early component was consistent with the well-known face-selective N170 (Bentin et al., 1996; Halgren et al., 2000; Jeffreys, 1996; Liu et al., 2000; Rossion et al., 2003) and its temporal onset appeared to be unaffected by task difficulty. The late component, appeared on average around 300 ms after the stimulus at the easiest condition and it systematically shifted later in time and became more persistent as a function of task difficulty. Both of these components showed substantial trial-to-trial variability (see Figure 10.2C) and

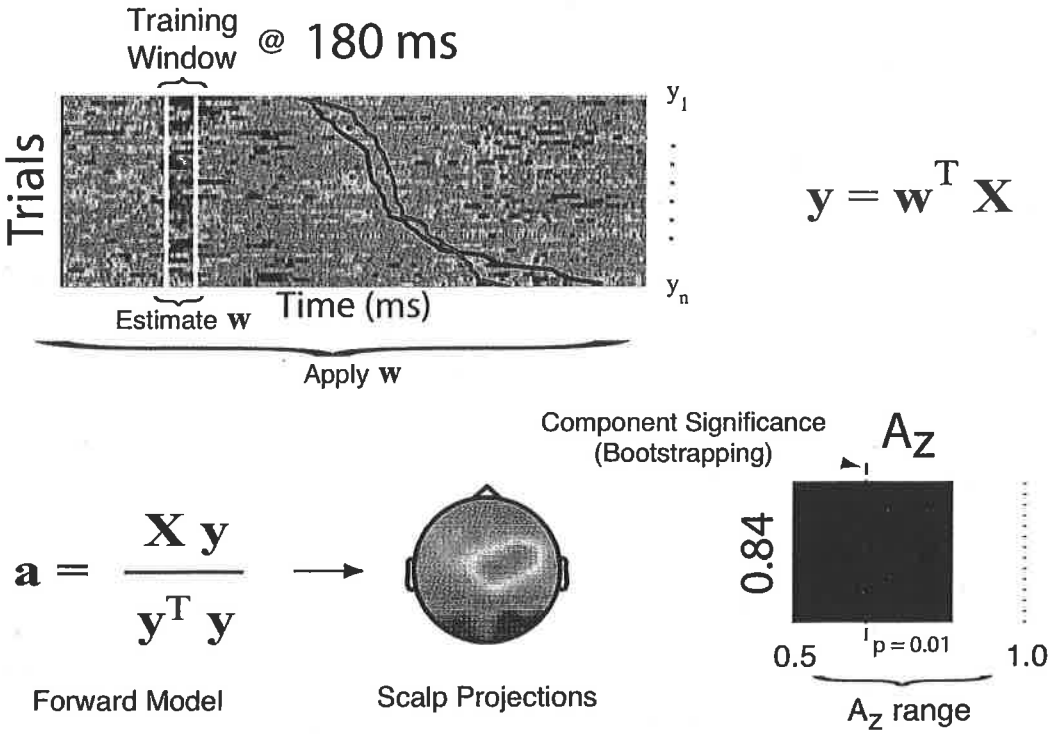

FIGURE 10.1 Summary of our linear discriminant (LD) methodology for extracting task-relevant components from single-trial analysis of the EEG. Each row of the discriminant component map represents a single trial across time. Discriminant components are represented by the **y** vectors. Trials are aligned to the onset of the stimulus (black vertical line) and are sorted by reaction time (sigmoidal curves). To construct this map we choose a training window, indicated by white vertical bars (for this example starting at 180 ms post-stimulus), during which we train the linear discriminator to estimate a weighting vector **w** across all sensors in **X**, such that **y** is maximally discriminating between the two experimental conditions (e.g. trial type 1 vs trial type 2). We use the forward model to project the discriminating component back to the sensors. An example scalp projection **a** is shown here and is used for interpreting the neuroanatomical significance of the components. To quantify the discriminator's performance we used ROC analysis and computed the area under the ROC curve (Az value).
Reproduced/adapted from Philiastides and Sajda, 2006).

they were both sensitive to decision-accuracy in that a high magnitude discriminator output value (i.e. **Y**) indicated an easy trial, whereas smaller values indicated more difficult decisions. Additional experimental manipulations enabled us to identify a third component, situated between the early and late components, around 220 ms post-stimulus, which systematically increased with increasing task difficulty. To rule out the possibility that this component is an artifact of the bottom-up processing of the stimulus we used a variant of our original paradigm where we colored the same images red or green and asked our subjects to either perform a simple color discrimination or the original face

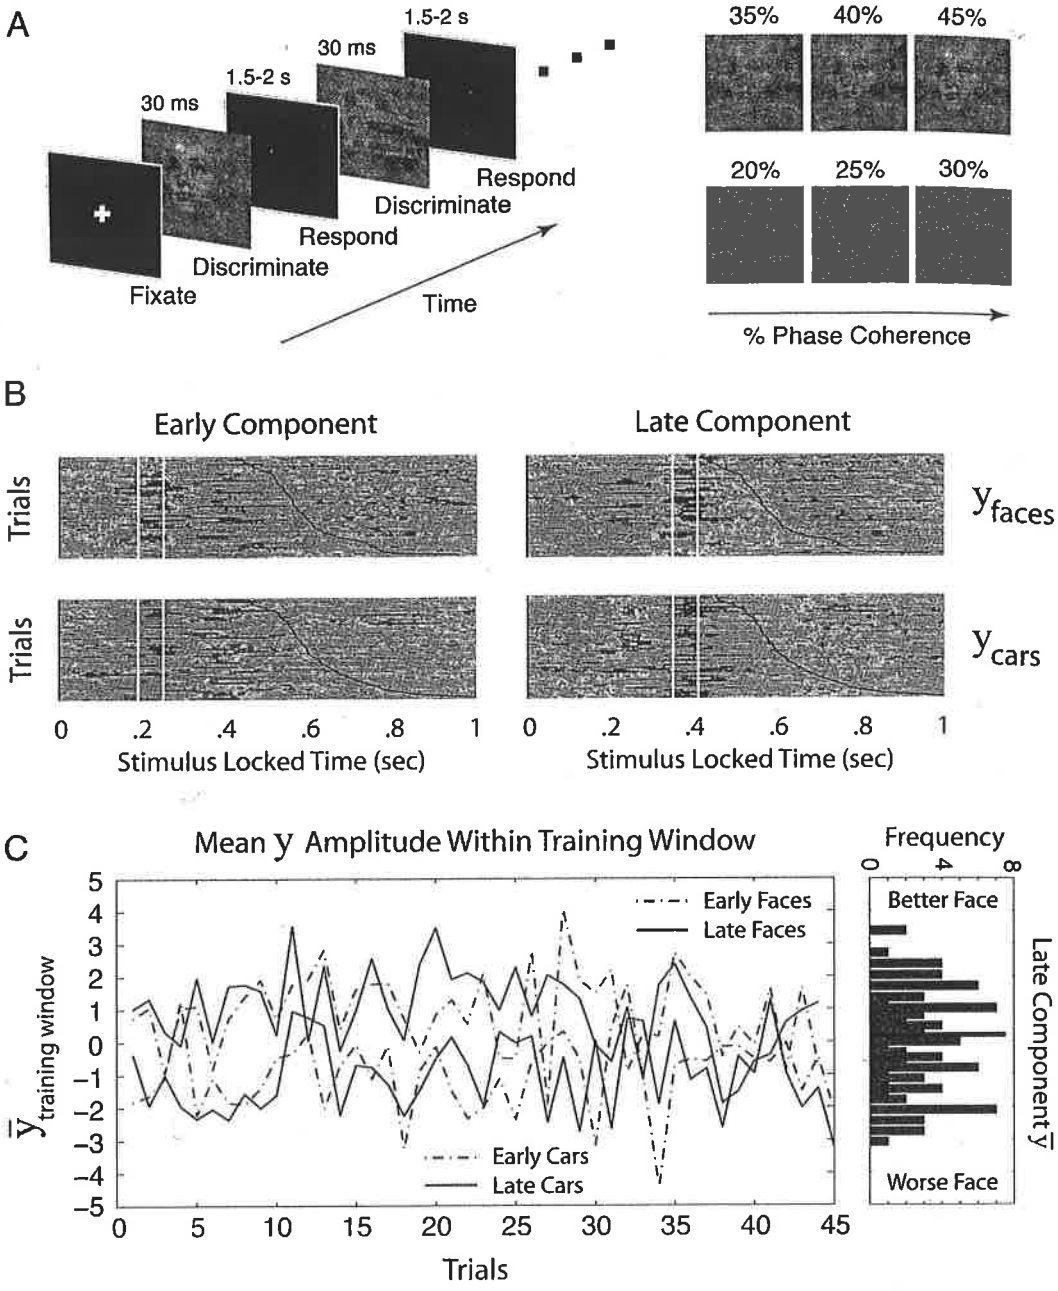

FIGURE 10.2 An example of a perceptual decision-making task and the task-relevant EEG components extracted using linear discrimination, together with their corresponding trial-to-trial variability. A) Behavioral paradigm (left) and sample face stimuli at difference levels of coherence (right). B) Discriminant component maps for one subject at 40% phase coherence. The four panels represent the face-vs-car discriminator output for the two EEG components (one "early" and one "late" relative to stimulus onset). Component maps are shown for both face and car trials using the training windows shown by the vertical white bars. Reaction time profiles are indicated by the black sigmoidal curves. The discriminator was designed to map face trials to positive (red) and car trials to negative values (blue). C) The mean EEG discriminator values within each of the training windows (y bar) for each trial and

FIGURE 10.2 (Continued) for each stimulus class. Shown are trial-by-trial values of the two components for faces and cars at 40% phase coherence only; note that these are not successive trials in the experiment, they are successive presentations of the stimulus condition to illustrate trial-to- trial variability. The amplitudes for the late component are also shown as a histogram (lower panel, right) with a cutoff to separate trials (into "better" face versus "worse" face) denoted by the thick black line. Reproduced from Ratcliff et al., 2009).

categorization task (Philiastides et al., 2006). This manipulation allowed us to keep the stimulus evidence unchanged while comparing the amplitude of the third component between a challenging face/car and a trivial color discrimination. We found that, for the same images, this component was significantly reduced when the subjects were merely discriminating the color of the stimulus confirming that this component is related to task difficulty.

These results taken together suggest that the different EEG components can be thought of as representing distinct cognitive events during perceptual decision making. Specifically, the early component appears to reflect the stimulus quality independent of the task (face/car or color discrimination) and is likely to represent early sensory processing. In contrast, the late component better represents information in the actual face/car decision process as it was shown to be a good predictor of overall behavioral performance during face categorization while its responses to the color discrimination were virtually diminished. Consistent with a top-down attentional control system, the difficulty component appears to be implicated in the recruitment of the relevant attentional and other neuronal resources required to make a difficult decision. Next we consider how the variability in these EEG components can be further interpreted within the context of a well-established and tested model of two-choice decision making.

Linking Trial-to-Trial Variability of EEG Components to Behavior Via the Diffusion Model

Advances in understanding decision processes in both psychology and neuroscience have produced models that require several sources of variability in order to fit experimental data. In psychological applications, the models attempt to fit accuracy and RT distributions for both correct and error responses. The need for assumptions about variability in the various components of processing come about because of the need to fit detailed behavior of error RTs. The major generalizations are that when accuracy is high and speed is stressed,

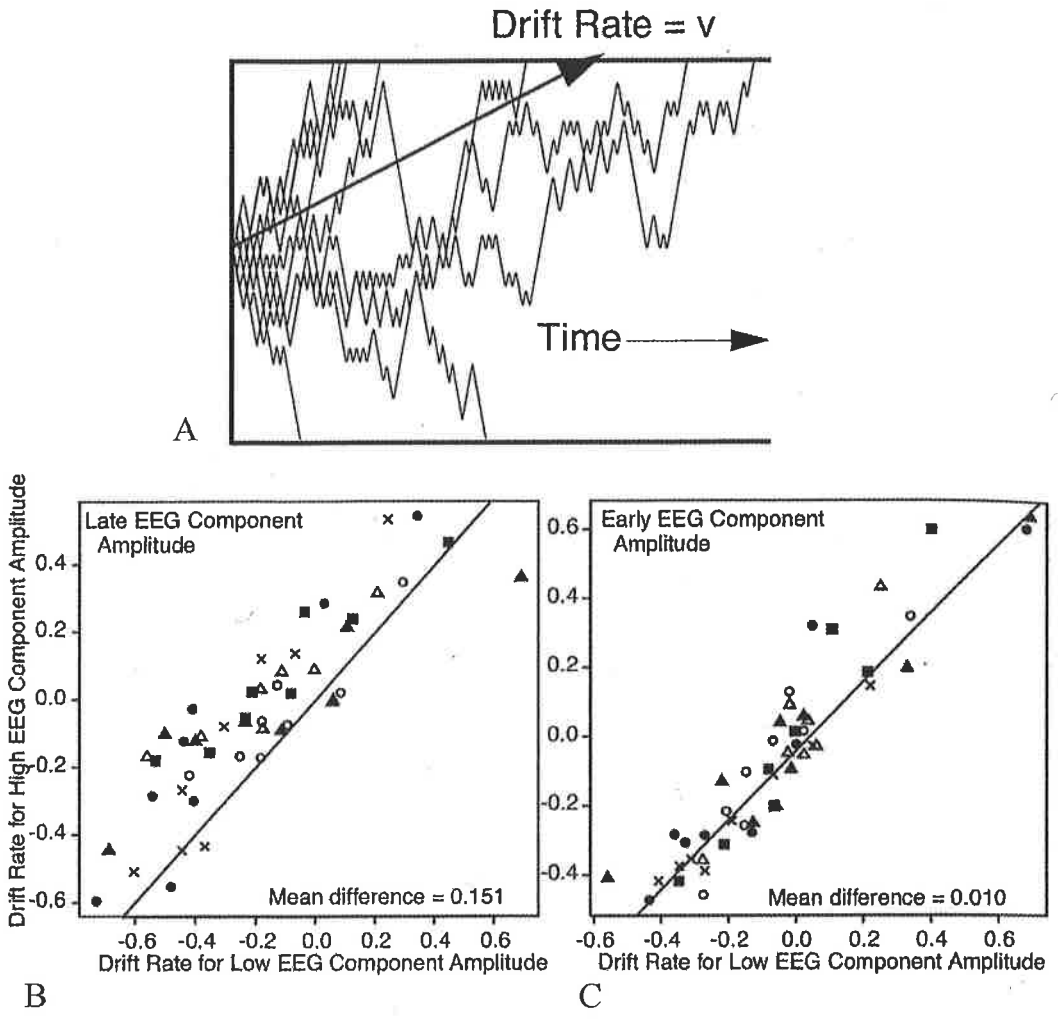

FIGURE 10.3 Linking EEG component variability to the diffusion model for two choice decision making. A) Simulated diffusion processes with the same mean drift rate, demonstrating that behavioral variability can be generated via the model. B) Fits for diffusion model drift rates for data sorted using "late" EEG component values. Drifts rates are systematically higher for EEG components having high values. D) When the "early" EEG component values are used to sort the data, there is no significant difference in the drift rates for high vs low EEG component amplitudes.
Adapted from Ratcliff et al., 2009.

errors are faster than correct responses, but when accuracy is lower and accuracy is stressed, errors are slower than correct responses. This pattern was very difficult to model and error RTs were ignored to a large degree in modeling until the mid 1990s.

In psychology, diffusion models have been shown to be able to successfully account for behavioral data in a range of experimental paradigms. These models assume a gradual accumulation of noisy evidence towards one of two decision criteria as in Figure 10.3A. The parameters of the model include the drift rate, or rate of accumulation of evidence.

In perceptual decision making, this represents the quality of the perceptual signal. The amount of evidence required to make a decision is represented by the distance between the two boundaries, and any bias towards one or other response is represented by any asymmetry between the starting point and either boundary. In addition to these parameters, the other components of processing, such as stimulus encoding and response output are represented by a single parameter which represents the duration of the non-decision components of processing.

The diffusion model, being a dynamic sequential sampling model, can be contrasted with signal detection theory (SDT) which is a static model. In SDT, there is only one source of variability, variability in perceptual strength across trials. This corresponds to variability in drift rate across trials in a diffusion model (see Figure 10.3A). Because there is only one source of variability in SDT, however, all sources of variability are collapsed into the one source. If we believe that there are multiple sources of variability, then SDT is clearly not adequate. We also know that subjects can trade speed for accuracy. A diffusion model analysis shows that to a good approximation, drift rates are invariant under such instructions with differences in accuracy and RT accounted for by a change in the parameter representing boundary separation. In contrast, SDT produces differences in discriminability as a function of speed-accuracy instruction manipulations.

Behavioral models of simple decision making require several sources of variability in order to fit data. Traditionally, behavioral measures (e.g. accuracy and RT) have been the sources of this variability. Neurophysiological measures, however, can also potentially be exploited. For instance, in terms of noninvasive neuroimaging, single-trial EEG offers the ability to track processing in a way that behavioral measures do not. EEG provides a millisecond by millisecond measure of the brain's electrical activity and therefore it is possible to link this activity to different parameters resulting from the psychological models of processing. (Philiastides et al., 2006) showed that the late component of processing correlated highly with drift rate from diffusion model fits to the experimental data. Because this late component provides an estimate of the quality of evidence on single trials we examined whether it could be used to index drift rate in a diffusion model analysis.

In our analysis, the data was divided in half as a function of the size of the late EEG component value. So for each trial, we decided whether the component values was greater or less than the mean for that condition (e.g., the histograms in Figure 10.2C) and then sorted the data into two groups for each condition for each subject. In the diffusion model, this would be equivalent to dividing the drift rates in each condition into two halves if the component value was an estimate of drift rate for each trial. The diffusion model was fit to the two groups of data for each subject, and results showed that for the

more face-like group of data, drift rates were more face-like than for the less face-like group of data (see Figure 10.3B). A similar analysis based on the early component of processing showed no difference (Figure 10.3C). In addition, the estimate of variability across trials in drift rate averaged over subjects was significantly lower than the value obtained by fitting all the data for each subject.

These results show a close connection between the variability of the late single trial EEG component value and drift rate in the diffusion model produced by fitting the behavioral data. This provides additional support for both the psychological reality of sequential sampling models. To better understand the neuronal origins of this variability however, one needs to use the single-trial information obtained from the EEG to inform the analysis of fMRI data collected for the same task. The next section describes our efforts of combining EEG and fMRI to describe decision making with high temporal as well as high spatial precision.

Coupling EEG to fMRI for Inferring Cortical Networks Underlying Perceptual Decision-Making

Despite significant progress made in understanding perceptual decision making in humans using EEG and fMRI in isolation, the spatial localization restrictions of EEG and the temporal resolution constraints of fMRI suggest that only a fusion of these modalities can provide a full spatiotemporal characterization of this process. Animal experiments have already demonstrated that hemodynamic signals are more closely coupled to synaptic than spiking activity and that changes in the fMRI BOLD signal can correlate tightly with synchronized oscillatory activity recorded from local field potentials (LFPs) (Logothetis, 2008; Logothetis and et al., 2001; Niessing et al., 2005; Viswanathan and Freeman, 2007). Under these premises, it is reasonable to assume that neural activity reflected in the EEG could also correlate well with the underlying BOLD hemodynamic response.

We have used two methods for coupling EEG and fMRI activity. One is to record EEG and fMRI simultaneously, and explicitly utilize the trial-to-trial fluctuations in the EEG components to construct regressors for use in the analysis of the fMRI activity. In a second approach, we use each modality separately and, given an appropriate experimental design, derive MRI regressors that are modulated by average amplitudes of EEG components associated with different experimental conditions. Specifically, we first perform single-trial LD to identify EEG components of interest (e.g. early, late and difficulty components). Assuming the discriminator is trained with T samples within each window (τ) of interest, the output (y_τ) has dimensions $T \times N$, where N is the total number

of trials. To achieve more robust single-trial estimates for $\mathbf{y}_\tau$, averaging across all training samples is performed:

$$\bar{\mathbf{y}}_{\tau,i} = \frac{1}{T}\sum_{j=1}^{T}\mathbf{y}_{\tau,i,j} \tag{2}$$

Where i is used to index trials and j training samples. For the first approach, $\bar{\mathbf{y}}_{\tau,i}$ is then used to modulate the amplitude of the different fMRI regressor events. Finally, the parametric regressor is convolved with a prototypical hemodynamic response function, which is used to model the fMRI data in the context of a general linear model (GLM). This process can be repeated for multiple windows/components (τ) each resulting in a separate regressor (see Figure 10.5 for a summary of this approach). Identifying the brain regions that correlate with each of these regressors will enable a comprehensive characterization of the cortical network involved in perceptual decision making.

In the absence of simultaneous EEG/fMRI measurements, the second method, namely using average values of EEG components to modulate regressors which are a function of the experiment condition (e.g. trial type). can be used instead. In this case the discriminator output associated with each component and each experimental condition is averaged across trials:

$$\bar{\mathbf{y}}^{c}_{\tau,i} = \frac{1}{N}\frac{1}{T}\sum_{i=1}^{N}\sum_{j=1}^{T}\mathbf{y}^{c}_{\tau,i,j} \tag{3}$$

where c is used to index the different experimental conditions. The average discriminator output per component and experimental condition obtained from equation 4 can be used to model the fMRI data. Importantly, $\bar{\mathbf{y}}^{c}_{\tau,i}$ is now a scalar - that is, in the absence of single-trial information during the fMRI session, all like trials will be modeled in the same way. Though inter-trial variability is ultimately concealed in this formulation, important information regarding the localization of each of the EEG components, that would otherwise be unattainable using EEG or fMRI alone, can now be obtained.

We first considered the latter approach for the perceptual decision making work presented above (Philiastides et al., 2006; Philiastides and Sajda, 2006). As highlighted earlier, the strength of our early EEG component was proportional to the stimulus evidence (i.e., stronger for easy than hard trials) and it remained unchanged during the face/car and color discriminations. The late EEG component also responded proportionally to the stimulus evidence during the face/car discrimination, but it was stronger across all difficulty levels relative to the early one. Unlike the early component, however, the strength of the late component was significantly reduced during the color discrimination.

In contrast to both the early and late components, the strength of the difficulty component was inversely proportional to the amount of stimulus evidence (i.e., stronger for hard rather than easy trials).

As a result of these observations, we constructed three parametric fMRI regressors, one for each of the early, difficulty, and late components in order to analyze the fMRI data collected for the same task. To modulate the heights of the corresponding regressor events we estimated the relative strengths of our components with respect to the difficulty (i.e., low [L] vs high [H] coherence) and the type of task (i.e., face vs car [FC] or red vs green [RG]) (i.e. $\bar{y}_{\tau,i}^{FC,L}$, $\bar{y}_{\tau,i}^{FC,H}$, $\bar{y}_{\tau,i}^{RG,L}$, $\bar{y}_{\tau,i}^{RG,H}$ where $\tau = (early, difficulty, late)$).

Figure 10.4 summarizes our findings in a form of a spatiotemporal diagram and demonstrates that a cascade of events associated with perceptual decision making takes place in a highly distributed neural network. These include early visual perception (early component), task/decision difficulty (difficulty component) and postsensory/decision-related events (late component). Clear is that by exploiting the variability across EEG component and trial types (i.e. the variability in the $\bar{y}_{\tau,i}^{FC,L}$, $\bar{y}_{\tau,i}^{FC,H}$, $\bar{y}_{\tau,i}^{RG,L}$, $\bar{y}_{\tau,i}^{RG,H}$) we were able to infer a more detailed picture of the cortical networks underlying perceptual decision making.

Next, we wanted to demonstrate the efficacy of the first method of exploiting the trial-to-trial variability measured during simultaneous acquisition of EEG and fMRI. Simultaneous EEG/fMRI is a relatively new neuroimaging modality that enables the simultaneous measurement of electrical and blood oxygenation level dependent (BOLD) activity. The electrical activity measured via EEG is temporally precise (millisecond resolution) and is a direct measure of neural activity whereas the BOLD activity measured via fMRI is more spatially localized (millimeter resolution) and represents an indirect measure of neural activity. We have conducted experiments that use such multimodal neuroimaging to correlate the trial-to-trial variability of temporally precise EEG components with simultaneously measured BOLD activity. The underlying hypothesis is that the trial-to-trial variability in the EEG components has information content that is meaningful, for example representing the dynamics of latent brain states that are unobservable via stimulus or behaviorally derived measures.

Given the technical challenges in acquiring EEG and fMRI simultaneously, we first focused on a very simple and classic perceptual detection paradigm, the auditory oddball task (Donchin and Coles, 1988; Picton, 1992; Polich, 2007). A detailed description of the paradigm and data acquisition can be found in (Goldman et al., 2009). Here we focus on the method for coupling the single-trial variability of the EEG components with the fMRI measured BOLD activity.

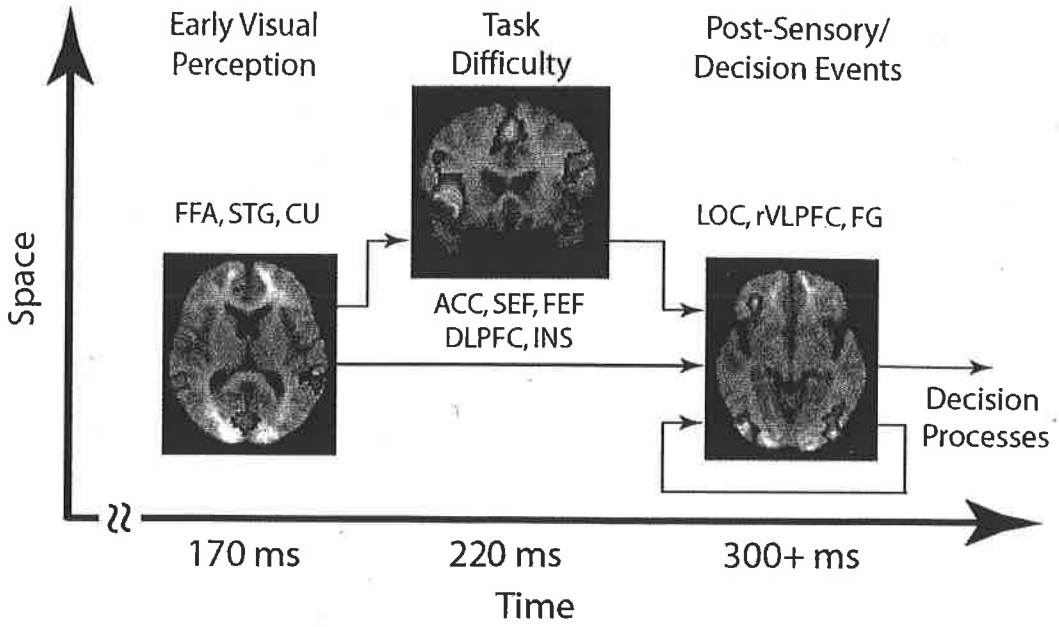

FIGURE 10.4 Spatio-temporal processing timing diagram resulting from an EEG-informed fMRI analysis For the early component, we see significant correlations with activity in areas implicated in early visual processing of objects/faces such as the fusiform face area (FFA) and the superior temporal sulcus (STS), (Allison et al., 1999; Haxby et al., 2000; Hoffman and Haxby, 2000; Kanwisher et al., 1997; Puce et al., 1998). The difficulty component is correlated with activity in the supplementary and frontal eye fields (SEF/FEF), the anterior cingulate cortex (ACC), the dorsolateral prefrontal cortex (DLPFC) and the anterior insula (INS). These observations are consistent with the interpretation that there exist an attentional control system that exerts top-down influence on decision making (Heekeren et al., 2004; Heekeren et al., 2008). Finally, the late component is correlated with activity in the lateral occipital complex (LOC) and in the right ventrolateral prefrontal cortex (rVLPFC). Aside from its involvement in object categorization (Grill-Spector et al., 2004; Grill-Spector et al., 2001; Grill-Spector et al., 1999; James et al., 2000, 2002), the LOC has been implicated in "perceptual persistence" (Ferber et al., 2002; Large et al., 2005), a process in which a percept assembled by lower-visual areas is allowed to remain in the visual system, via feedback pathways, as a form of iconic memory (Coltheart, 1980; DiLollo, 1977; VanRullen and Koch, 2003).
Adapted from Philiastides and Sajda, 2007.

Figure 10.5 shows specifically how we utilize the trial-to-trial variability in the EEG components to construct novel regressors for correlating with the BOLD activity. In this example two EEG components are identified after stimulus onset, one at 250ms and the other at 400ms. Both of these components discriminate target from nontarget stimulus and both are stimulus locked. Note that though both components have discriminative power in terms of

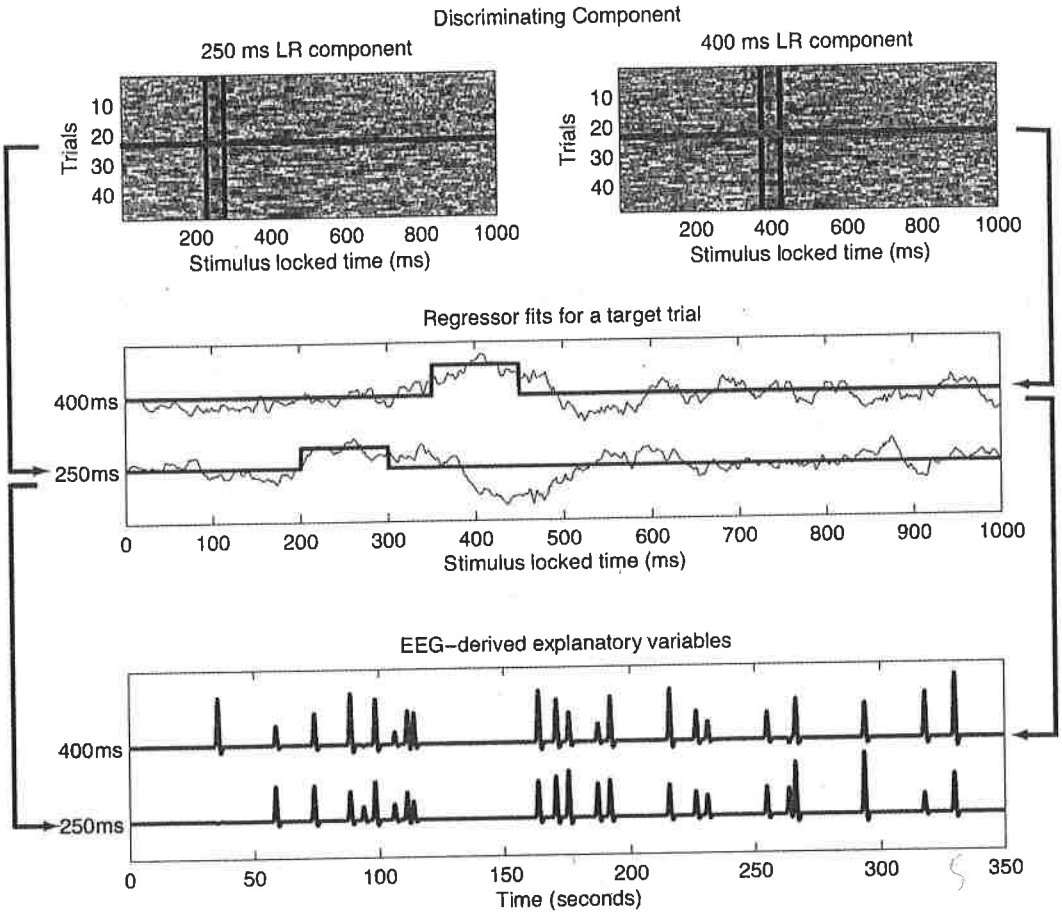

FIGURE 10.5 Our approach for coupling the trial-to-trial variability of the EEG with the BOLD signal, given simultaneous EEG/fMRI acquisition. (top) Output for all trials of the single-trial EEG discriminator for two stimulus-locked 50 ms windows (data between black vertical bars) centered at 250 ms and 400 ms post stimulus-onset. Hot to cold color scale indicates positive to negative values of the discriminator output (i.e. positive and negative correlations). (middle) EEG discriminator output for a single target trial for each of the two components (black curves), showing the fMRI event model amplitude as the average of the discriminator output for each window, 250 ms (blue) and 400 ms (red), for the trial. (bottom) Single-trial fMRI model for target trials across the entire session for the 250 ms and 400 ms windows. Note that the event timing for each of the two windows is the same, but the event amplitudes are different. Reproduced from Goldman et al. (2009).

target from non-target trials, their trial-to-trial variability is not 100% correlated, and therefore their individual trial-to-trial variability in principal could capture different aspects of the decision making process. For each of these components, we construct regressors by using the amplitude of the component on each trial to modulate the amplitude of a boxcar regressor. A regressor is constructed for each EEG component and then convolved with the hemodynamic

response function. This leads to two explanatory variables that capture the trial-to-trial variability of the EEG components and can be used in a general linear model (GLM) analysis to correlate with the BOLD activity. Finally we orthogonalize these EEG-derived explanatory variables with respect to traditional stimulus and behaviorally derived regressors, thus ensuring the activity they capture is purely correlation with the EEG component trial-to-trial variability.

Figure 10.6 shows results for statistically significant fMRI activations resulting from trial-to-trial variability in the EEG during an auditory oddball paradigm. (Additional details can be found in (Goldman et al., 2009). Several interesting observations can be made. The first is that the locations of the fMRI

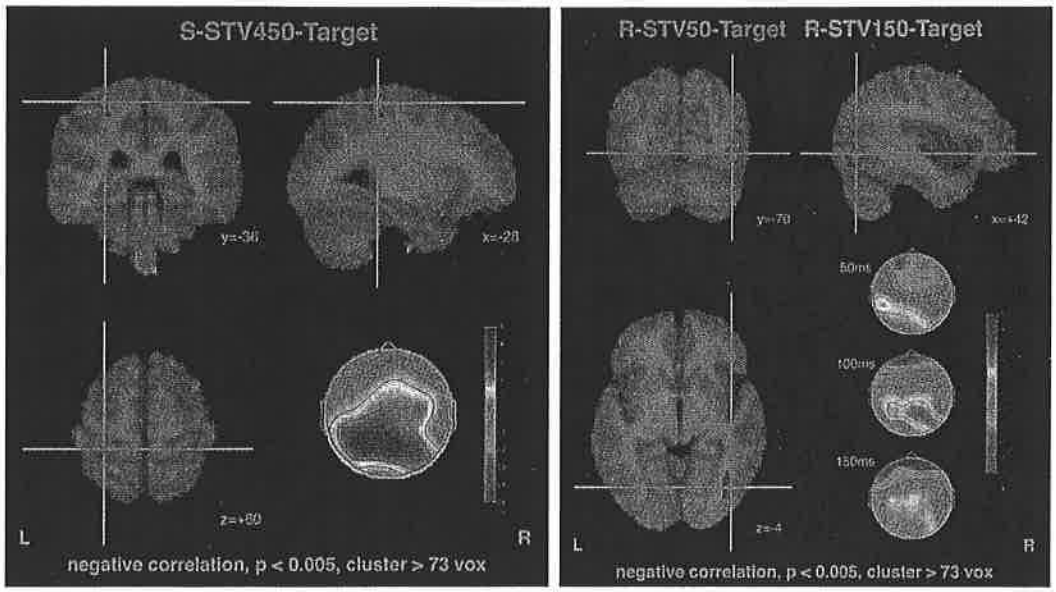

A B

FIGURE 10.6 (A) fMRI activations for the stimulus-locked single-trial analysis showing regions with significant BOLD signal correlation ($p<0.005$, cluster>73 voxels, negative correlation) with single-trial variability to targets for the 450ms window, S-STV450-Targ. For target tones, only the stimulus-locked 450ms window passed both the EEG and fMRI thresholds. Shown also is the scalp topography of the corresponding 450ms window stimulus-locked EEG discriminating component (arbitrary units). (B) fMRI activations for the response-locked single-trial analysis showing regions with significant BOLD signal correlation ($p<0.005$, cluster>73 voxels, negative correlation) to response-locked single-trial variability to target tones. The response-locked 50ms window, R-STV50-Targ (blue), and 150ms window, R-STV150-Targ (green), passed both the EEG and fMRI thresholds. The scalp maps of the output of the EEG discriminator for the 3 windows from 50-150ms response-locked are also shown (arbitrary units). Reproduced from Goldman et al. (2009).

activations that arise from the trial-to-trial variability are remarkably consistent with the scalp projections of the corresponding EEG components. Since only the trial-to-trial variability of the EEG components is used in the GLM analysis, no information about the locations of the EEG components (i.e., scalp projections) is used to constrain the spatial location of the fMRI activations. Thus, the fact that that EEG trial-to-trial variability leads to consistent localizations in both modalities supports the interpretation that that this variability is neurophysiologically meaningful. A second observation pertains to the specific locations of the activated regions. Figure 10.6A shows activations localized to somatosensory cortex with Figure 10.6B showing activations in lateral occipital complex. Given that the paradigm is an auditory oddball task, why would two cortical areas selective to other modes of sensory input (somatosensation and vision) be activated? The third observation provides a clue, namely that these activation are negatively correlated with the EEG component trial-to-trial variability. In other words, on a trial, when the EEG component increases the fMRI BOLD signal decreases, and vice versa. A possible interpretation of these three observations is that the trial-to-trial variability represents an attentional "polling" of the sensory inputs, with attention allocated in a "push-pull" fashion—i.e. when attention is directed to one modality it is pulled from the other sensory modalities. We are further investigating this hypothesis with additional experiments that consider activations elicited when the oddball paradigm is performed via a different sensory modality (e.g., visual or somatosensory). In general, these findings suggest that the trial-to-trial variability in EEG components may reflect latent brain states that, when combined with fMRI, can yield novel insight into perceptual decision making. Currently, we are conducting experiments for the face/car paradigm, within the context of this type of simultaneous EEG/fMRI analysis. We believe that this will enable us better elucidate additional details of the underlying cortical networks (i.e. improve upon the accuracy of the spatio-temporal diagram of Figure 10.6).

Conclusions

Our efforts using spatially and temporally precise neuroimaging, machine learning and signal processing, and cognitive modeling to measure and analyze neuronal variability have so far been aimed at understanding how we make very simple decisions. The ultimate challenge is to understand the cortical circuits involved in making typical, everyday decisions such as "Should I take the subway or walk to work today?" or "Should I read this chapter given how busy I am?" The neuronal variability underlying these everyday decisions may

in fact be the key to understanding what makes each of us unique—i.e,. what differentiates individuals from one another. In addition, analysis of such neuronal variability may be critical for identifying precursors to behavioral changes, including pathological changes that are associated with cognitive deficits and disease. Thus the imaging, analysis and modeling methods we have described can be seen as a suite of tools, to be used in concert, for measuring and analyzing neuronal variability associated with both normal and abnormal decision making in the human brain.

Acknowledgments

This research was supported by funding from the NIH (grants R33-EB004730 and R01-MH085092) and DARPA (contract NBCHC080029).

REFERENCES

Allison, T., Puce, A., Spencer, D.D., McCarthy, G. (1999) Electrophysiological studies of human face perception: potentials generated in occipitotemporal cortex by face and nonface stimuli. *Cereb Cortex, 9,* 415–430.

Bentin, S., Allison, T., Puce, A., Perez, A., McCarthy, G. (1996) Electrophysiological studies of face perception in humans. *Journal of Cognitive Neuroscience, 8,* 551–565.

Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., Movshon, J.A. (1996) A relationship between behavioral choice and visual responses of neurons in macaque MT. *Vis Neurosci, 14,* 87–100.

Britten, K.H., Shadlen, M.N., Newsome, W.T., Movshon, J.A. (1992) The analysis of visual motion: A comparison of neuronal and psychophysical performance. *J of Neuroscience, 12,* 4745–4765.

Coltheart, M. (1980) The persistences of vision. *Philos T Roy Soc B, 290,* 57–69.

DiLollo, V. (1977) Temporal characteristics of iconic memory. *Nature, 267,* 241–243.

Donchin, E., Coles, M.G. (1988) Is the P300 component a manifestation of context updating? *Behavioral and Brain Sciences, 11,* 357–374.

Ferber, S., Humphrey, G.K., Vilis, T. (2002) The lateral occipital complex subserves the perceptual persistence of motion-defined groupings. *Cereb Cortex, 35,* 793–801.

Goldman, R.I., Wei, C.-Y., Philiastides, M.G., Gerson, A.D., Friedman, D., Brown, T.R., Sajda, P. (2009) Single-trial discrimination for integrating simultaneous EEG and fMRI: Identifying cortical areas contributing to trial-to-trial variability in the auditory oddball task. *Neuroimage, 47,* 136–147.

Grill-Spector, K., Knouf, N., Kanwisher, N. (2004) The fusiform face area subserves face perception, not generic within-category identification. *Nat Neurosci, 7,* 555–562.

Grill-Spector, K., Kourtzi, Z., Kanwisher, N. (2001) The lateral occipital complex and its role in object recognition. *Vision Res, 41,* 1409–1422.

Grill-Spector, K., Kushnir, T., Edelman, S., Avidan, G., Itzchak, Y., Malach, R. (1999) Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron, 24,* 187–203.

Guger, C., Ramoser, H., Pfurtscheller, G. (2000) Real-time EEG analysis with subject-specific spatial patterns for a brain computer interface BCI. IEEE *Trans Rehabil Eng, 8,* 441–446.

Halgren, E., Raij, T., Marinkovic, K., Jousmaki, V., Hari, R. (2000) Cognitive response profile of the human fusiform face area as determined by MEG. *Cereb Cortex, 10,* 69–81.

Haxby, J.V., Hoffman, E.A., Gobbini, M.I. (2000) The distributed human neural system for face perception. *Trends Cogn Sci, 4,* 223–233.

Heekeren, H.R., Marrett, S., Bandettini, P.A., Ungerleider, L.G. (2004) A general mechanism for perceptual decision-making in the human brain. *Nature, 431,* 859–862.

Heekeren, H.R., Marrett, S., Ungerleider, L. (2008) The neural systems that mediate human perceptual decision making. *Nat Rev Neurosci, 9,* 467–479.

Hoffman, E., Haxby, J.V. (2000) Distinct representation of eye gaze and identity in the distributed human neural system for face perception. *Nat Neurosci, 3,* 80–84.

Hyvarinen, A., Karhunen, J., Oja, E. (2001) *Independent component analysis.* New York: John Wiley and Sons.

James, T.W., Humphrey, G.K., Gati, J.S., Menon, R.S., Goodale, M.A. (2000) The effects of visual object priming on brain activation before and after recognition. *Curr Biol, 10,* 1017–1024.

James, T.W., Humphrey, G.K., Gati, J.S., Menon, R.S., Goodale, M.A. (2002) Differential effects of viewpoint on object-driven activation in dorsal and ventral streams. *Neuron, 35,* 793–801.

Jeffreys, D.A. (1996) Evoked studies of face and object processing. *Visual Cognition, 3,* 1–38.

Jung, T.P., Makeig, S., McKeown, M.J., Bell, A.J., Lee, T.W., Sejnowski, T.J. (2001) Imaging brain dynamics using independent component analysis. *Proc IEEE, 89,* 1107–1122.

Kanwisher, N., McDermott, J., Chun, M.M. (1997) The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Jl of Neurocience, 17,* 4302–4311.

Lal, T.N., Schroder, M., Hinterberger, T., Weston, J., Bogdan, M., Birbaumer, N., Scholkopf, B. (2004) Support vector channel selection in BCI. *IEEE Trans Biomed Eng, 51,* 1003–1010.

Large, M.E., Aldcroft, A., Vilis, T. (2005) Perceptual continuity and the emergence of perceptual persistence in the ventral visual pathway. *J of Neurophysiology, 93,* 3453–3462.

Liu, J., Higuchi, M., Marantz, A., Kanwisher, N. (2000) The selectivity of the occipitotemporal M170 for faces. *Neuroreport, 11,* 337–341.

Logothetis, N.K. (2008) What we can do and what we cannot do woth fMRI. *Nature, 453,* 869–878.

Logothetis, N.K., et al., J.P. (2001) Neurophysiological investigation of the basis of the fMRI signal. *Nature, 412,* 150–157.

Makeig, S., Westerfield, M., Jung, T.P., Enghoff, S., Townsend, J., Courchesne, E., Sejnowski, T.J. (2002) Dynamic brain sources of visual evoked responses. *Science, 295,* 690–694.

Niessing, J., Ebisch, B., Schmidt, K.E., Niessing, M., Singer, W., Galuske, R.A.W. (2005) Hemodynamic signals correlate tightly with synchronized gamma oscillations. *Science, 309,* 948–951.

Onton, J., Westerfield, M., Townsend, J., Makeig, S. (2006) Imaging human EEG dynamics using independent component analysis. *Neurosci Biobehav Rev, 30,* 802–822.

Parra, L., Alvino, C., Tang, A., Pearlmutter, B., Young, N., Osman, A., Sajda, P. (2002) Linear spatial integration for single-trial detection in encephalography. *Neuroimage, 17,* 223–230.

Parra, L.C., Spence, C.D., Gerson, A.D., Sajda, P. (2005) Recipes for the linear analysis of EEG. *Neuroimage, 28,* 326–341.

Philiastides, M.G., Ratcliff, R., Sajda, P. (2006) Neural representation of task difficulty and decision-making during perceptual categorization: a timing diagram. *J of Neuroscience, 26,* 8965–8975.

Philiastides, M.G., and Sajda, P. (2006) Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cereb Cortex, 16,* 509–518.

Philiastides, M.G., and Sajda, P. (2007) EEG-informed fMRI reveals spatiotemporal characteristics of perceptual decision making. *J of Neuroscience, 27,* 13082–13091.

Picton, T.W. (1992) The P300 wave of the human event-related potential. *J Clin Neurophysiol, 9,* 456–479.

Polich, J. (2007) Updating P300: an integrative theory of P3a and P3b. *Clin Neurophysiol, 118,* 2128–2148.

Puce, A., Allison, T., Bentin, S., Gore, J.C., McCarthy, G. (1998) Temporal cortex activation in humans viewing eye and mouth movements. *J of Neuroscience, 18,* 2188–2199.

Ramoser, H., Müller-Gerking, J., Pfurtscheller, G. (2000) Optimal Spatial Filtering of Single Trial EEG During Imagined Hand Movement. *IEEE Trans. Rehab. Eng, 8,* 441–446.

Ratcliff, R., Philiastides, M.G., Sajda, P. (2009) Quality of evidence for perceptual decision making is indexed by trial-to-trial variability of the EEG. *Proceedings of the National Academy of Science, 106,* 6539–6544.

Rossion, B., Joyce, C., Cottrell, G.W., Tarr, M.J. (2003) Early laterization and orientation tuning for face, word, object processing in the visual cortex. *Neuroimage, 20,* 1609–1624.

Thulasidas, M., Guan, C., Wu, J. (2006) Robust classification of EEG signals for brain-computer interface. *IEEE Trans Neural Syst Rehabil Eng, 14,* 24–29.

VanRullen, R., Koch, C. (2003) Visual selective behavior can be triggered by a feed-forward process. *Journal of Cognitive Neuroscience, 15,* 209–217.

Viswanathan, A., Freeman, R.D. (2007) Neurometabolic coupling in cerebral cortex reflects synaptic more than spiking activity. Neurometabolic coupling in cerebral cortex reflects synaptic more than spiking activity. *Nat Neurosci, 10,* 1308–1312.