UNIVERSITY OF CALIFORNIA,
IRVINE

Refining understanding of human decision making by testing integrated neurocognitive models of EEG, choice and reaction time

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Psychology

by

Michael Dawson Nunez

Dissertation Committee:
Professor Ramesh Srinivasan, Chair
Associate Professor Joachim Vandekerckhove
Distinguished Professor Barbara A. Dosher

2017

DEDICATION

In memory of Chris D. Bush, a friend and brother.

TABLE OF CONTENTS

LIST OF FIGURES

Page

LIST OF TABLES

ACKNOWLEDGMENTS

This work did not occur in a vacuum. I am immensely appreciative for Josh Tromberg's help with data collection and EEG preprocessing. Josh was instrumental to the beginning of my research in the Human Neuroscience Lab. Aishwarya Gosai, Sean O'Reilly-Jones, Kitty Lui, Michelle Cheung, and Kiana Scambray were equally helpful with data collection, EEG artifact mitigation, and other various tasks around the lab. Bill Winter is thanked for his initial help when I was a new graduate student. Cort Horton is well appreciated for creating useful MATLAB functions for EEG visualization and artifact removal. Siyi Deng and Sam Thorpe are thanked for their contributions to the FEM solutions and anatomical fMRI image generation. Current and past members of the Human Neuroscience Lab and the Cognition and Individual Differences Lab are thanked for their useful comments and good company. The anonymous reviewers of Chapter 3 are appreciated for their constructive suggestions and comments. Mentors and committee members include Paul L. Nunez, Michael D. Lee, Mark Steyvers, Hernando Ombao, Emre Neftci, and Barbara A. Dosher. I am immensely grateful to my research advisors and collaborators Joachim Vandekerckhove and Ramesh Srinivasan. Special thank you to my mother, Kirsty D. Nunez.

The following work was supported by (1) the Cognitive Sciences Department, at the University of California, Irvine, (2) NIH grant 2R01MH68004, (3) NSF grant #1230118 from the Methods, Measurements, and Statistics panel, (4) grant #48192 from the John Templeton Foundation, and (5) NSF grant #1658303 from the Cognitive Neuroscience panel.

Frontiers in Psychology and Elsevier (through The Journal of Mathematical Psychology) are appreciated for publishing Chapters 2 and 3 respectively.

CURRICULUM VITAE

Michael Dawson Nunez

EDUCATION

Doctor of Philosophy in Psychology w/ cont. in Cog. Neuro. **2017**
University of California, Irvine *Irvine, California*

Master of Science in Cognitive Neuroscience **2017**
University of California, Irvine *Irvine, California*

Master of Science in Statistics **2015**
University of California, Irvine *Irvine, California*

Bachelor of Science in Mathematics and Economics **2010**
Tulane University *New Orleans, Louisiana*

RESEARCH EXPERIENCE

Graduate Research Assistant **2012–2017**
University of California, Irvine *Irvine, California*

Research Assistant **2011–2012**
Tulane University *New Orleans, Louisiana*

TEACHING EXPERIENCE

Teaching Assistant **2012–2016**
University of California, Irvine *Irvine, California*

PUBLICATIONS

How attention influences perceptual decision making: Single-trial EEG correlates of drift-diffusion model parameters **2017**

Michael D. Nunez, Joachim Vandekerckhove, Ramesh Srinivasan

Journal of Mathematical Psychology

Individual differences in attention influence perceptual decision making **2015**

Michael D. Nunez, Ramesh Srinivasan, Joachim Vandekerckhove

Frontiers in Psychology

Electroencephalography (EEG), neurophysics, experimental methods, and signal processing **2016**

Michael D. Nunez, Paul L. Nunez, Ramesh Srinivasan

Handbook of Neuroimaging Data Analysis

SOFTWARE

artscreenEEG `https://github.com/mdnunez/artscreenEEG/`

MATLAB repository to perform basic artifact correction on electroencephalographic (EEG) data

ABSTRACT OF THE DISSERTATION

Refining understanding of human decision making by testing integrated neurocognitive models of EEG, choice and reaction time

By

Michael Dawson Nunez

Doctor of Philosophy in Psychology

University of California, Irvine, 2017

Professor Ramesh Srinivasan, Chair

The cognitive process and time course of quick human decision making was evaluated using reaction time, choice distributions, and human electrophysiology as recorded by EEG. These data were used to evaluate drift-diffusion models, a class of decision-making models that assume a stochastic accumulation of evidence on each trial, within hierarchical Bayesian frameworks. The first goal was to elucidate the effect of visual attention on decision making. To this aim two studies were performed. In the first study it was found that individual differences in evidence accumulation rates and non-decision time (preprocessing and motor response times) can be explained by attentional differences as measured by steady-state visual evoked potential (SSVEP) responses to the flicker frequency of signal and noise components of the visual stimulus. Participants who were able to suppress their SSVEP response to visual noise in high frequency bands were able to accumulate correct evidence faster and had shorter non-decision times, leading to more accurate responses and faster response times. In the second study it was found that measures of attention obtained from simultaneous EEG recordings can explain per-trial evidence accumulation rates and perceptual preprocessing times during a visual decision making task. That is, single-trial evoked EEG responses, P200s to the onsets of visual noise and N200s to the onsets of visual signal, explain single-trial evidence accumulation and preprocessing times. The second goal was obtain inference

about the time course of quick decision making. A method of estimating and verifying individuals' visual encoding time is proposed using traditional event-related potential (ERP) measures. The possibility of using single-trial N200 and trial-averaged N200 ERP latencies as estimates of human visual encoding time is explored using both simple linear regression and complex hierarchical Bayesian modeling. Posterior distributions of linear-effect parameters suggest that EEG responses to the onset of visual stimuli reflect stimulus encoding times. The possibility of using a verifiable EEG measure of the time course of motor preparation is also explored. Finally, a theoretical cognitive framework for quick decision making is proposed which assumes differential mechanisms of visual encoding, drift-diffusion evidence accumulation, and motor response.

Chapter 1

In search of a simple explanation of human decision making

1.1 The prevalence of quick decision-making

You will make many quick decisions tomorrow. You could make a decision to buy an apple or a piece of chocolate at the convenience store, make a decision to walk right or left when approaching a tree in your path, make a decision to hug someone you love or to kiss them, or even make a decision to either jump out of the way of a vehicle moving towards you or stand in the same spot and hope that the vehicle veers away. Human beings make many types of these decisions every day, and quick decisions are diverse. Some decisions are more important to us than others, some our integral to our survival, some quick decisions are made faster and more accurate by certain people, sometimes it takes awhile for us to recognize that a decision needs to be made.

This type of speeded decision making that occurs within approximately two seconds is the topic of this dissertation. The goal of my thesis is to illuminate the decision-making process

with a simple explanation (e.g a quantitative mathematical"model") of both the time course and attentional process of decision-making. This model explains observations we make about human behavior and the human brain, presented in terms of what is thought to occur within the human mind (i.e. the "cognition" of the mind).

1.1.1 A guiding example of quick decision-making

There are many situations on the road when the driver of a vehicle must decide to stop or accelerate through an intersection by observing a traffic light. The presence of a green arrow for an adjacent lane (i.e. the distractor or "noise") can be distracting for the driver whose light is red (i.e. the "signal"). The presence of the distractor affects the reaction time and choice of the driver. However the driver can suppress their attention to the green arrow and/or attend to the correct red light in their lane. The decision to stop or accelerate is an example of a perceptual decision. Perceptual decision making is the process of making quick decisions based on objects' features observed with the senses. As shown in the stoplight example, attention is highly influential in the perceptual decision making process. When distracting objects exist in visual space, one must attend only to the relevant objects and actively ignore distracting objects. Each time an individual reaches a stop light, they will be more likely to make a safer decision if they suppress distracting visual input and enhance relevant visual input. The effect of visual attention on the decision-making process will be discussed in Chapters 2 and 3. How individual differences in humans' visual attention processing explain individual differences in decision making will be explored in Chapter 2. How short-term, within-person changes in visual attention processing affect decision making ability will be explored in Chapter 3.

To make a decision to stop or accelerate, the driver must first recognize that a decision must be made by recognizing the visual input (the traffic light). The amount of time this

visual preprocessing (i.e. visual "encoding") takes depends on factors such as the amount of distracting stimuli and the arousal of the driver. The driver must then make a decision by accumulating evidence, the amount of time depending upon the visual attention and ability of the driver. Finally the driver must make a motor response, and the amount of motor preparation required and physical action (e.g. depressing the brake pedal) all take time. While all cognitive processing will happen in a matter of milliseconds, those milliseconds are important for the safety of the driver. The chronometry (i.e. the estimation of timing) of decision making will be discussed in Chapters 4 and 5. Chapter 4 will cover the possibility of tracking the within-person timing of visual encoding with evoked brain activity. Chapter 5 will explore the possibility of tracking within-person motor preparation time with evoked brain activity.

1.2 Background

1.2.1 Electroencephalography (EEG)

Electroencephalography (EEG) is the measurement of the electric potentials on the scalp surface generated (in part) by neural activity originating from the brain. EEG is sensitive to changes in brain activity on a millisecond time scale and is thus an convenient tool when studying quick decision making. This is a major advantage of EEG over other brain imaging modalities such as functional magnetic resonance imaging (fMRI) or near-infrared spectroscopy (NIRS) that operate on time scales in the seconds to minutes range. Over the past 100 years, neuroscientists and clinical neurologists have made use of EEG to obtain insight into cognitive or clinical disease state by applying a variety of signal processing and statistical analyses to EEG time series (for an extended introduction to EEG for a general audience see Nunez et al., 2016). In this thesis, I make use of EEG to directly measure

potentials from the human cortex in order to draw inference about cognitive events. In particular this thesis will be focused on *evoked* EEG. That is, EEG time-locked to certain events or visual stimuli.

1.2.2 Behavioral models of decision making

Drift-diffusion models are a widely-used class of models used to jointly predict humans' choices and reaction times (RT) during two-choice decision making (Link and Heath, 1975; Ratcliff, 1978; Ratcliff and McKoon, 2008). This family of models has been useful in explaining between- and within-participant variability in simple decision making experiments (Vandekerckhove and Tuerlinckx, 2008; Vandekerckhove et al., 2011). "Neural" drift-diffusion models have also successfully incorporated functional magnetic resonance imaging (fMRI) and EEG recordings into hierarchical models of choice-RT (e.g. Mulder et al., 2014; Turner et al., 2015; Nunez et al., 2015, 2017; van Ravenzwaaij et al., 2017). While other similar models of choice-RT have successfully predicted behavior during visual decision making, such as the simpler linear ballistic accumulator model (Brown and Heathcote, 2008) or a more complicated drift-diffusion model that intrinsically accounts for trial-to-trial variability of cognitive parameters (Ratcliff, 1978; Ratcliff and McKoon, 2008), in this work a diffusion model was chosen that allows for tests of specific predictions from models of attention (i.e. Smith and Ratcliff, 2009; Lu and Dosher, 1998) and specific cognitive interpretations of evoked-EEG measures, while being simple-enough to fit with reasonable computational demands given hierarchical forms of the model (discussed below).

Diffusion models add to the analyses of humans' behavior by assuming underlying cognitive processes which have some empirical validation (Voss et al., 2004). In the drift-diffusion model it is assumed that humans accumulate evidence for one choice over another (usually modeled as a correct versus incorrect choice, as in this work, without loss of generality) in a

random walk evidence accumulation process with an infinitesimal time step until sufficient evidence is accumulated to exceed the threshold for one of the two choices. That is, evidence E_t accumulates following a Wiener process (i.e. Brownian motion) with drift rate δ and instantaneous variance ς^2 (Ross, 2014) such that

$$\frac{dE_t}{dt} \sim \mathcal{N}(\delta, \varsigma^2) \tag{1.1}$$

Thus the drift rate δ describes mean rate of evidence accumulation within a trial and the diffusion coefficient ς influences the variance of evidence accumulation within one trial, with the true variance of the current evidence at any particular time t being $\varsigma^2 t$. This process stops once α relative evidence is accumulated for one choice over another. While neural coding may be more sequential in nature, the infinitesimal approximation should hold true for small time steps. Graphical representations of the diffusion model are provided in **Figure 1.1** and **Figure 4.7**.

A few additional parameters describe the simple diffusion model of speeded choice and reaction time that is discussed in this work. The boundary separation α is equal to the amount of relative evidence required to make choice A over choice B (or make a correct decision over an incorrect decision). The boundary separation has been shown to be manipulated by speed vs. accuracy strategy trade-offs (Ratcliff et al., 2001; Voss et al., 2004). The non-decision time τ is equal to the amount of time within the reaction time of each trial that is not dedicated to the decision making process. Typically non-decision time is assumed to be equal to the sum of preprocessing time ("encoding" time) before the evidence accumulation process and motor response time after a decision has been reached. The relative contribution of these two non-decision times is not identifiable from behavior alone and therefore rarely explicitly modeled. However their identification with EEG is discussed in Chapters 4 and

5. A parameter β reflects the starting position of evidence, i.e. the bias towards one choice over another. However the bias parameter is often assumed to be equal to .5 reflecting no bias when modeling correct versus incorrect choices, as in this work.

All three of the parameters related to evidence accumulation are not identifiable with behavioral data alone (i.e. drift rate δ, the diffusion coefficient ς, and the boundary separation α). Only two of the three parameters can be assumed to vary across subjects and trials (e.g. multiplying ς by two and dividing both α and δ by two would result in the same fit of choice-RT) (Ratcliff and McKoon, 2008; Wabersich and Vandekerckhove, 2014). Previous studies have typically chosen to fix the diffusion coefficient ς to 1 or 0.1 (Vandekerckhove et al., 2011; Wabersich and Vandekerckhove, 2014). However due to the predictions made by Dosher and Lu (2000b), in that internal noise is suppressed by attention to the signal, we choose to leave ς to vary. In studies covered in Chapters 2 and 3, the boundary separation α was fixed at 1 for all trials and subjects. In Chapter 2, individual-differences in internal noise as measured by ς was explicitly explored and therefore α was fixed to 1. In Chapter 3, our primary analysis focused on the trial-to-trial variability in the evidence accumulation process due to fluctuations in attention from trial-to-trial within individuals. Although trial-to-trial speed-accuracy trade-offs can be experimentally introduced to find neural correlates of the boundary separation (e.g. van Maanen et al., 2011) or may exist due to per-trial performance feedback (Dutilh et al., 2012), we have no reason to believe that the boundary separation will vary considerably from trial-to-trial within a subject due to changes in *attention*. In Chapters 3 and 4, the majority of analyses fixed ς to 1 and explored condition-level and individual differences in α, δ, and τ.

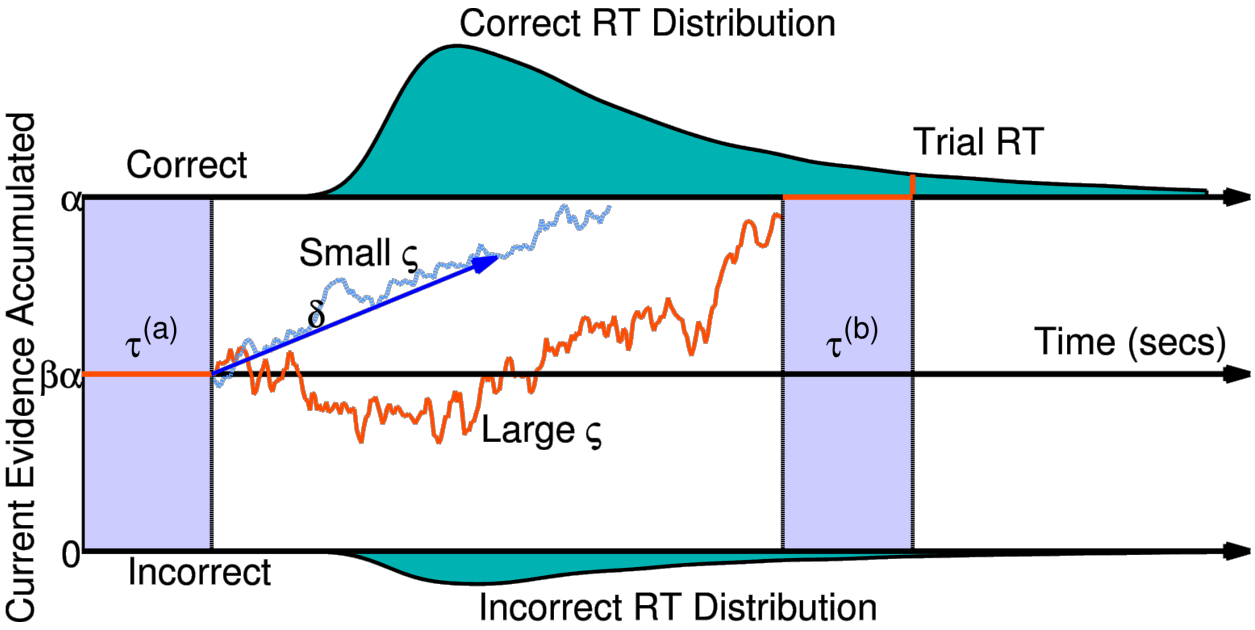

Figure 1.1: A visual representation of the diffusion model. The orange line represents the participant's stochastic evidence accumulation process during one trial. When a participant accumulates enough evidence over time for a correct or incorrect response (graphically represented by the top and bottom boundaries at 0 and α respectively) a decision is made. The drift rate δ is the mean rate of evidence accumulation (evidence units per second) during the participant's decision time on one trial. The bias parameter β represents a bias of the participant towards one choice or the other (set to $\frac{1}{2}$ when the model parameters are expressed in terms of correct over incorrect evidence instead of choice A over choice B evidence). The non-decision time τ is the portion of the participant's reaction time (RT) during the trial not associated with decision making, equal to the sum of encoding/preprocessing time $\tau^{(a)}$ and motor response time $\tau^{(b)}$, which are not separable with behavioral data alone. The boundary separation parameter α represents the amount of relative evidence needed to make a decision. Another parameter is the variability in the evidence accumulation process, the diffusion coefficient ς. A larger ς indicates that a participant would have closer to chance performance (i.e., accuracy of β) and more variable reaction time distributions. In this trial the diffusion coefficient is large in comparison to a smaller diffusion coefficient as shown by the light blue dashed line. The teal shaded areas represent the correct (top) and incorrect (bottom) reaction time distributions. In this example the systematic component of the decision making process is positive $\delta > 0$ indicating a mean trend towards correct responses. However, incorrect responses can still be reached due to the random component of the evidence accumulation process.

1.2.3 Hierarchical Bayesian Models

Discovering the relationship of EEG data and cognitive parameters in models of human behavior yields EEG measurements that are psychologically interpretable. EEG data can also provide new and additional information about the cognitive process that cannot be discerned with just behavior alone. Thus a flexible framework is needed for building and testing theoretical models of the relationship of electrical observations from the human cortex (EEG), human cognition, and human behavior. The hierarchical Bayesian modeling framework is ideally suited for this joint analysis of multiple modes of data (Lee, 2011; Turner et al., 2013). Bayesian inference refers to underlying probability theory and methods used to obtain conclusions about data (for an entertaining introduction to Bayesian inference see Etz and Vandekerckhove, 2017).

Hierarchical modeling refers to the mathematical procedure of assuming statistical relationships between multiple levels of data description, and recent advances in mathematical psychology have introduced hierarchical Bayesian versions of cognitive models (Rouder et al., 2005; Vandekerckhove et al., 2011; Lee and Wagenmakers, 2014). Hierarchical modeling often yields better estimates of parameters due to shrinkage, a phenomenon whereby parameters are better estimated (and data are better described) because hierarchical relationships enforce similarity across similar parameters. For instance, the condition-level mean accuracies in an experimental task could statistically describe observed subject-level mean task accuracies though a normal distribution, and yield more predictive estimates of future subject-level ability. Hierarchical Bayesian modeling also allows discovering complex relationships between multiple data types within cognitive neuroscience (see Turner et al., 2016) by allowing simultaneous estimation of posterior distributions of multiple parameters. Fitting procedures produce samples from probability distributions that display knowledge (i.e. "uncertainty") about parameter estimates and thus certainty about the effects of cognition or EEG data in specific theoretical models. While relating neural data to human behavior

can be performed in multiple ways, with multiple steps of analysis (see Turner et al., 2017, for a review of these methods), the method used predominantly in the EEG domain so far has been the joint modeling approach. The joint modeling approach relates cognitive parameters of behavioral models to either summary measures of EEG or cognitive parameters derived from EEG measures.

1.3 Combining EEG and modeling techniques to discern human cognition

1.3.1 Single-trial analysis

Single-trial cognitive parameter estimation is often not possible when summarizing cognitive parameters from behavioral data distributions over many experimental trials, such as in drift-diffusion models. But task related cortical activity as measured by EEG have been used in hierarchical Bayesian modeling to find estimates of cognitive parameters on single-trials and to draw inference about the function of the electrocortical activity itself (e.g. Frank et al., 2015; Nunez et al., 2017). Single-trial estimates of evidence threshold in quick decision making tasks can be found by using theta-band power (4 to 8 Hz) of the prefrontal cortex, thought to be a measure of strategic control (Cavanagh et al., 2011; Frank et al., 2015). On trials when theta-band power of the frontal cortex is increased, it is expected that those trials contain longer decision-making times due to a willingness by the subject to be less fast but more accurate. This is thought to be a marker of a top-down cognitive control of strategy during decision making. Corroborating evidence of frontal-lobe decision strategy control was found by (Herz et al., 2017). Herz et al. (2017) found that theta-band (2-8 Hz) local field oscillations in the subthalamic nucleus (STN; coupled to theta-band activity in the frontal EEG electrode Fz) also explained trial-to-trial differences in decision strategy. It

is expected that future work will yield trial-to-trial estimates of the decision strategy (e.g. boundary separation α) parameter as assumed by drift-diffusion models of decision making by obtaining trial-to-trial power estimates of the frontal theta-band.

Single-trial estimates of evidence accumulation rate during quick decision making and non-decision time (time in milliseconds of a human reaction time not related to a decision) can also be obtained using hierarchical Bayesian modeling with event-related potentials (ERP) estimates on single trials, time-locked to the onset of visual stimuli. Such work is described in Chapter 3. It was found that ERP measures described trial-to-trial differences in visual encoding time (a component of non-decision time during reaction time) and trial-to-trial differences in evidence accumulation rate, as described by trial-level estimates of the drift parameter. However stronger EEG correlates of cognitive parameters should be found in order to better describe variation in reaction times across trials. A growing body of work indicates that rising amplitudes over parietal cortex during speeded decision-making may describe the evidence accumulation process itself (e.g. O'Connell et al., 2012; Twomey et al., 2015; van Ravenzwaaij et al., 2017; Pisauro et al., 2017). Therefore single-trial estimates of such amplitudes (such as single-trial estimates of P300 ERP potentials) would be beneficial for understanding how cognitive processing on particular trials operates within a subject during decision-making.

1.3.2 Subject-level analysis

Sometimes there is reason to believe that certain EEG measures are strong correlates of a particular cognitive process. These EEG measures can inform the cognitive model. EEG correlates of additional cognitive processes, such as visual attention, add inference about the overall human cognitive process when used in combination of behavioral modeling. For example, steady-state visual evoked potentials (SSVEPs) are the product of a frequency-

tagging paradigm where EEG amplitude is measured at the temporal frequency of a flickering visual stimulus, such that SSVEPs can be used to measure visual attention. Nunez et al. (2015, ; Chapter 2) found that individual differences in SSVEPs predicted individual differences in drift rates and diffusion coefficients (a measure of within-trial evidence accumulation variance) across subjects which in turn explained individual differences in reaction time distributions and accuracy measures. That is, in Chapter 2 it is found that differences in experimental participants attention (both visual noise suppression and visual signal enhancement) as measured by SSVEPs related to some specific differences in participants cognition during decision-making.

Another approach to combining EEG measures is the direct input approach as described by (Turner et al., 2017). The direct input approach would replace a cognitive parameter in a Bayesian model directly with an observed EEG measure. While no published material was found using this technique with EEG and hierarchical Bayesian methods, single-neuron data has been embedded in models of perceptual decision making to estimate the evidence accumulation process (e.g. Palmeri et al., 2015). The extension to EEG and hierarchical Bayesian methods is a simple extension and may be expanded in the future if solid theoretical ground exists for its use. O'Connell et al. (2012) discovered a parietal EEG measure time-locked to the response, thought to track the evidence accumulation process during perceptual decision making. If verified, this EEG measure could be used to discern the rate of evidence accumulation per subject, freeing estimation of the evidence boundary and within-trial variance parameter often ignored in the literature (by fixing the parameter to a certain value). In addition, initial work in Chapter 4 suggests that visual encoding time may be separable from decision and motor preparation time, traditionally unidentifiable with behavior alone, using fixed relationships between early evoked potentials time-locked to visual stimuli and visual encoding time.

1.3.3 Words of caution

Caveats to the approach of embedding EEG in hierarchical Bayesian models are given by Hawkins et al. (2017). Theoretically informed modeling remains important because the cognitive correlates in observed EEG measures are often not clear. EEG measures must have low measurement noise variance, otherwise many trials are needed to differentiate cognitive models. EEG is also very prone to muscle artifact influence in certain frequency bands, so optimal artifact correction is important before inputting EEG measures into hierarchical Bayesian models (see Section 7 of Nunez et al., 2016). Finally, theoretical generative models of EEG must be developed in the future to learn useful combined generative models of observed cortical dynamics, cognition, and human behavior.

1.3.4 Publicly available software

Modern software allows hierarchical Bayesian models to be easily created, built, and tested with both behavioral and EEG data using multiple types of Markov Chain Monte Carlo sampling techniques. Although still being developed and improved, JAGS (Plummer, 2003), Stan (Carpenter et al., 2016), and PyMC3 (Salvatier et al., 2016) are all recommended tools for these steps. A more specific toolbox is HDDM: a Python software package that can perform linear regression between calculated EEG signals on single-trials (and on the condition- and subject-levels) and parameters of drift-diffusion models of accuracy and reaction time data (Wiecki et al., 2013). Hawkins et al. (2017) created R code and examples for sampling from hierarchical drift-diffusion models (HDDM) with neural inputs and single-trial neural regressors, located here: `https://osf.io/ws3fn`. MATLAB and JAGS example code performing some hierarchical drift-diffusion model analyses with single-trial neural inputs can be found at `https://github.com/mdnunez/mcntoolbox/` (Nunez et al., 2017). While the current literature and software are biased towards models of quick decision-making. The

tools needed to extend these methods to models of other types of human cognition combined with EEG is readily available.

Chapter 2

Individual differences in attention influence decision making

2.1 Introduction

The joint analysis of physiological and behavioral data has been a topic of recent interest. In a string of publications, a number of research groups (Cassey et al., 2014; Forstmann et al., 2010; Turner et al., 2013) have presented work in which neurophysiological data are linked to parameters of cognitive or behavioral process models (see also Palmeri et al., 2017). The goal of these modeling exercises is not only to evaluate the predictive power of brain activity for behavior, but also to elucidate the nature of this prediction. The use of cognitive models with neural data and cognitive parameters permits more psychologically interpretable labeling of the neurophysiological measurements, providing links between brain activity, cognition, and behavior.

In the present paper, we apply a cognitive model constrained by EEG data to fit accuracy and response times of multiple individuals from a perceptual decision making task. The

goal of the model fit is twofold: 1) to demonstrate the superior generalizability of such a model as compared to model variants without neural input components and 2) to evaluate the hypothesis that individual differences in enhancement or suppression of visual attention, as measured by EEG, contribute to individual differences in cognition and thus to individual differences in accuracy and/or reaction time in the task.

In order to show out-of-sample generalizability, we first fit the model to a training set of participants and obtain the requisite (population-level) linking parameters, and then make predictions about the behavior of a new participant to which the model was not trained. In the sections that follow, we will describe 1) the cognitive process model that we have chosen, 2) the task to which it is applied and the EEG data that we collected, 3) a series of three models of increasing complexity, of which the model with external attentional EEG covariates is the most complex, 4) the results of the generalization exercise and 5) evaluation of the hypothesis.

2.1.1 Steady-state visual evoked potentials as a measure of attention

In this study, we will demonstrate how attentional mechanisms can explain individual differences in perceptual decision making as estimated by a cognitive model. In a typical visual attention experiment, the signal stimulus is attended and preferentially processed while competing stimuli (i.e., visual noise) are not further processed. A number of studies have demonstrated that a measure of the deployment of attention can be obtained by using flickering stimuli and electroencephalographic (EEG) recordings of the (frequency tagged) steady-state visual evoked potentials (SSVEPs) (Morgan et al., 1996; Müller et al., 1998; Ding et al., 2006; Bridwell and Srinivasan, 2012; Garcia et al., 2013). SSVEPs are narrow band responses at the visual flicker frequencies and flicker harmonics of a stimulus (Regan,

1977). When a stimulus is attended, the SSVEP is enhanced, and when a stimulus is not attended or suppressed, the SSVEP is diminished. This approach has been used to investigate individual differences in attention strategy in detection and discrimination tasks. Bridwell et al. (2013) found that only a subset of participants could deploy the optimal attention strategy and modify their strategy by the task demands. An SSVEP approach has also been used to show that individuals are trained by their own experiences. Individuals with attentional training due to a history of fast-action video gaming have been found to preferentially suppress noise rather than enhance the signal, and those individuals performed better at vigilance tasks (Krishnan et al., 2013).

2.1.2 The case for hierarchical Bayesian models

The hierarchical Bayesian process modeling framework is ideally suited for the joint analysis of multiple modes of data—Turner et al. (2013) describe three such joint modeling strategies and Vandekerckhove (2014) describes a fourth. One strategy afforded by hierarchical Bayesian models involves constraining the estimation of cognitive process models by introducing the brain data as (fixed) covariate information. This strategy carries the disadvantage that it does not by default allow for measurement variance on the neurophysiological side, but has the advantage of being relatively straightforward to implement in a computationally efficient fashion. By conditioning the estimation of the cognitive parameters on brain data (or other external covariates), it is expected that unexplained variability between participants can be reduced, and consequently that such a model should perform better in generalization tests.

Interindividual variability (i.e. variability in the participant-level cognitive parameters; changes over subscript j) in diffusion models has been previously analyzed by fitting a diffusion model to each participant individually then comparing parameters across model

fits. The individual differences were then gauged by statistical analyses on the models' resulting maximum likelihood parameter estimates (Ratcliff et al., 2001; Wagenmakers et al., 2008). Some limitations to this technique are that large sample sizes are needed for diffusion model parameter estimation, that shared condition-level differences across individuals cannot be easily evaluated (Wagenmakers, 2009; Vandekerckhove et al., 2011), and that statistical uncertainty is not propagated across stages of the analysis. Hierarchical Bayesian methods along with Monte Carlo sampling techniques allow for the estimation of complex models. These methods have been used to explain individual differences in the diffusion model and other cognitive models without the need for large sample sizes (Lee, 2008; Lee and Newell, 2011; Vandekerckhove et al., 2011). Additionally, the hierarchical framework allows for between-participant variability to be explained when each participant's diffusion model parameters are functionally related to known exogenous data (e.g., physiological data).

2.1.3 Constraining model parameters with EEG data

We assume that brain activity compels cognition, which in turn drives participant behavior. Assuming attention constrains one or more of the cognitive processes in perceptual decision making, then as a consequence of attentional mechanisms we expect SSVEPs to help explain between-participant variability in the parameters of the diffusion model and thus between-participant variability in RT and accuracy. In one study, an occipital SSVEP amplitude was shown to track visual sensory evidence over the time course of a trial, suggesting that SSVEPs can reflect the evidence accumulation process itself (O'Connell et al., 2012). The experimental stimulus used in this study involves a flickering signal overlayed on time-varying visual noise, designed to evoke separate SSVEP responses to the signal and the visual noise, which we expect will explain individual differences in the model parameters and behavior.

We hypothesize increased within-trial evidence accumulation rates, reflected by increased

drift rates, for those subjects who suppressed attention to the visual noise. We further hypothesize that another benefit of attention for RT and accuracy is a result of reduced within-trial variability in the accumulation of evidence. Thus, we predict an across-individuals relationship between enhanced attention to the signal and decreased diffusion coefficients.

As mentioned above, one of the parameters of the diffusion model must be fixed rather than estimated (either diffusion coefficient ς, drift rate δ, or boundary separation α). For the present study a variable boundary separation across conditions is not a valid interpretation of the data since the changes between conditions occur unannounced, leaving the participant with no opportunity to adapt strategies (e.g., switch between a speed or accuracy strategy) in response to stimulus changes. In our parameterization, we leave the diffusion coefficient ς free to vary, set α to one evidence unit, and assume no bias ($\beta = \frac{1}{2}$) towards correct responses. The joint density f of RT t and accuracy w of this simplified diffusion model is given in Equation 2.1. The density is derived from the limiting approximation given by Ratcliff (1978) where $\alpha = 1$ and $z = 0.5$.

$$\begin{cases} f(t, w = 0 \mid \varsigma^2, \tau, \delta) = \pi\varsigma^2 e^{-\frac{1}{2}\left[\frac{\delta}{\varsigma^2} + \frac{\delta^2}{\varsigma^2}(t-\tau)\right]} \sum_{k=1}^{+\infty} \left[k \sin\left(\frac{1}{2}\pi k\right) e^{-\frac{1}{2}k^2\pi^2\varsigma^2(t-\tau)}\right] \\ \\ f(t, w = 1 \mid \varsigma^2, \tau, \delta) = f(t, w = 0 \mid \varsigma^2, \tau, -\delta) \end{cases} \tag{2.1}$$

In what follows, we will use the effect of attention, as measured by SSVEPs, to constrain diffusion model parameter estimates (in our case δ_j, ς_j, and τ_j). In particular, we assume that, on each trial, a participant's attention is reflected in phase locking (i.e. SSVEPs) to the attended visual signal and decreased phase locking to the unattended visual noise.

We will demonstrate that the hierarchical Bayesian SSVEP-driven diffusion model has pre-

dictive ability as well as descriptive ability—more specifically, that our ability to predict each participant's accuracy and RT behavior is improved by including the SSVEP measures of attention processes.

2.2 Material & Methods

2.2.1 Participants

The following study was approved by the University of California, Irvine Institutional Review Board and was performed in accordance with APA standards. Informed consent was obtained from each of the seventeen participants (8 females and 9 males) who took part in the study. The mean age of sixteen of the participants was 25 with an age range of 21 to 30. Another participant was over 45 years of age. Sixteen participants self identified as being right handed while another identified as being left or ambidextrous. All participants had at least 20/30 vision or corrected vision as measured by a visual acuity chart available on the internet (Olitsky et al., 2013). No participants reported any history of neurological disorder. Each participant completed the experiment in one session within 2.5 hours.

2.2.2 Experimental Stimulus

The participants were given a two-alternative forced-choice perceptual decision making task in which they were asked to differentiate the mean rotation of bars within a circular field of bars that deviated randomly from mean rotation. One half of the trials had a mean bar rotation of 45° while the other half had a mean rotation of 135°. The bar field was flickered against a time-varying noise pattern.

The participants viewed each trial of the experimental stimulus on a monitor in a dark room.

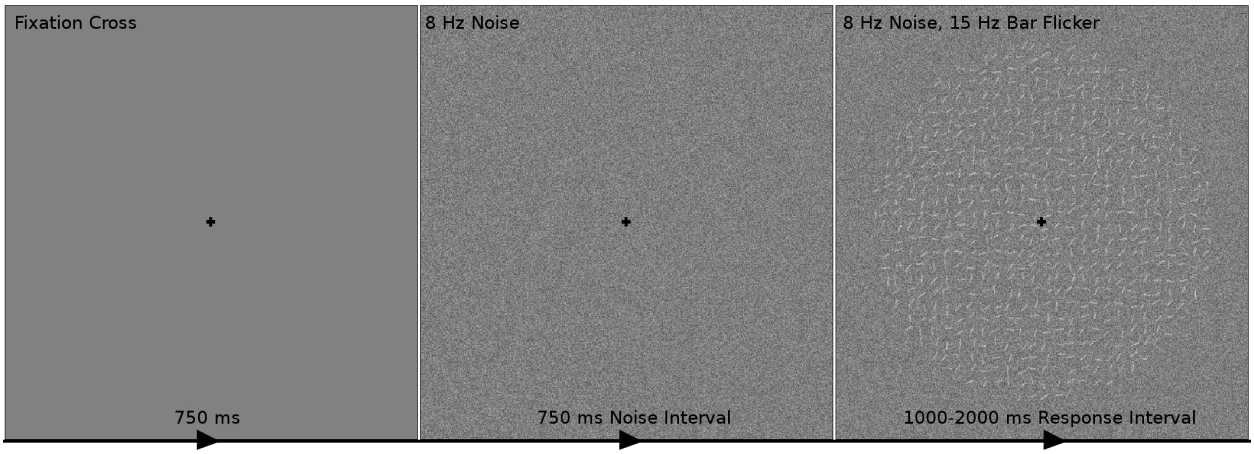

Figure 2.1: The time course of one trial of the experimental stimulus. The participant first fixated on a black cross for 750 ms indicating the beginning of a trial. The participant then observed visual contrast noise changing at 8 Hz for 750 ms while maintaining fixation. A circular field of small oriented bars flickering at 15 Hz overlaid on the changing visual noise was then shown to the participant for 1000 to 2000 ms. The task was to indicate during this response interval whether the bars were on average oriented towards the "top-right" (45° from the horizontal line; as in this example) or "top-left" (135°) corners. It was assumed that the participant's decision making process began at the start of the response interval. After the response interval, the fixation cross was shown in isolation for 250 ms to alert the participant that the trial was over and to collect remaining responses.

The time course of one trial is shown in **Figure 2.1**. Participants were positioned such that the entire circular field of small oriented bars had a visual angle of 9.5°. Within each trial the participant first saw a black cross for 750 ms in the middle of the screen on which they were instructed to maintain fixation throughout the trial. The participant then observed visual contrast noise changing at 8 Hz for 750 ms; this time period of the trial will be referred to later in this paper as the *noise interval*. The participant then observed a circular field of small oriented bars flickering at 15 Hz overlaid on the square field of visual noise pattern changing at 8 Hz and responded during this time frame, henceforth referred to as the *response interval*. The visual noise and bar field are modulated at constant rates (8 and 15 Hz respectively) to evoke frequency-tagged signal and noise responses in the cortex which we measured as steady-state visual evoked potentials (SSVEPs). The SSVEP responses at the signal frequencies (15 Hz and its harmonics) and at the contrast noise frequencies (8 Hz and its harmonics) were used to measure the effect of attention to the signal stimulus and noise stimulus. The display time of the response interval was sampled between 1000 and 2000 ms from a uniform distribution. After this display period the black fixation cross was shown in isolation for 250 ms to alert the participant the trial was over and to collect any delayed responses.

Three levels of variance of bar rotation and three levels of contrast noise were used to modulate the task difficulty. In the first level of bar rotation variance, each bar was drawn from a uniform $\mathcal{U}(-30°, 30°)$ distribution centered on the mean angle. In the two other levels, the rotations of each bar were drawn from $\mathcal{U}(-35°, 35°)$ and $\mathcal{U}(-40°, 40°)$ respectively. The three levels of contrast noise were 30% contrast noise, 45% contrast noise and 60% contrast noise. The 30% contrast noise condition was obtained by the addition of a random draw from a $\mathcal{U}(-15\%, 15\%)$ distribution to the luminance of each pixel in a square field. Baseline luminance was 50%. The other contrast noise conditions were obtained similarly. Each participant was shown 90 trials from each bar rotation-noise condition combination.

The bar rotation (BR) variance manipulation was hypothesized to modulate each participant's diffusion coefficient since the participant would have more variable information in harder trials. Considering each bar's rotation as a unit of information contributing to a 'left' or 'right' response, information would be more variable in trials that sampled the bar rotations from wider uniform distributions. It was thought that contrast noise would degrade the amount of information each bar gave to the decision process thus leading to smaller drift rates in trials with higher noise contrast.

2.2.3 Behavior and EEG Collection

Participants first completed a training session of 36 trials each. Participants were asked to complete a second training set if their percentage accuracy was subjectively judged by the experimenter to not converge to a stable value. Each participant completed 6 blocks of 90 trials each for a total of 540 trials with breaks between each block of variable time. Each trial lasted randomly (uniformly) from 2.75 to 3.75 seconds. Participants were asked to respond during the 1 to 2 second response interval as accurately as possible, with no-answer trials considered as incorrect. To maintain participant performance, auditory feedback was given after the response interval to the alert the participant if they were correct or incorrect. Performance feedback was also provided between blocks by displaying on the screen the percentage of trials answered correctly in that block. The behavioral data consists of each participant's accuracy and reaction time during each trial.

High-density electroencephalography (EEG) was collected using Electrical Geodesics, Inc.'s 128-channel Geodesic Sensor Net and Advanced Neuro Technology's amplifier with electrodes sitting on the participant's scalp throughout the duration of the experiment. Electrical activity from the scalp was recorded at a sampling rate of 1024 samples per second with an online average reference using Advanced Neuro Technology's digitization software. The

EEG data was then imported into MATLAB for offline analysis.

Linear trends were removed from the EEG data. As we were only interested in 1 to 50 Hz EEG, the following filters were applied to each channel: 1) A high pass Butterworth filter with a 1 Hz pass band with 1 dB ripple and 0.25 Hz stop band with 10 dB attenuation, 2) a stopband Butterworth filter with 59 and 61 Hz pass bands with 1 dB ripple and 59.9 to 60.1 Hz stop band with 10 dB attenuation (to remove power-line noise), and 3) a low pass Butterworth filter with a 50 Hz pass band with 1 dB ripple and 60 Hz stop band with 10 dB attenuation. Artifactual data thought to be generated by phenomena outside of the cortex were removed from the EEG data using a paradigm involving Independent Component Analysis (ICA): First, any trials or channels were rejected that had time-courses unusual for cortical activity and/or had properties that ICA is deemed to not extract well, such as trials with high frequency activity indicative of muscle activity, trials or channels with high 60 Hz amplitude indicative of power-line noise suggesting poor electrode-to-skin connection, or trials with sudden high amplitude peaks that cannot be generated by cortical activity (Delorme et al., 2007). Second, ICA was used to remove linear mixtures of channel time-courses that did not subjectively correspond to EEG data in spatial map on the scalp, in power spectrum, and/or in event-related potential (ERP). Typical artifactual components include: those components with spatial maps of highly weighted electrodes near the eyes suggestive of eye movements, those components with high amplitudes at high frequencies and low amplitudes at low frequencies suggestive of muscle activity, and spatial maps of highly weighted singular electrodes suggestive of poor electrode-scalp connectivity. A final cleaning step was performed by rejecting any trials that had high amplitudes not typical of cortical electrical activity.

For each participant, steady-state visual evoked potentials (SSVEPs) to the visual noise and signal (the circular bar field) were found at each electrode. In this experiment a steady-state response was defined by the consistency in phase at the frequencies of the stimulus

(8 and 15 Hz) and the harmonic frequencies of the stimulus (16, 24, 32, 40, 48, 30 and 45 Hz). The uniformity of phase across trials was measured by the Phase Locking Index (PLI) across trials. The PLI is a statistical characterization of phase synchronization resulting from an experimental stimulus and has been shown to be successful in characterizing cortical signals (Rosenblum et al., 1996; Sazonov et al., 2009). The phase locking index ignores signal amplitude and ranges from 0 (all trials out-of-phase) to 1 (all trials in-phase) (Tallon-Baudry et al., 1996). The equation used for PLI is provided in Equation 2.2. PLI is the average of ≈ 540 trials of amplitude normalized Fourier coefficients of the time interval. For each electrode e and participant j, PLI is defined as a function of frequency f.

$$\mathrm{PLI}_{ej}(f) = \left| \frac{1}{540} \sum_{i=1}^{540} \frac{F_{iej}(f)}{\left|F_{iej}(f)\right|} \right| \tag{2.2}$$

The steady-state responses to the visual noise were analyzed based on both the 750 ms noise interval and the first 1000 ms of the response interval while the steady-state responses to the signal were analyzed based only on the first 1000 ms of the response interval. Because steady-state responses located in parietal electrodes have been successfully related to attentional mechanisms in past studies (Ding et al., 2006; Bridwell and Srinivasan, 2012), electrical activity at parietal electrodes was hypothesized to be most descriptive of cognitive processes in the visual decision making task. The subject mean PLI at all frequencies averaged over parietal channels is shown in **Figure 2.2**. Topographic maps of the distribution of the PLI are shown at the fundamental and first two harmonics for signal and noise frequencies. It is clear that the SSVEP is broadly distributed over frontal, parietal, and occipital networks, as has been found in other studies (Ding et al., 2006; Bridwell and Srinivasan, 2012; Krishnan et al., 2013). The mean PLIs over prefrontal, frontal, central, parietal, and occipital electrode groups for each of the evoked frequencies were used as predictors in the model.

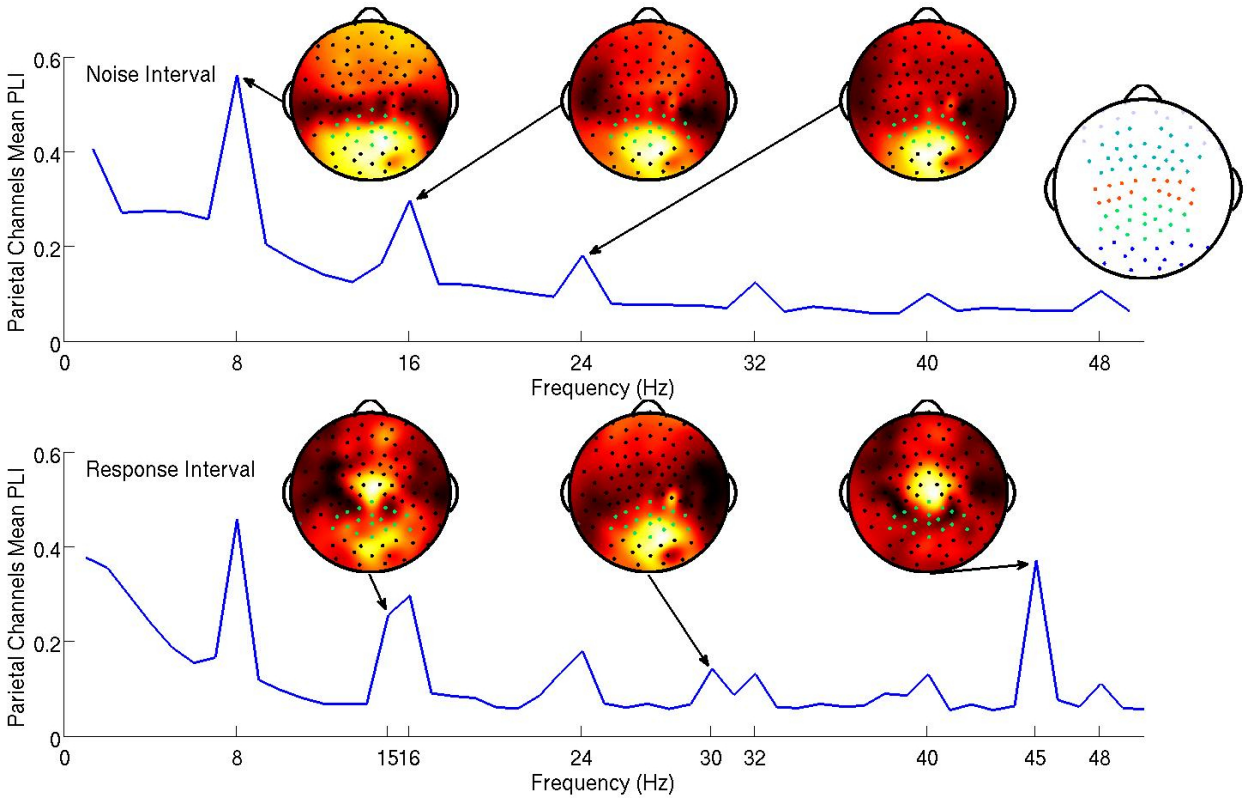

Figure 2.2: The subject mean parietal channel PLI at all frequencies in the noise interval (top panel) and the response interval (bottom). The resolution of the the PLI spectra in the top plot is approximately 1.3 Hz due to the PLI being a function of Fourier transforms of 750 ms epochs. The resolution of the PLI spectra in the bottom plot is 1 Hz as the Fourier transforms are of 1000 ms epochs. The 15 and 16 steady-state responses during the response interval are separable when using 1000 ms epochs. Also shown are subject mean PLI topographies (at 8, 16, and 24 Hz during the noise interval and 15, 30, and 45 Hz during the response interval, each on a standardized scale) indicating where the maximum subject mean PLI is located on the scalp in relation to the parietal electrodes (highlighted green). It is clear from these topographies that using only parietal electrodes will not capture all of steady-state response information. An index of electrode locations is also provided in the top right. Prefrontal, frontal, central, parietal, and occipital electrode groups are colored light blue, teal, orange, green, and blue respectively.

We expect the evoked cortical networks to change dependent upon the flicker frequencies of the stimulus (Ding et al., 2006; Bridwell and Srinivasan, 2012), as shown by the stimulus response in **Figure 2.2** where the spatial distributions of the fundamental and harmonic responses are quite different. However we do not expect the behavior of these harmonics to be uncorrelated. To avoid multicollinearity, we performed two principal components analyses (PCAs; on the noise and signal frequencies separately) to obtain a smaller number of PLI measures from uncorrelated cortical networks. The first PCA reduced 60 PLI variables (5 cortical locations by 6 noise harmonics in both the noise and response intervals) to 16 principal components. The second PCA transformed 15 PLI variables (5 cortical locations by 3 signal harmonics) to 15 principal components. Our criteria for which principal components to include in the hierarchical Bayesian models were 1) based upon the improvement of in-sample predictive power as we increased the number of principal components, resulting in candidate principal components and 2) then based upon the out-of-sample predictive power of the candidate principal components.

2.2.4 Hierarchical Bayesian models

All trials from every participant were used for model fitting except those trials in which there was deemed to be EEG artifact and those trials during which the participant made no response or responded more than once. Since our models do not account for non-decision making trials, exceedingly fast trials (faster than 250 ms) were excluded as well.

The marginal likelihood for the model—that is, the predicted distribution of the data conditional on all parameters—is the first passage time distribution of a Wiener process with constant drift. We call this probability density function the *Wiener distribution*. For each trial i, subject j, and condition k, the observed accuracy w_{ijk} and reaction time t_{ijk} were combined in a two-element vector $\mathbf{y}_{ijk}$. These values were then assumed to be drawn from a

joint distribution:

$$\mathbf{y}_{ijk} \sim \mathcal{W}(\delta_{ijk}, \varsigma_{ijk}, \tau_{ijk}). \tag{2.3}$$

We applied a sequence of three models—each adding a new feature—to the data.

2.2.5 Model 1: No individual differences

We assumed in **Model 1** that all three diffusion model parameters were constant across participants (i.e., that all participants were *identical*), and depended only on the experimental condition k. The diffusion model was fit to the reaction time and accuracy data of all 17 participants under the assumption that all participants had the same drift rate δ_k, diffusion coefficient ς_k, and non-decision time τ_k that were variable across condition k but not variable across participant j. Here k denotes both the particular bar rotation condition and the particular contrast noise condition-level, $k = 1, \ldots, 9$. A graphical representation of **Model 1** is provided in **Figure 2.3(a)**.

The assumptions of the model, together with the prior distributions for the parameters, appear below. The priors for the drift rate δ_k and non-decision time τ_k were truncated normal distributions due to the knowledge of the natural constraints of the diffusion model and prior knowledge of acceptable values for similar tasks. Note that the second parameter of the normal distributions below represent the variance.

$$\begin{aligned}
\delta_{jk} &= \delta_k, \quad \delta_k \sim \mathcal{N}(0.0, 5) \in (-9, 9) \\
\varsigma_{jk} &= \varsigma_k, \quad \varsigma_k \sim \mathcal{N}(0.5, 4) \\
\tau_{jk} &= \tau_k, \quad \tau_k \sim \mathcal{N}(0.3, 4) \in (0, 1)
\end{aligned}$$

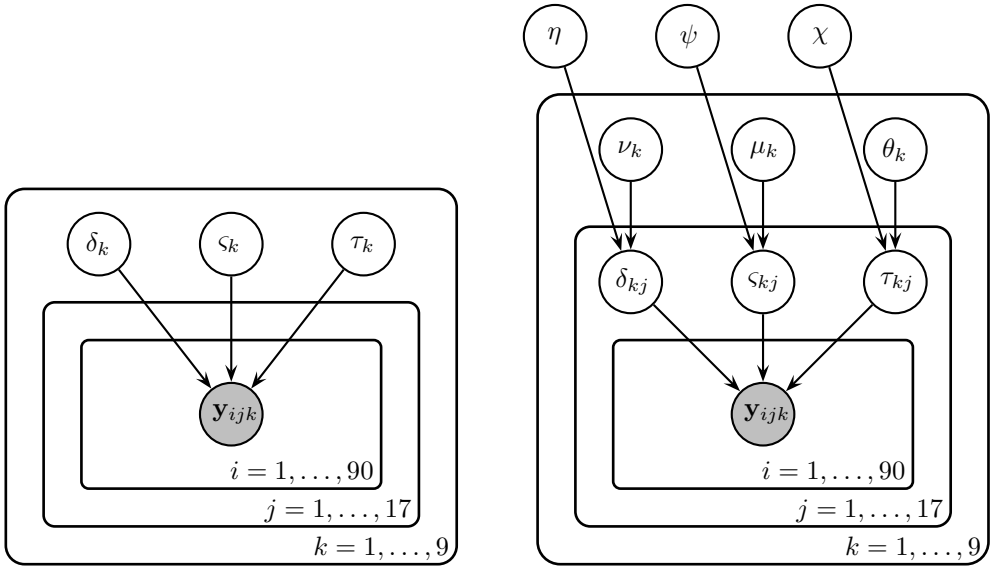

(a) (a) (b) (b)

Figure 2.3: A graphical representation of **Model 1** (a) and **Model 2** (b). In **Model 1**, drift rates δ_k, diffusion coefficients ς_k, and non-decision times τ_k were assumed to vary over conditions k but remain invariant across participants j and trials i. There were three bar rotation conditions and three contrast noise conditions. Here k denotes each bar rotation and contrast noise pair. In **Model 2**, drift rates δ_{jk}, diffusion coefficients ς_{jk}, and non-decision times τ_{jk} were assumed to vary over both conditions and participants. Each of these parameters are in turn assumed to be drawn from normal distributions with means that varied over conditions k and with variances that did not vary across conditions.

2.2.6 Model 2: Individual differences

In **Model 2** we assumed that participants differ but are draws from a single superordinate population (i.e., participants are *exchangeable*). Consequently, the drift rate δ_{jk}, diffusion coefficient ς_{jk}, and non-decision time τ_{jk} varied by both subject j and condition k. Subject-level parameters were assumed to be drawn from normal distributions with means that were variable over condition only. Variances were assumed to be invariant across conditions to maintain model simplicity (i.e., the model assumes *homoskedasticity* in the parameters). The prior distributions of the parameters are listed below.

$$
\begin{aligned}
(\delta_{jk} \mid \nu_k, \eta) &\sim \mathcal{N}(\nu_k, \eta) \in (-9, 9), \quad \nu_k \sim \mathcal{N}(0.0, 5), \quad \eta \sim \Gamma(6, 0.10) \\
(\varsigma_{jk} \mid \mu_k, \psi) &\sim \mathcal{N}(\mu_k, \psi) \qquad\qquad , \quad \mu_k \sim \mathcal{N}(0.5, 4), \quad \psi \sim \Gamma(4, 0.05) \\
(\tau_{jk} \mid \theta_k, \chi) &\sim \mathcal{N}(\theta_k, \chi) \in (0, 1) \quad , \quad \theta_k \sim \mathcal{N}(0.3, 4), \quad \chi \sim \Gamma(5, 0.01)
\end{aligned}
$$

A graphical representation of **Model 2** is provided in **Figure 2.3(b)**.

2.2.7 Model 3: Individual differences with neural correlates

With **Model 3**, we will attempt to explain any individual differences in cognitive parameters by introducing the neural data as explanatory variables. The model is similar to **Model 2**, but additionally includes a regression structure to explain variability in subject-level model parameters with steady-state PLI values.

In order to avoid multicollinearity, PLIs were first subjected to a principal component analysis (PCA), and the resultant independent components were used as predictors. The PCA was performed on the noise and signal frequencies separately. The first PCA reduced 60 PLI variables to 16 principal components and the second PCA transformed 15 PLI variables into 15 components. The criterion used to determine which principal components to include was the out-of-sample predictive power of each model. Predictive power was measured as

R^2_{pred}, a measure of the percentage of total between-subject variance explained, in this case of the correct-RT medians of each condition. The equation used for R^2_{pred} is provided in the appendix.

Subject-level drift rates δ_{jk}, diffusion coefficients ς_{jk}, and non-decision times τ_{jk} were assumed to be drawn from normal distributions with means of the form $\alpha_k + \boldsymbol{x_j^\intercal \gamma}$ where α_k is condition k's effect on the subject-level cognitive parameter, $\boldsymbol{x}_j$ is a vector of principal components, and $\boldsymbol{\gamma}$ is a vector of regression coefficients (i.e., the effect of each principal component on the cognitive parameter). The graphical representation of the model is provided in **Figure 2.4**. The priors of the variance parameters are the same as in **Model 2**. Weakly informative prior distributions of $\mathcal{N}(0.0, 10)$ were given to the weight variables that make up the vectors $\boldsymbol{\gamma_{(\delta)}}$, $\boldsymbol{\gamma_{(\varsigma)}}$, and $\boldsymbol{\gamma_{(\tau)}}$. The other hyperpriors and priors were:

$$
\begin{aligned}
\left(\delta_{jk} \mid \alpha_{(\delta)k}, \boldsymbol{\gamma_{(\delta)}}, \eta\right) &\sim \mathcal{N}\left(\alpha_{(\delta)k} + \boldsymbol{x_j^\intercal \gamma_{(\delta)}}, \eta\right) \in (-9, 9), \quad \alpha_{(\delta)k} \sim \mathcal{N}\left(0.0, 5\right) \\
\left(\varsigma_{jk} \mid \alpha_{(\varsigma)k}, \boldsymbol{\gamma_{(\varsigma)}}, \psi\right) &\sim \mathcal{N}\left(\alpha_{(\varsigma)k} + \boldsymbol{x_j^\intercal \gamma_{(\varsigma)}}, \psi\right) \qquad\quad , \quad \alpha_{(\varsigma)k} \sim \mathcal{N}\left(0.5, 4\right) \\
\left(\tau_{jk} \mid \alpha_{(\tau)k}, \boldsymbol{\gamma_{(\tau)}}, \chi\right) &\sim \mathcal{N}\left(\alpha_{(\tau)k} + \boldsymbol{x_j^\intercal \gamma_{(\tau)}}, \chi\right) \in (0, 1) \ , \quad \alpha_{(\tau)k} \sim \mathcal{N}\left(0.3, 4\right)
\end{aligned}
$$

2.2.8 Posterior sampling

We used the JAGS software (Plummer, 2003) to analyze the data by drawing samples from the joint posterior distribution of the parameters of the hierarchical models. To compute the likelihood function associated with the assumed decision making process (the Wiener distribution), we used the *jags-wiener* module (Wabersich and Vandekerckhove, 2014). This allowed us to explain accuracy and response time distributions within conditions and across subjects. For each model, samples from the posterior distributions of the parameters were found by running JAGS with six Markov Chain Monte Carlo (MCMC) chains of length 21000, with 1000 burn-in (discarded) samples and a thinning parameter of 10 (keeping only every 10th sample) resulting in six joint posterior distribution estimates of 2000 samples

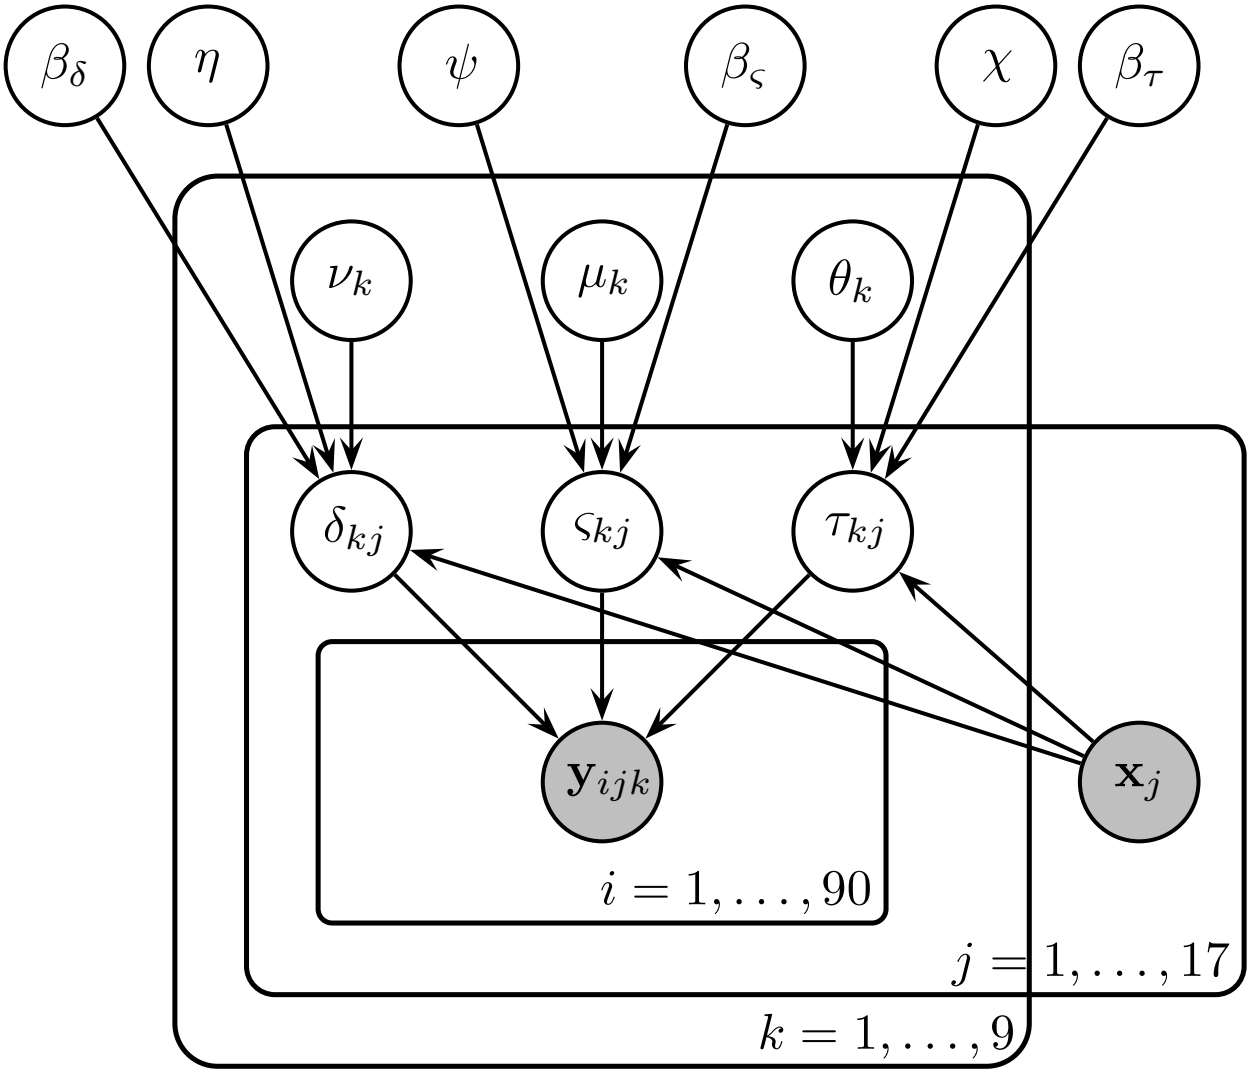

Figure 2.4: Graphical representation of **Model 3**. Drift rates δ_{jk}, diffusion coefficients ς_{jk}, and non-decision times τ_{jk} were assumed to vary over both conditions and participants. Each of these parameters are assumed to be drawn from normal distributions with means of the form $\alpha_k + \boldsymbol{x}_j^\intercal \boldsymbol{\gamma}$, where $\boldsymbol{x}_j$ is the vector of SSVEP responses of subject j, and with variances that did not vary across conditions. As an example, $\alpha_{(\tau)k}$ is the condition effect on the non-decision time and $\gamma_{(\tau)}$ reflects the change in non-decision time (seconds) due to a one SSVEP unit difference across two participants.

each. We used the $\hat{R}$ statistic to compare within-chain variance to between-chain variance in order to assess convergence of the MCMC algorithm (Gelman and Rubin, 1992).

2.2.9 Posterior Predictive Distributions

To quantify model fit, in-sample posterior predictive distributions of accuracy-RTs from 5000 simulated experiments were estimated by sampling from the posterior distributions of subject-level parameters for each of the three models. That is, $s = 1, \ldots, 5000$ samples were randomly drawn from the subject-level posterior distributions of the model parameters producing 5000×1 column vectors for each drift rate $\boldsymbol{\delta}_{jk}^{(s)}$, diffusion coefficient $\boldsymbol{\varsigma}_{jk}^{(s)}$, and non-decision time $\boldsymbol{\tau}_{jk}^{(s)}$. The samples $\left(\boldsymbol{\delta}_{jk}^{(s)}, \boldsymbol{\varsigma}_{jk}^{(s)}, \boldsymbol{\tau}_{jk}^{(s)}\right)$ were used to generate accuracy-RT samples from the Wiener distribution (with the rejection sampling algorithm described in Tuerlinckx et al., 2001).

In order to find candidate PLI predictors for **Model 3** and also to gauge the ability of each model type to predict new subjects' behavioral data, *in-sample* and *out-of-sample* posterior predictive distributions were generated using the PLI coefficients and posterior distributions of the *condition*-level parameters to find predictive distributions of the *subject*-level parameters. This procedure does not use samples from the subject-level posterior distributions directly, but estimates the subject-level parameters from the posteriors of the condition-level parameters and EEG covariates before finding a posterior predictive distribution of accuracy-RTs. Samples from the posterior predictive distribution of subject j's mean drift rate on a trial in condition k are drawn from a normal distribution with mean $\boldsymbol{\alpha}_{(\delta)k}^{(s)} + \boldsymbol{G}_{(\delta)}^{(s)}\boldsymbol{x_2}$ where $\boldsymbol{x_2}$ is the vector of subject j's principal component PLI values, $\boldsymbol{\alpha}_{(\delta)k}^{(s)}$ are samples from the posterior distribution of condition k's effect on drift rate, and $\boldsymbol{G}_{(\delta)}^{(s)}$ is a matrix consisting of samples from the posterior distributions of the PLI coefficients for drift rate. For in-sample prediction, we fit different possible forms of **Model 3**, with different numbers of principal

components, 17 times each to generate in-sample posterior distributions to find candidate principal components. Then for out-of-sample prediction, we fit different possible forms of **Model 3**, with the resulting candidate principal components, 17 times with each participant removed from the data set. In the previously mentioned example, both the condition effect on drift rate and PLI coefficients are estimated from the model with all subjects except j for out-of-sample prediction.

2.3 Results

For all models and all parameters, convergence of the Monte Carlo chains was satisfactory: $\hat{R} \leq 1.01$ for all parameters ($\hat{R} \geq 1.10$ is conventionally taken as evidence for non-convergence; Gelman and Rubin, 1992).

2.3.1 Model 1: No individual differences

Marginal posterior distributions of the parameters of **Model 1** are plotted in **Figure A.1** of the appendix. The variability of evidence units gained per second ς_k increased as bar rotation variance grew. Evidence units gained per second, drift rate δ_k, was found to decrease both with larger contrast noise and larger bar rotation. The parameter estimates seem to show a complex interaction effect of bar rotation and contrast noise on non-decision time τ_k. However, the results from **Model 2** will indicate that **Model 1** is sufficiently misspecified that this interaction cannot be interpreted in a meaningful way.

2.3.2 Model 2: Individual differences

The marginal posterior distributions of the condition-level parameters are shown in **Figure A.1** of the appendix. At the condition level, the effects of the experimental manipulations on drift rate and the diffusion coefficient remain similar to the results of **Model 1**: Mean drift rates ν_k were found to decrease as bar rotation variance grew, smaller mean drift rates were observed in the high visual noise condition, and mean diffusion coefficients μ_k increased as bar rotation variance grew. Main effects on the condition-level non-decision time not clearly observable in **Model 1** were found in **Model 2**. Mean non-decision time θ_k was slow when the bar rotation variance was high, and participants were estimated to have quick non-decision times in low visual noise conditions.

The complex interactive pattern of non-decision times obtained in **Model 1** no longer appears.

By adding subject-level parameters, the current model not only provides a clearer picture of condition-level behavior of all participants, but describes the *individual differences* of the participants modeled by the subject-level parameters, δ_{jk}, ς_{jk}, and τ_{jk}. Posterior distributions for the subject-level parameters of the easiest condition ($\pm 30^\circ$ bar rotation and 30% noise) are provided in **Figure A.2**. Due to subject-level parameters deviating from the condition-level parameter's means, this model is able to predict within-sample data well compared to the previous model. Percent variances explained (R^2_{pred}) of correct-RT subject medians by within-sample posterior prediction are provided in **Table 2.1**. **Model 2** explains at least 86.3% of median correct-RT between-subject variance in each condition.

2.3.3 Model 3: Individual differences with neural correlates

The results of **Model 2** clearly demonstrate differences between participants' cognition in the perceptual decision making task. We were further able to explain the differences in the cognitive variables using the neural data: **Model 3** was fit in a similar manner to **Model 2**, but additionally included principal components of the steady-state PLIs as regressors, as represented by the vector $\boldsymbol{x_j}$, on the subject-level model parameters.

We generated in-sample posterior predictive distributions using condition-level parameter posterior distributions (as opposed to in-sample posterior prediction from subject-level parameters), PLI coefficient posterior distributions, and PLI variables from each subject to find principal components that best predicted correct reaction time distributions. A plot of in-sample unexplained median correct-RT between-subject variance as a decreasing function of number of principal component (PC) regressors included in the model is provided in **Figure A.3** of the supplemental materials. Based on this analysis, principal components 2, 4, and 7 of both the noise and signal sets were tested further to find the model that best predicted out-of-sample reaction time of correct responses.

Model 3 was the model that best predicted out-of-sample correct-RT distributions by using noise component 2 and signal component 7 as exogenous PLI regressors on the diffusion model parameters. It should be noted that the amount of variance of the original PLI data explained by each PC is not reflective of each PC's out-of-sample predictive power, just as the amount of variance of the original data explained by each PC is not reflective of its contribution to the model (Jolliffe, 1982). A table of percent between-subject variance of median correct-RT explained (R^2_{pred}) by out-of-sample prediction is provided in **Table 2.1**. Tables of percent between-subject variance of mean, 25th percentile, and 75th percentile correct-RT explained by out-of-sample prediction are provided in the appendix. A new paricipant's correct-RT distribution in each condition can be more accurately predicted using the participant's EEG

in **Model 3**'s framework than by using **Model 1**'s or **Model 2**'s framework. 31.9% of the between-subject variance of the easiest condition's median correct-RT is explained by out-of-sample prediction.

To aid in interpretation, the posterior distributions of the regression coefficients for each PC were projected into the PLI coefficient space by multiplying the matrix of PC coefficient posterior samples **G** by the inverse-weight matrix **V** from the PCA algorithm which projects the principal components into the PLI data space. The result **GV** are samples from the posterior distributions of the regression coefficients for each PLI variable. This transformation was performed once for each of the noise and signal variable sets.

The posterior distributions of the signal PLI coefficients are provided in **Figure 2.5** with means, medians and 95% and 99% credible intervals. From the principal component coefficient and PLI coefficient posteriors, it was clear that there is a complex signal response at multiple frequencies and cortical locations on the diffusion coefficient and non-decision time. Participants with larger signal occipital 15 and 45 Hz PLIs are expected to have smaller variances in the evidence accumulation process (diffusion coefficients) than those participants with smaller occipital signal PLIs. However, the opposite effect is found in the frontal electrodes with large 15 and 45 Hz PLIs being associated with larger evidence accumulation variances. Larger signal responses at 30 and 45 Hz in parietal electrodes is also associated with larger diffusion coefficients. The effect of signal response on non-decision time is also complex but closely related to the effect of signal response on the diffusion coefficient. No evidence of an association between participants' differences in signal response to differences in evidence accumulation rates (drift rates) was found.

The posterior distributions of the noise PLI coefficients from the response interval are provided in **Figure 2.6**. The posterior distributions of the noise PLI coefficients from the noise interval are provided in **Figure A.4** of the appendix. In all noise harmonic frequencies during the noise interval and most harmonic frequencies (16, 24, 32, and 48 Hz) during the

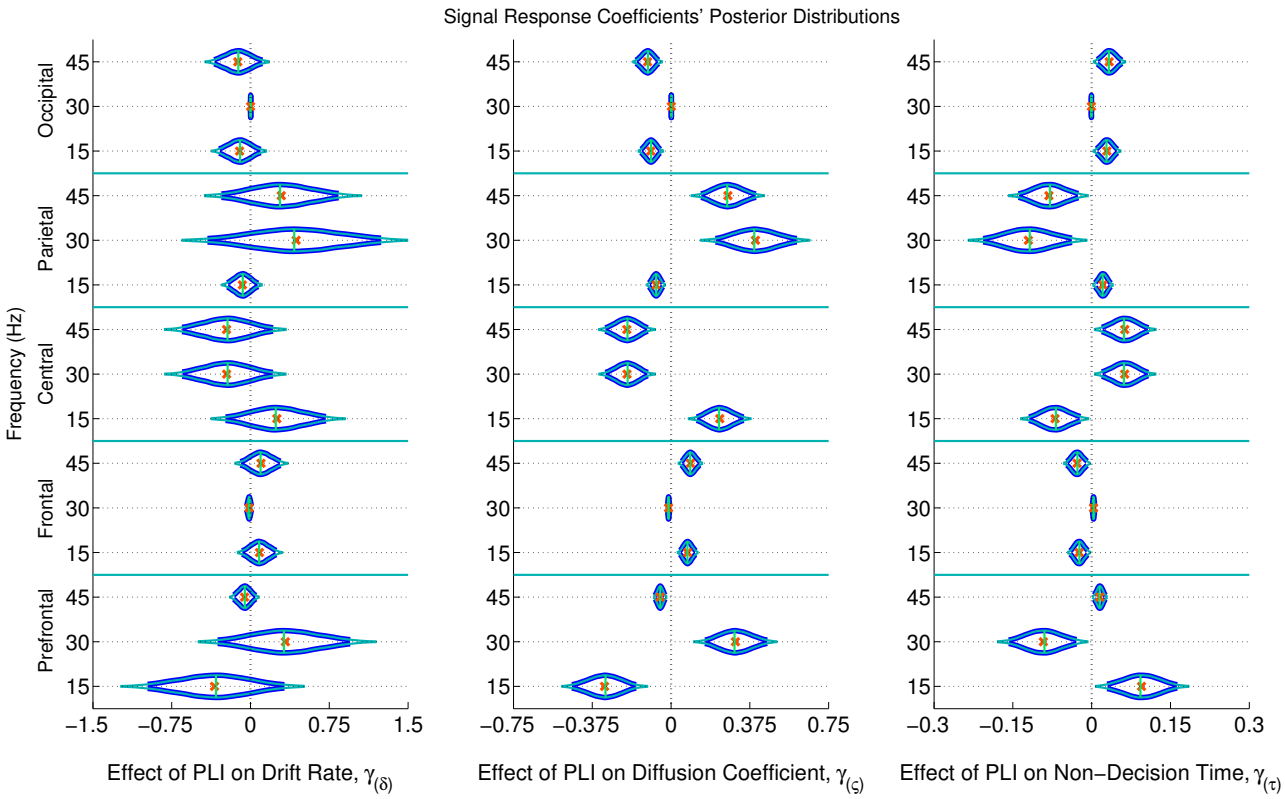

Figure 2.5: The marginal posterior distributions of the signal PLI coefficients (i.e., the effects of signal enhancement) as measured by a steady-state phase-locking index (PLI), on the evidence accumulation rate (drift rate; in evidence units per second), variance in the evidence accumulation process (the diffusion coefficient; in evidence units per second), and non-decision time during the response interval (in seconds). Dark blue posterior density lines indicate 95% credible intervals while smaller teal lines indicate 99% credible intervals. Small horizontal green lines embedded in density curves indicate the median of the posterior distributions while the orange crosses indicate posterior means. There is an effect of signal response on the diffusion coefficient and non-decision time that is complex across frequencies and scalp location. A participant whose PLI responses at all locations and frequencies are 0.2 units greater than another participant's responses is expected to have 0.061 evidence units per second larger evidence accumulation variances (where $\alpha = 1$ evidence unit is required to make a decision) and have 18 ms faster non-decision times, leading to faster but less accurate responses. There was no evidence of an effect of attention to the signal on evidence accumulation rate (the drift rate).

response interval, those subjects who had smaller PLIs at all electrode locations had faster evidence accumulation rates (drift rates). This finding suggests that those subjects who better suppressed the stimulus noise accumulated correct evidence faster. Furthermore, a similar effect was found on non-decision time. Noise suppression in the harmonic frequencies was associated with smaller non-decision times across subjects. However smaller PLIs at 8 Hz were associated with slower evidence accumulation and faster non-decision times. Looking at these effects as a whole, those subjects with more suppressed responses to the noise at all frequencies had larger drift rates and smaller non-decision times leading to faster, more accurate responses. As a plausible but oversimplified example, a participant whose PLI responses at all frequencies and locations was suppressed 0.2 units more than another participant during both the noise and response intervals is expected to accumulate 0.418 evidence units per second faster than another participant and have a 70 ms faster non-decision time. There was little to no evidence of an effect of individual variation in brain responses to noise on within-trial evidence accumulation variability (the diffusion coefficient).

2.4 Discussion

We have shown that a Bayesian diffusion model framework with hierarchical participant-level parameters is useful in describing individual differences in the rate of evidence accumulation, variance in evidence accumulation process, and preprocessing and/or motor response time in a novel perceptual decision making paradigm. Assuming the model describes the relationship between cognition and behavior sufficiently well, we are able to infer cognitive differences among participants. Furthermore, we have shown that differences in participants' attention as measured by steady-state visual evoked potentials relate to some of these differences in participants' cognition.

Individual differences in the rates of evidence accumulation (drift rates) were partially ex-

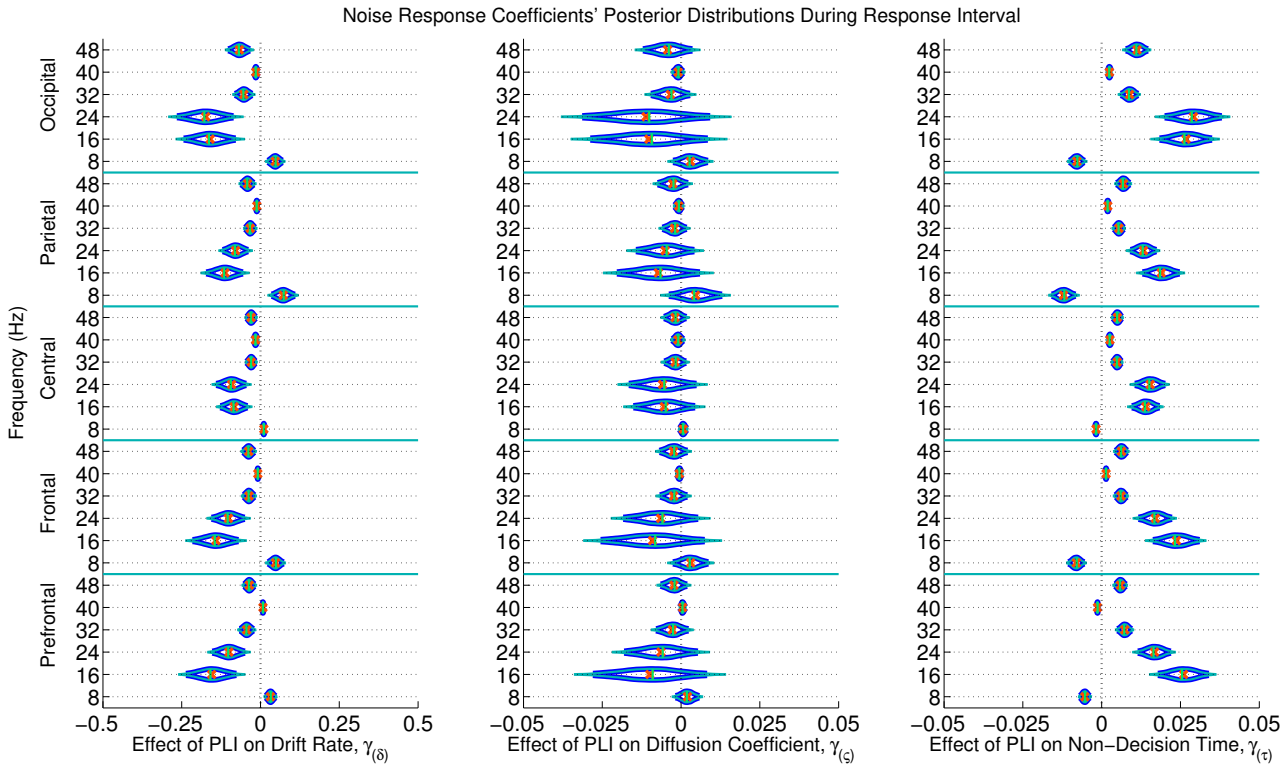

Figure 2.6: The marginal posterior distributions of the noise PLI coefficients (i.e., the effects of noise suppression) as measured by a steady-state phase-locking index (PLI), on the evidence accumulation rate (drift rate; in evidence units per second), variance in the evidence accumulation process (the diffusion coefficient; in evidence units per second), and non-decision time (in seconds) during the response interval. Dark blue lines indicate 95% credible intervals, smaller teal lines indicate 99% credible intervals, horizontal green lines indicate posterior medians, and the orange exes indicate posterior means. At noise harmonic frequencies (16, 24, 32, and 48 Hz) during the response interval, those subjects who suppressed noise had faster evidence accumulation rates; this effect was found at all electrode groups. However, noise enhancement at 8 Hz was associated with slower evidence accumulation. Furthermore, those subjects who better suppressed noise at the same harmonic frequencies had faster non-decision times. For example, a participant whose PLI responses were suppressed 0.2 units more than another participant's responses at all locations and frequencies during the response interval is expected to accumulate 0.288 evidence units per second faster (where $\alpha = 1$ evidence unit is required to make a decision) and have 48 ms faster non-decision times, leading to faster and more correct responses. There was no evidence of an effect of attention to the visual noise on variance in evidence accumulation (the diffusion coefficient).

plained by individual differences in noise suppression as measured by SSVEPs. Participants who better suppressed noise at high frequencies during the both the preparatory period (noise interval) and the decision period (response interval) were able to accumulate correct evidence faster, which led to more accurate, faster response times. Furthermore, those individuals who better suppressed noise in the same frequency bands and locations had faster non-decision times (preprocessing and/or motor response speed). This effect on non-decision time is hypothesized to be reflective of faster preprocessing time in subjects who better suppressed noise since we do not expect noise suppression to effect motor response speed. Both findings suggest a role of noise suppression in beta and gamma EEG frequency bands on the speed of evidence accumulation and preprocessing prior to evidence accumulation in perceptual decision making tasks.

Enhancement of signal was found to describe individual variation in "randomness" of evidence accumulation within trials (as measured by the diffusion coefficient). Participants who did not properly enhance signal in occipital, central, and pre-frontal electrodes had the most variable evidence accumulation processes. There is also evidence that a participant's enhancement of signal may have affected their preprocessing time in a complex way across frequencies and cortical locations. This suggests that signal enhancement in beta and gamma EEG frequency bands effect within-trial evidence accumulation variance and preprocessing in perceptual decision making.

In summary, from the results of the modeling procedure it was found that some individual variation in evidence accumulation speed (drift rate) is explained by noise suppression, some individual variation in evidence accumulation variance (diffusion coefficient) is explained by signal enhancement, and some individual variation in non-decision time (presumably preprocessing time) is explained by both noise suppression and signal enhancement.

The usefulness of the model with SSVEP attention measures as regressors is not only in its descriptive ability, but also in its predictive ability. New subject correct-RT behavior

was not accurately described by the model without individual differences nor the model with individual differences. But by explicitly including individual differences with neural covariates in hierarchical models, the correct reaction time distributions of new subjects with known neural measures are more accurately predicted. We expect the addition of the phase-locking index of SSVEPs to be predictive of behavior in any perceptual decision making paradigm, especially if used in a hierarchical Bayesian framework. Theoretically the hierarchical EEG-diffusion model will also be able to predict the PLI measures of a missing participant given a participant's behavioral data. We will explore the practicality of such predictions in future studies. Possible applications of behavioral and neural data prediction include: (a) the ability to interpolate data from incomplete behavioral data sets (b) the ability to interpolate data from incomplete neural data sets (c) more powerful statistical inference through simultaneous accounting for changes in behavior and neural data.

In the future for both hypothesis testing and response-RT prediction, latent variables linearly or non-linearly related to the EEG covariates can be included with the cognitive model in a hierarchical Bayesian framework (see Vandekerckhove, 2014, for details). The benefits of such an analysis would be: to choose neural covariates maximally descriptive or predictive of the data, choose electrodes and frequencies maximally descriptive or predictive of the data, reduce the number of covariates, and reduce the multicollinearity of the covariates by assuming there exist underlying variables related to multiple EEG covariates. In the present study, the problems of multicollinearity and variable overabundance were overcome with two principal component analyses (PCAs). PCAs do not extract mixtures of the data which are most descriptive or predictive of the model parameters but instead extract mixtures of the data which are uncorrelated. A shortcoming of this study is that we did not pick frequencies and cortical locations that were maximally predictive of behavior as exogenous variables. Cortical locations naively based upon large non-focal groupings were chosen. Instead of performing a non-Bayesian PCA before submitting the neural data to the Bayesian algorithm, a linear mixture of neural data that best describes the cognitive model parameters could be

extracted from the Bayesian algorithm itself, analogous to a partial least squares regression in a non-Bayesian approach (see Krishnan et al., 2013, for an example). In order to use this latent variable technique, the model must be run on a training set using a subset of the EEG data and then run on a test set to measure out-of-sample model predictive ability. This would result in a data reduction of the EEG that best predicts behavior in the context of the model.

		In-sample Prediction			Out-of-sample Prediction		
Rotation	Noise	M1	M2	M3	M1	M2	M3
±30°	30%	−0.1%	94.3%	95.0%	−13.5%	−11.8%	31.9%
±35°	30%	−0.1%	95.6%	95.8%	−12.3%	−11.7%	27.6%
±40°	30%	−0.5%	92.2%	92.1%	−12.5%	−11.5%	19.9%
±30°	45%	−1.2%	86.3%	87.4%	−15.1%	−11.7%	29.4%
±35°	45%	−0.2%	92.3%	91.6%	−12.0%	−13.6%	22.8%
±40°	45%	0.2%	92.6%	91.9%	−11.9%	−15.0%	28.0%
±30°	60%	−0.7%	93.1%	92.8%	−12.9%	−13.0%	18.6%
±35°	60%	−2.3%	92.5%	92.6%	−14.7%	−13.5%	13.3%
±40°	60%	−0.6%	90.9%	91.2%	−13.8%	−18.0%	26.2%

Table 2.1: Percentage of between-subject variance in correct-RT medians explained by in-sample and out-of-sample prediction (R^2_{pred}) for each experimental condition. The in-sample predictive ability of the no-individual differences **Model 1** was unsurprisingly poor, while the in-sample predictive ability of individual differences models (with and without EEG regressors, **Model 2** and **Model 3** respectively) explained most of the variance of correct-RT subject medians. Out-of-sample prediction was performed by using an iterative leave-one-subject-out procedure, first by obtaining posterior distribution estimates for each parameter by modeling all but one participant's behavior and EEG data and then estimating the left-out participant's correct-RT distribution using the resulting model fit and the left-out participant's EEG. Models without EEG regressors (i.e. **Model 1** and **Model 2**) are poor choices for new participant behavior prediction. The model with a noise principal component and a signal principal component of the phase-locked EEG as covariates of diffusion model parameters (**Model 3**) more accurately predicts new participants' correct-RT behavior. Negative values indicate overdispersion of the model prediction (due to posterior uncertainty) relative to the real data.

Chapter 3

Within-subject changes in attention influence decision making

3.1 Introduction

The goal of this study was to evaluate whether attention could predict different components of the decision making process on each trial of a visual discrimination experiment. We make use of high-density electroencephalographic (EEG) recordings from the human scalp to find single-trial evoked potentials (EPs) to the onset of visual signal and to the onset of a distractor (mask) to measure the deployment of attention to task-relevant features. We found that on each trial, modulations of the evoked potentials by attention were predictive of specific components of a drift-diffusion model of the decision making process.

3.1.1 Visual attention and decision making

Attention is beneficial for decision making because relevant features of the environment can be preferentially processed to enhance the quality of evidence. During visual tasks individuals may deploy different attention strategies such as: enhancing the signal, suppressing external noise (distractors), or suppressing internal noise (Lu and Dosher, 1998; Dosher and Lu, 2000b). These strategies are thought to change based on the signal to noise ratio of the stimulus, such that individuals will enhance sensory gain to both signal and noise during periods of low noise and sharpen attention to only signal during periods of high noise (Lu and Dosher, 1998), although specific strategies have been shown to differ across subjects (Bridwell et al., 2013; Krishnan et al., 2013; Nunez et al., 2015). Multiple groups have proposed models of visual attention and decision making that yield diverse reaction time and choice distributions dependent upon attentional load (Spieler et al., 2000; Smith and Ratcliff, 2009). Attention can be deployed to the features and/or location of a stimulus, and attention can benefit decision making when the subject is cued to the location or features of the stimulus (Eriksen and Hoffman, 1972; Shaw and Shaw, 1977; Davis and Graham, 1981).

Event-related potentials (ERPs) are trial-averaged EEG responses to external stimuli. Visual ERPs (also labeled Visual Evoked Potentials; VEPs) have been shown to track visual attention to the onset of stimuli (Harter and Aine, 1984; Luck et al., 2000). That is, amplitudes of the peaks of the ERP waveform (i.e., ERP "components") within certain millisecond-scale time windows are shown to be larger when subjects encounter task-relevant stimuli in the expected location in visual space. Two components of particular interest are the *N1* (or *N200*) and *P2* (or *P200*) components. The terms *N1* and *P2* refer the order of negative and positive peaks in the time series respectively, and the more general alternative names *N200* and *P200* refer to their approximate latencies in milliseconds. Changes in N200 latencies have been shown to correlate with attentional load (Callaway and Halliday, 1982), and N200 measures have even been used in Brain-Computer Interfaces (BCI) that make use

of subjects' attention to specific changing stimuli, such as letters in a BCI speller (Hong et al., 2009). Findings in these trial-averaged EEG (ERP) studies suggest that information is also available in single-trials of EEG that can be used to evaluate the relationship between attention and decision making. In this paper we will use the alternative names *P200* and *N200* because 1) the exact time windows of components vary across studies, 2) components in this study were both localized to around 200 milliseconds, and 3) components in this study were found on *single-trials* as opposed to in the trial-average.

3.1.2 Single-trial EEG measures of attention

EEG correlates of attention and decision making have been found using classification methods. One group has shown that the amplitude of certain EEG components in the time domain track type and duration of two-alternative forced choice responses and then showed that these components' amplitudes tracked evidence accumulation rate (Philiastides et al., 2006; Ratcliff et al., 2009). However the EEG components in these studies were found by finding the maximum predictors of the behavioral data and thus had no a priori interpretation. Another group has found that that single-trial amplitude in a few frequency bands, especially the 4-9 Hz theta band, predicts evidence accumulation rates (van Vugt et al., 2012). However these oscillations were found using canonical correlation analysis (CCA; Calhoun et al., 2001), a data driven algorithm that found any EEG channel mixtures that contained correlations with the drift diffusion model parameters. While the results were confirmed using cross validation, the set of EEG identified by this method also did not have an a priori explanation. These studies point us in directions of exploration and perform well at prediction, but we have little information as to whether the EEG information reflected attention, the decision process itself, or some other correlate of evidence accumulation. In this study, we introduce a simple procedure that is informed by ERPs known to be related to attention, and we make use of single-trial ERP estimates to model behavior on single trials.

3.1.3 Hypothesized attention effects

An integrated model of visual attention, visual short term memory, and perceptual decision making by Smith and Ratcliff (2009) predicts that attention operates on the encoding of the stimulus, and that enhanced encoding increases drift rate during the decision making process. Furthermore, the model predicts that visual encoding time (i.e. visual preprocessing) will be reduced by attention which is reflected in the non-decision time parameter. However, this model of visual attention only considers the detection of a stimulus in an otherwise blank field—that is, a field with no visual noise. Thus, it does not have predictions for the distinct processes of noise suppression and signal enhancement, as in the Perceptual Template Model (Lu and Dosher, 1998). Signal enhancement during the evidence accumulation process is predicted to reduce the diffusion coefficient ς because the mechanism by which signal enhancement takes place, according to the Perceptual Template Model, is additive *internal noise* reduction[1] (Dosher and Lu, 2000a); this mechanism is predicted to be most effective in low noise conditions since decreasing internal noise will lead to better processing of *both* the visual signal and external visual noise. External noise suppression, on the other hand, is expected to reflect the encoding of the stimulus by manipulation of a perceptual template, increasing the average rate of evidence accumulation δ by improving the overall quality of evidence on a trial. The Perceptual Template Model predicts this mechanism is most effective in high noise conditions.

In a previous study we showed that *individual differences* in noise suppression predicts individual differences in evidence accumulation rates and non-decision times (Nunez et al., 2015). We also showed that differences across individuals in signal enhancement predict individual differences in non-decision times and evidence accumulation variance (i.e. the diffusion coefficient), which we assume tracks internal noise in the subject. The findings of

[1] "In stimulus enhancement, attention increases the gain on the stimulus, which is formally equivalent to reducing internal additive noise. This can improve performance only in low external noise stimuli, since external noise is the limiting factor in high external noise stimuli." (Dosher and Lu, 2000a, p. 1272)

signal enhancement effects on evidence accumulation variance and noise suppression effects on evidence accumulation rate seem to correspond closely to predictions made by the Perceptual Template Model. However the Perceptual Template Model does not make explicit predictions about attention effects on non-decision times. The previous study did not explore how trial-to-trial variation in attention affected trial-to-trial cognitive differences. Individual differences in attention could be found that are not detected to be changing within a subject, and/or trial-to-trial variability in attention could occur that does not change across individuals. In this study, we show that within-subject, trial-to-trial variability in attention to *both* noise and signal predict variability in drift rate and non-decision times, corresponding closely to predictions made by the model of Smith and Ratcliff (2009) that predicts speeded encoding time and increased evidence accumulation rate due to enhanced attention. The two studies together suggest that within-trial evidence accumulation variances ς varied across individuals, but we did not find evidence that this measure varied within individuals due to changes in trial-to-trial attention.

3.2 Methods

3.2.1 Experimental stimulus: Bar field orientation task

Reported in a previous study, behavioral and EEG data were collected from a simple two-alternative forced choice task to test *individual differences* in attention during visual decision making (Nunez et al., 2015). Here, we reanalyze these data to explore *per-trial* attention effects on the decision making process. Subjects were instructed on each trial to differentiate the mean rotation of a field of small bars that were either oriented at 45 deg or 135 deg from horizontal on average. Two representative frames of the display and the time course of a trial are provided in **Figure 3.1**. The circular field of small bars was embedded in a square

field of visual noise that was changing at 8 Hz. The bar field was flickering at 15 Hz. These frequencies were chosen to evoke steady-state visual evoked potentials (SSVEPs), stimulus frequency-tagged EEG responses that were useful in the previous study but will not be used in this study. Stimuli were built and displayed using the MATLAB Psychophysics toolbox (Psychtoolbox-2; www.psychtoolbox.org).

Subjects viewed each trial of the experimental stimulus on a monitor in a dark room. Subjects sat 57 cm away from the monitor. The entire circular field of small oriented bars was 9.5 cm in diameter, corresponding to a visual angle of 9.5°. Within each trial subjects first observed a black cross for 750 ms in the center of the screen, on which they were instructed to maintain fixation throughout the trial. Subjects then observed visual noise for 750 ms. This time period of the stimulus will henceforth be referred to as *cue interval*, with the onset EEG response at the beginning of this interval being the response to the *noise* (or "distractor") stimulus. Subjects then observed the circular field of small oriented bars overlaid on the square field of visual noise for 1000 to 2000 ms and responded during this interval. Subjects were instructed to respond as accurately as possible while providing a response during every trial. Because evidence required to make a decision only appeared during this time frame, the decision process was assumed to take place during this interval. This interval is referred to as the *response interval*, and attentional onset EEG measures during this time period are referred to as responses to the *signal* stimulus. After the response interval the fixation cross was again shown for 250 ms to alert the subjects that the trial was over and to collect any delayed responses.

Three levels of variance of bar rotation and three levels of noise luminance were used to modulate the task difficulty. However only the noise luminance manipulation is relevant for the analysis presented here. Average luminance of the noise was 50% and the luminance of the bars was 15%. In the low noise condition, that luminance was drawn randomly at each pixel from a uniform distribution of 35% to 65% luminance. In the medium and high noise

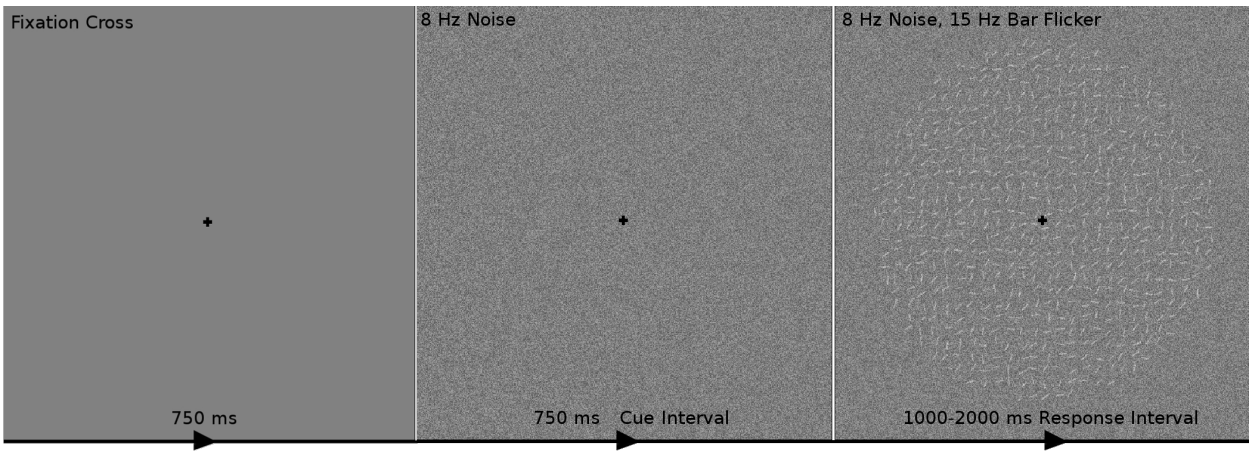

Figure 3.1: The time course of one trial of the experimental stimulus. One trial consisted of the following: 1) 750 ms of fixation on a black cross on a gray screen, then 2) visual contrast noise changing at 8 Hz for 750 ms while maintaining fixation (dubbed the *cue interval*) and 3) a circular field of small oriented bars flickering at 15 Hz overlaid on the changing visual noise for 1000 to 2000 ms while maintaining fixation (dubbed the *response interval*). The subjects' task was to indicate during the response interval whether the bars were on average oriented towards the "top-right" (45° from horizontal; as in this example) or the "top-left" (135° from horizontal). It was assumed that subjects' decision making process occurred only during the response interval but could be influenced by both onset attention to the visual noise during the cue interval and onset attention to visual signal during the response interval.

conditions, noise luminance was drawn randomly at each pixel from a uniform distribution of 27.5% to 72.5% and 20% to 80% luminance respectively. Each subject experienced 180 trials from each noise condition, interleaved, for a total of 540 trials split evenly over 6 blocks. The total duration of the visual experiment for each subject was approximately one hour and 15 minutes including elective breaks between blocks. More details of the experiment can be found in our previous publication (Nunez et al., 2015).

Behavioral and EEG data were collected concurrently from 17 subjects. Subjects performed accurately during the task. The across-subject mean, standard deviation, and median of accuracy were $90.1\% \pm 5.8\%$, $\widetilde{y} = 91.6\%$, while the across-subject mean, standard deviation, and median of average reaction time were 678 ± 106 , $\widetilde{t} = 670$ ms. Individual differences in behavior existed across subjects with the most accurate subject answering 98.3% of trials correctly and the least accurate subject answering 78.5% of trials correctly. Two different subjects were the fastest and slowest with mean RTs of 502 ms and 866 ms respectively.

3.2.2 Single-trial EEG predictors

Electroencephalograms (EEG) were recorded using Electrical Geodesics, Inc.'s high density 128-channel Geodesic Sensor Net and Advanced Neuro Technology's amplifier. Electrical activity from the scalp was recorded at a sampling rate of 1024 samples per second with an online average reference using Advanced Neuro Technology software. The EEG data was then imported into MATLAB for offline analysis. Linear trends were removed from the EEG data, and the data were band pass filtered to a 1 to 50 Hz window using a high pass Butterworth filter (1 Hz pass band with a 1 dB ripple and a 0.25 Hz stop band with 10 dB attenuation) and a low pass Butterworth filter (50 Hz pass band with 1 dB ripple and a 60 Hz stop band with 10 dB attenuation).

EEG *artifact* is broadly defined as data collected within EEG recordings that does not

originate from the brain. Electrical artifact can be biological (e.g. from the muscles-EMG or from the arteries-EKG) or non-biological (e.g. temporary electrode dislocations, DC shifts, or 60 Hz line noise). Contribution of muscle and electrical artifact was reduced in recordings by using an extended Infomax Independent Component Analysis algorithm (ICA; Makeig et al., 1996; Lee et al., 1999). ICA algorithms are used to find linear mixtures of EEG data that relate to specific artifact. Components that are indicative of artifact typically have high spatial frequency scalp topographies, high temporal frequency or a $^1/_f$ frequency falloff, and are present either in only a few trials or intermittently throughout the recording. These properties are not shared by electrical signals from the brain as recorded on the scalp (Nunez and Srinivasan, 2006). Using these metrics, components manually deemed to reflect artifact were projected into EEG space and subtracted from the raw data. Components deemed to be a mixture of artifact and brain activity were kept. More information about using ICA to reduce the contribution of artifact can be found in Jung et al. (2000).

Event-related potential (ERP) components have been shown to index attention (Callaway and Halliday, 1982; Harter and Aine, 1984; Luck et al., 2000), in particular the P200 and N200 latencies and amplitudes, and these values were used as independent measures of attention in the following analyses. Event-related potentials (ERPs) are EEG responses that are time-locked to a stimulus onset and are typically estimated by aligning and averaging EEG responses across trials. They usually cannot be directly measured on each trial from single electrodes. Raw EEG signals could be used as a single-trial measures but typically have very low signal-to-noise ratios (SNR) for task-specific brain responses. Since the goal of this analysis was to explore single-trial effects of attention on visual decision making, a single-trial estimate of the ERP was developed.

Because the signal-to-noise ratio (SNR) in ongoing EEG increases when adjacent electrodes of relevant activity are summed (Parra et al., 2005), we anticipated that the SNR of the response to the visual stimulus would be boosted on individual trials by summing over the

mixture of channels that best described the average visual response. Traditional ERPs at each channel (represented by a matrix of size $T \times C$ where T is the length of a trial in milliseconds and C is the number of EEG channels) were calculated separately for each subject. One ERP was calculated for the response to the visual signal and another was found for the response to the visual noise by averaging a random set of two-thirds of the trials across all conditions for each subject in each window. This random set of trials was the same set used as the training set for cross validation, to be discussed later. The test sets of trials were not used to calculate the traditional ERPs.

Singular value decomposition (SVD; analogous to principal component analysis) of the *trial averages* were then used to find linear mixtures of channels that explained the largest amount of the variance in the ERP data (i.e. the first right-singular vectors $\mathbf{v}$, explaining a percentage of variance from 39.4% to 91.9% and 45.0% to 93.2% across subjects in the cue and response intervals respectively). The first right-singular vectors were then used as weights to mix the raw EEG data into a brain response biased toward the maximum response to the visual stimuli, yielding one time course of the EEG per trial for both the cue and response intervals. A visual representation of the simple procedure for a single trial is provided in **Figure 3.2**. The raw data matrix $\mathbf{E}$ of dimension $N \times C$ was multiplied by the first right singular vector $\mathbf{v}$ (a $C \times 1$ vector of channel weights) to obtain a $N \times 1$ vector $\mathbf{Ev} = \mathbf{e}$, which could then be split up into epochs of length $T \times 1$ representing the response to the stimulus on each trial. Note that the voltage amplitudes of the ERP measures calculated based on this method will differ from traditional single electrode ERP amplitudes since the single-trial estimates are a weighted sum of potentials over all electrodes.

Not only did this method boost the SNR of the EEG measures, but this method also reduced EEG measures of size $T \times C$ on each trial to one latent variable that varies in time of size $T \times 1$. Thus the correlation of the EEG as inputs to the model was drastically reduced and the interpretability of model parameters was increased compared to analyses with highly

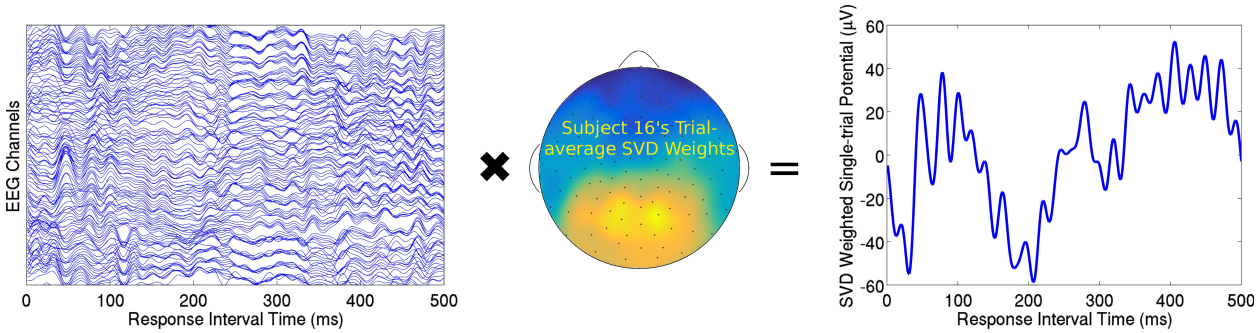

Figure 3.2: A visual representation of the singular value decomposition (SVD) method for finding single-trial estimates of evoked responses in EEG. The EEG presented here is time-locked to the signal onset during the response interval, such that the single-trial ERP encoded the response to the signal onset. A single trial of EEG from Subject 16 (Left) can be thought of as a time by channel ($T \times C$) matrix. The first SVD component explained the most variability (79.9%) in Subject 16's ERP response to the signal across all trials in the training set. SVD weights $\mathbf{v}$ ($C \times 1$) are obtained from the ERP response (i.e. trial-averaged EEG; $T \times C$) and can be plotted on a cartoon representation of the human scalp with intermediate interpolated values (Middle). This specific trial's ERP (Right) was obtained by multiplying the time series data from each channel on this trial by the associated weight in vector $\mathbf{v}$ and then summing across all weighted channels.

correlated model inputs. The weight vector $\mathbf{v}$ for each subject in both the cue and response intervals also yields a scalp map when the values of the weights are interpolated between electrodes. Channel weights calculated using SVD on subject's ERPs to the noise onset (during the cue interval) are shown in topographic plots for each subject in **Figure 3.7**. Channel weights calculated using SVD on subject's ERPs to the signal onset (during the response interval) are shown in topographic plots for each subject in **Figure 3.6**. While raw EEG on single trials from single electrodes may have large enough SNRs to be informative for our analysis, we would not obtain an idea of the locus of activation or the pattern of activation over the scalp.

We focused our analysis on the windows 150 to 275 ms post stimulus-onset in the cue and response intervals. These windows were found to contain P200 and N200 ERP components. On each trial, we measured the peak positive and negative amplitudes, and the latency at which these peaks were observed. We used these 8 single-trial measures to predict single-trial

diffusion model parameters. However in this paper we will focus only on the results of models with 4 single-trial measures: the amplitude and latency of the peak positive deflection (P200) during the *cue interval* and the amplitude and latency of the peak negative deflection (N200) during the *response interval*, because very weak evidence, if any, was found for the effects of the other attention measures on diffusion model parameters in models with all 8 single-trial measures. It should be noted that single-trial measures of EEG spectral responses at SSVEP frequencies (see Nunez et al., 2015) were briefly explored but future methods must be developed to increase signal-to-noise ratios of SSVEP measures on single-trials.

3.2.3 Hierarchical Bayesian models

Hierarchical models of visual decision making were assumed and placed into a Bayesian framework. Bayesian methods yield a number of benefits compared to other inferential techniques such as traditional maximum likelihood methods. Rather than point estimates of parameters, Bayesian methods provide entire distributions of the unknown parameters. Bayesian methods also allow us to perform the model fitting procedure in a single step, maintaining all uncertainty about each parameter through each hierarchical level of the model.

One downside of Bayesian methods is that creating sampling algorithms to find posterior distributions of Bayesian hierarchical models can be time consuming and cumbersome. However Just Another Gibbs Sampler (JAGS; Plummer, 2003) is a program that uses multiple sampling techniques to find estimates of hierarchical models, only requiring the form of the model, data, and initial values as input from the user. In order to find posterior distributions, we have used JAGS with an extension that adds a diffusion model distribution (without intrinsic trial-to-trial variability) as one of the distributions to be sampled from (Wabersich and Vandekerckhove, 2014). Similar software packages to fit hierarchical diffusion models

have been developed independently in other programming languages such as Python (Wiecki et al., 2013).

In order to evaluate the benefit to prediction of adding EEG measures to hierarchical diffusion models, three different models were compared. Model 3 assumed that evidence accumulation rates, evidence accumulation variances, and non-decision times were each equal to a linear combination of EEG measures on each trial. Because we found no effect of the observed single-trial EEG measures on single-trial evidence accumulation variances, we also fit Model 2, where single-trial evidence accumulation rates and non-decision times were influence by EEG, but single-trial evidence variances were not. Model 1 did not assume any EEG contribution to any parameters. This model assumed that parameters not varying with EEG would change based on subject and condition, drawn from a condition level distribution. Graphical representations of the hierarchical Bayesian models are provided in **Figure 3.3** following the convention of Lee and Wagenmakers (2014).

For Model 1 (**Figure 3.3a**), prior distributions were kept mostly uninformative (i.e. parameters of interest had prior distributions with large variances) so that the analyses would be data-driven. The prior distributions of parameters for each subject j and condition k free from EEG influence had the following prior and hyperprior structure

$$\begin{aligned}
\left(\delta_{jk} \mid \mu_{(\delta)k}, \sigma_{(\delta)}\right) &\sim \mathcal{N}(\mu_{(\delta)k}, \sigma^2_{(\delta)}), \quad \mu_{(\delta)k} \sim \mathcal{N}(1.5, 4^2) \in (-9, 9), \quad \sigma_{(\delta)} \sim \Gamma(5, 0.20) \\
\left(\tau_{jk} \mid \mu_{(\tau)k}, \sigma_{(\tau)}\right) &\sim \mathcal{N}(\mu_{(\tau)k}, \sigma^2_{(\tau)}), \quad \mu_{(\tau)k} \sim \mathcal{N}(0.3, 1^2) \in (0, 3) \quad , \quad \sigma_{(\tau)} \sim \Gamma(5, 0.05) \\
\left(\varsigma_{jk} \mid \mu_{(\varsigma)k}, \sigma_{(\varsigma)}\right) &\sim \mathcal{N}(\mu_{(\varsigma)k}, \sigma^2_{(\varsigma)}), \quad \mu_{(\varsigma)k} \sim \mathcal{N}(0.6, 2^2) \in (0, 4) \quad , \quad \sigma_{(\varsigma)} \sim \Gamma(5, 0.05)
\end{aligned}$$

Where the normal distributions $\mathcal{N}$ are parameterized with mean and variance respectively and the gamma distributions Γ are parameterized with shape and scale parameters respectively.

In Models 2 (**Figure 3.3b**) and 3 (**Figure 3.3c**), to ensure noninterference by the prior distributions, uninformative priors were given for both the effects $\boldsymbol{\gamma}_{jk}$ of EEG on the param-

eters of interest and the linear intercepts η_{jk}. Note that the effect of EEG $\boldsymbol{\gamma}_{jk}$ is a vector with one element per EEG regressor and each effect of EEG is assumed to be statistically independent from the others. If a drift-diffusion model parameter was assumed to be equal to a linear combination of EEG inputs then the following two lines replaced the priors of the respective parameter above.

$$\begin{aligned}
\left(\eta_{jk} \mid \mu_{(\eta)k}, \sigma_{(\eta)}\right) &\sim \mathcal{N}(\mu_{(\eta)k}, \sigma^2_{(\eta)}) \quad , \quad \mu_{(\eta)k} \sim \mathcal{N}(0, 100^2) \quad , \quad \sigma_{(\eta)} \sim \Gamma(5,5) \\
\left(\boldsymbol{\gamma}_{jk} \mid \boldsymbol{\mu}_{(\gamma)k}, \sigma_{(\gamma)}\right) &\sim \mathcal{MVN}(\boldsymbol{\mu}_{(\gamma)k}, \sigma^2_{(\gamma)}I), \quad \boldsymbol{\mu}_{(\gamma)k} \sim \mathcal{MVN}(\mathbf{0}, 100^2 I), \quad \sigma_{(\gamma)} \sim \Gamma(5,5)
\end{aligned}$$

In Models 2 and 3, the parameter on each trial was assumed to be equal to a simple linear combination of the vector of single-trial EEG inputs $\mathbf{x}_{ijk}$ on that trial i with η_{jk} and $\boldsymbol{\gamma}_{jk}$ as the intercept and slopes respectively:

$$\begin{aligned}
\delta_{ijk} &= \eta_{(\delta)jk} + \boldsymbol{\gamma}^{\intercal}_{(\delta)jk}\mathbf{x}_{ijk} \\
\tau_{ijk} &= \eta_{(\tau)jk} + \boldsymbol{\gamma}^{\intercal}_{(\tau)jk}\mathbf{x}_{ijk} \\
\varsigma_{ijk} &= \eta_{(\varsigma)jk} + \boldsymbol{\gamma}^{\intercal}_{(\varsigma)jk}\mathbf{x}_{ijk}
\end{aligned}$$

Where the first two equations refer to the structure of Model 2 and all three equations refer to the structure of Model 3. Note that the $p * 1$ vector of effects $\boldsymbol{\gamma}_{jk}$ of EEG on each parameter could include the intercept term η_{jk} to create a $(p+1) * 1$ vector of effects $\boldsymbol{\gamma}^*_{jk}$ (and the EEG vector $\mathbf{x}_{ijk}$ would include a value of 1 to be multiplied by the intercept term). We use this notation in **Figure 3.3** for simplicity.

Because not all trials are believed to actually contain a decision-making process (i.e. the subject quickly presses a random button during a certain percentage of trials reflecting a "fast guess"), reaction times below a certain threshold were removed from analysis and cross-validation. Cutoff reaction times were found for each subject by using an exponential moving average of accuracy after sorting by reaction time (Vandekerckhove and Tuerlinckx, 2007). The rejected reaction times were all below 511 ms with a mean cutoff of 410 ms across

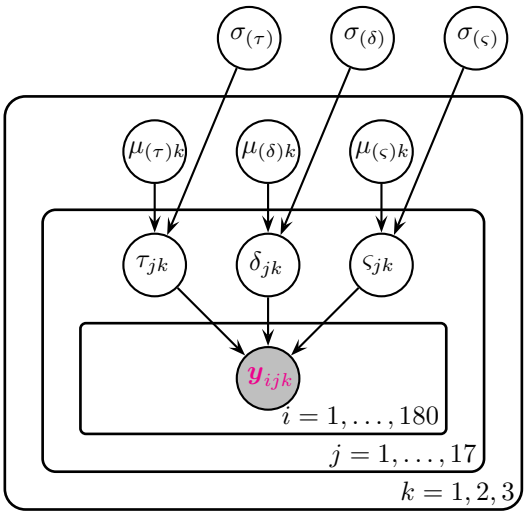

(a) Model 1

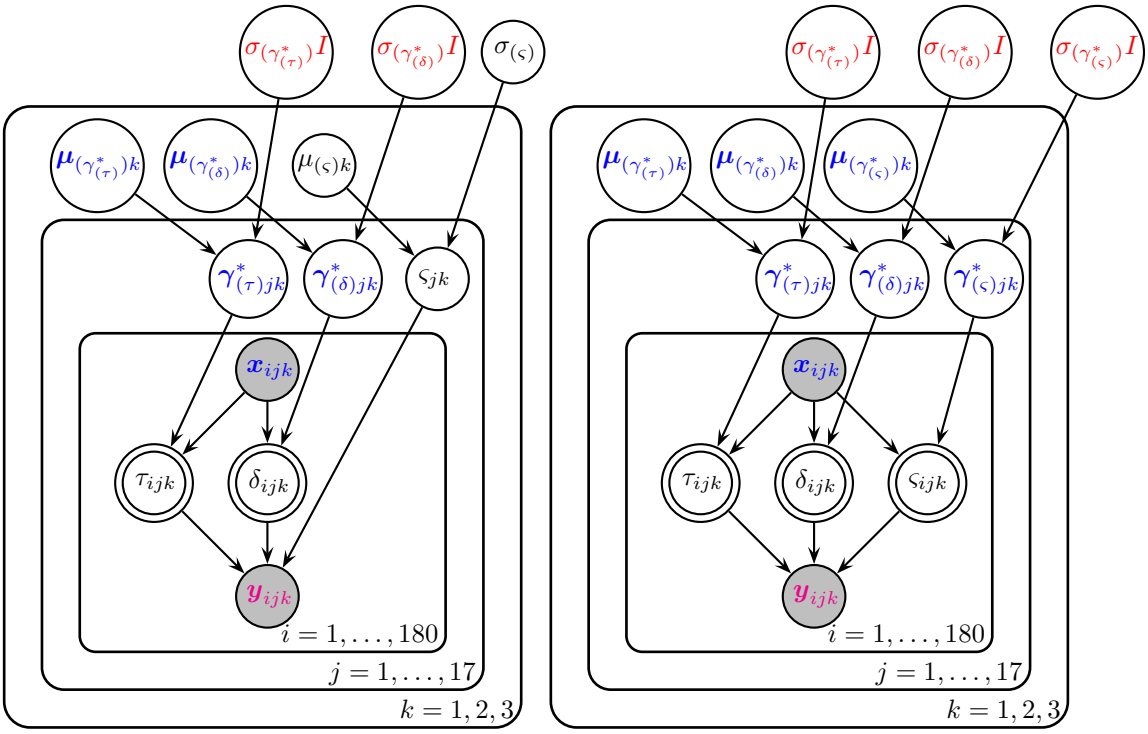

(b) Model 2 (c) Model 3

Figure 3.3: Graphical representations of the three hierarchical Bayesian models following the convention of Lee and Wagenmakers (2014). Each node represents a variable in the model with arrows indicating what variables are influenced by other variables. The magenta $2 * 1$ vector of reaction time and accuracy $\boldsymbol{y}_{ijk}$ and the blue $(p + 1) * 1$ vector of p EEG regressors (+1 intercept) $\boldsymbol{x}_{ijk}$ are observed variables, as indicated by the shaded nodes. Bolded blue variables indicate $(p + 1) * 1$ vectors, such as the subject j level effects $\boldsymbol{\gamma}^*_{jk}$ of each EEG regressor and the condition k level effects $\boldsymbol{\mu}_{(\gamma^*)k}$ of each EEG regressor. In Model 3 for each trial i, values of non-decision time τ_{ijk}, drift rate (evidence accumulation rate) δ_{ijk}, the diffusion coefficient (evidence accumulation variance) ς_{ijk} are deterministic linear combinations of single-trial EEG regressors $\boldsymbol{x}_{ijk}$ and the effects of those regressors $\boldsymbol{\gamma}^*_{jk}$ that vary by subject and condition.

subjects. This resulted in an average rejection rate of 1.4% of trials across subjects with a maximum of 6.3% of trials rejected for one subject and a minimum of 0.7% of trials rejected for 11 of the 17 subjects.

Each model was fit using JAGS with six Markov Chain Monte Carlo (MCMC) chains run in parallel (Tange, 2011) of 52,000 samples each with 2,000 burn-in samples and a thinning parameter of 10 resulting in 5,000 posterior samples in each chain. The posterior samples from each chain were combined to form one posterior sample of 30,000 samples for each parameter. All three models converged as judged by $\hat{R}$ being less than 1.02 for all parameters in each model. $\hat{R}$ is a statistic used to assess convergence of MCMC algorithms (Gelman and Rubin, 1992).

Posterior distributions were found for each free parameter in the three models. Credible intervals of the found posterior distributions were then calculated to summarize the findings of each model. EEG regressor effects were deemed to have weak evidence if the 95% credible interval between the 2.5^{th} and 97.5^{th} percentiles of the subject mean parameter $\mu_{(\gamma)j}$ was non-overlapping zero. Effects were deemed to have strong evidence if the 99% credible interval between the 0.5^{th} and 99.5^{th} percentiles was non-overlapping zero.

3.2.4 Cross-validation

All trials from all subjects were used during initial exploration of the data. However once it was decided that the signal onset response was a candidate predictor of drift rate, cross-validation was performed using a training and test set of trials. Out-of-sample performance for both *known* and *unknown* subjects were found by randomly assigning two-thirds of the trials from each subject in a random sample of *subjects* (i.e. 13 of 17 subjects were chosen at random) as the training set and one-third of the trials from the 13/17 subjects and all trials from the remaining 4/17 subjects as the test set. After fitting the model with the

training set, posterior predictive distributions of the accuracy-RT data were found for each subject. Posterior predictive distributions were calculated by drawing from the subject-level posteriors of the *known* subjects and by drawing from the condition-level posteriors of the *unknown* subjects. The posterior predictive distributions were then compared to the sample distribution of the test set.

In some recent papers, evaluation of models' prediction ability has been left to the readers with the aid of posterior predictive coverage plots (e.g. see figures in Supplementary Materials). Here we formally evaluate the similarity of the posterior predictive distributions to the test samples via a "proportion of variance explained" calculation. Specifically, we calculated R^2_{pred} of subjects' accuracy and correct reaction time 25th percentiles, medians, and 75th percentiles across subjects. R^2_{pred} is a measure of percentage variance in a statistic T (e.g. accuracy, correct-RT median, etc.) explained by in-sample or out-of-sample prediction. In this paper, R^2_{pred} is defined as the percentage of total *between-subject variance* of a statistic T explained by out-of-sample or in-sample prediction. It is a function of the mean squared error of prediction (MSEP) and the sample variance of the statistic T based on a sample size of $J = 13$ or $J = 4$ subjects for known and unknown subject calculations respectively. This measure also allows comparisons across studies with similar prediction goals. Mathematically, R^2_{pred} is defined as

$$R^2_{\text{pred}} = 1 - \frac{\sum_{j=1}^{J}(T_j - T_{(\text{pred})j})^2/(J-1)}{\sum_{j=1}^{J}(T_j - \bar{T})^2/(J-1)} = 1 - \frac{\text{MSEP}_T}{\widehat{\text{Var}[T]}} \tag{3.1}$$

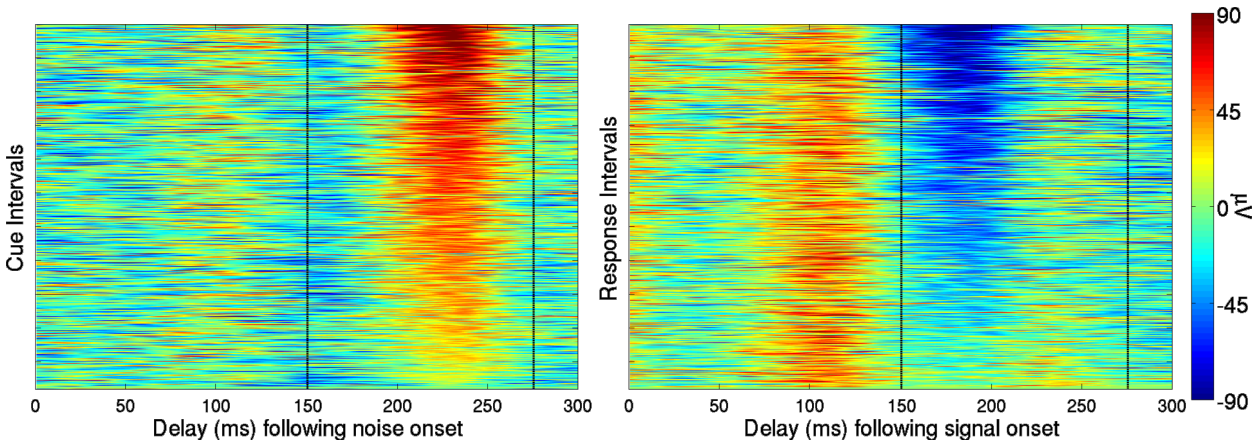

Figure 3.4: Single-trial evoked responses of an example subject, Subject 12, to the visual noise during the cue interval (Left) and single-trial evoked responses to the visual signal during the response interval (Right). Single-trial P200 and N200 magnitudes were found by finding peak amplitudes in 150 to 275 ms time windows (as indicated by the vertical dashed lines) of the SVD-biased EEG data in both the cue and response intervals. The first 300 ms of the intervals are sorted by single-trial P200 magnitudes in the cue interval and single-trial N200 magnitudes in the response interval. Latencies of the single-trial P200 and N200 components correspond to known latencies of P2 and N1 ERP components.

3.3 Results

The single-trial EEG measures "regressed" on diffusion model parameters were the peak positive and negative amplitudes and latencies (corresponding to P200 and N200 peaks respectively) in the 150 to 275 ms windows post noise-onset in the cue interval and post signal-onset in the response interval. However the magnitude and latency of the peak negative deflection (N200) in response to the noise stimulus and the magnitude and latency of the peak positive deflection (P200) in response to the signal stimulus were not informative (i.e. most condition-level effect posteriors of these measures overlapped zero significantly in models with all P200 and N200 measures included as regressors). For simplicity we only discuss results of models with P200 measures following the noise stimulus in the cue interval and N200 measures following the signal stimulus in the response interval. Example single trial amplitudes of these P200 and N200 peaks for Subject 12 are shown in **Figure 3.4**.

Since no effects of explored measures were found on within-trial evidence accumulation variance in Model 3 (i.e. posterior distributions of $\boldsymbol{\gamma}_{(\varsigma)jk}$ were centered near zero), the only EEG effects that will be discussed are those on evidence accumulation rate and non-decision time from a fit of Model 2. The posterior distributions of EEG effects on evidence accumulation rate $\boldsymbol{\gamma}_{(\delta)jk}$ and non-decision time $\boldsymbol{\gamma}_{(\tau)jk}$ did not differ significantly from Model 2 to Model 3. Moreover, Model 2 produced better out-of-sample prediction than Model 3 for new subjects. A graphical example of the effects found with Model 2 in two representative trials are given in **Figure 3.5**.

3.3.1 Intercept terms of evidence accumulation rate and non-decision time

The intercept term of each variable gives the value of each variable not explained by a linear relationship to N200 and P200 amplitudes and latencies. That is, the intercept gives the value of each parameter that remains constant from trial to trial, with the between-trial variability of the parameter being influenced by the changing trial-to-trial EEG measures. Model 2's posterior medians of the condition level evidence accumulation rate intercepts $\mu_{(\eta_\delta)j}$ and non-decision time intercepts $\mu_{(\eta_\tau)j}$ are reported. In low noise conditions, evidence accumulation rate intercepts were 1.46 evidence units per second (i.e. if there was no behavioral effect of EEG on each trial and no variance in the evidence accumulation process, it would take the average subject 343 ms to accumulate evidence since a decision is reached when $\alpha = 1$ evidence unit is accumulated and subjects start the evidence accumulation process with .5 evidence units). In medium and high noise conditions, evidence accumulation rate intercepts were 1.30 and 0.86 evidence units per second respectively. Non-decision time intercepts were 340 ms in low noise conditions, 425 ms in medium noise conditions, and 440 ms in high noise conditions.

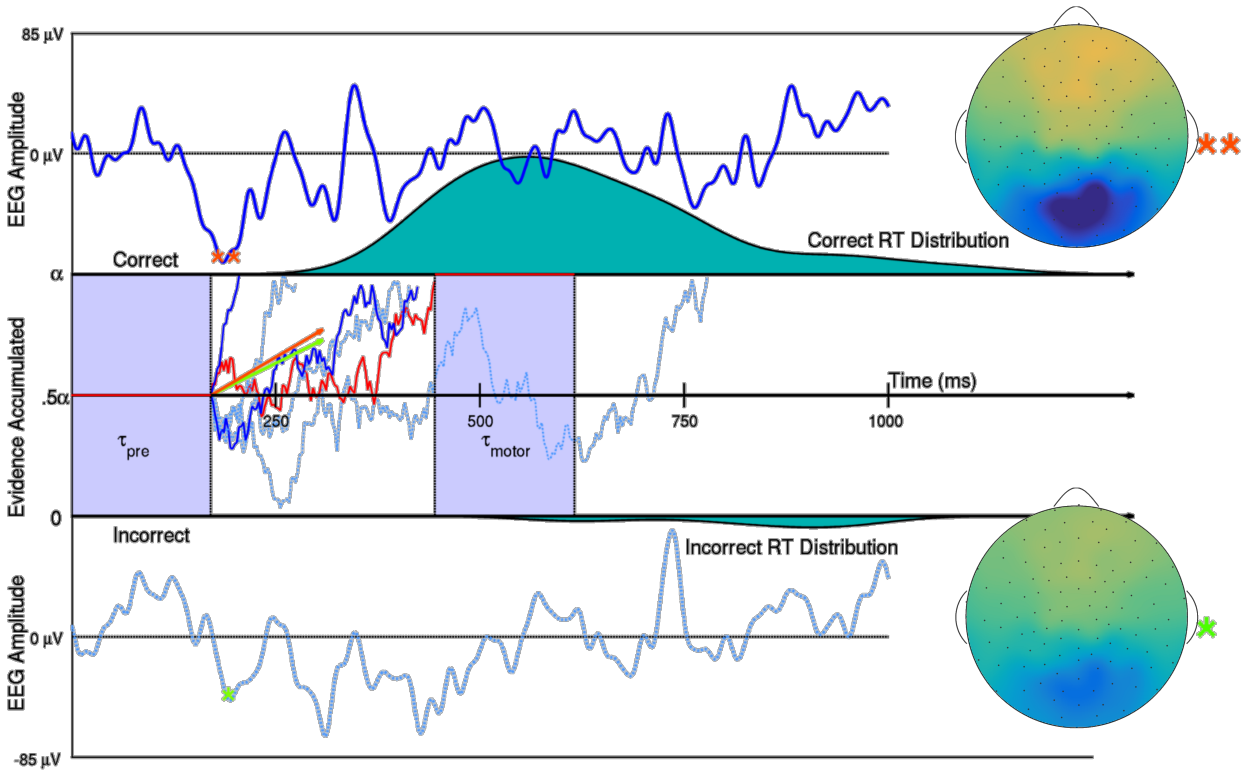

Figure 3.5: Two trials of Subject 10's SVD weighted EEG (Top and Bottom with bounds -85 to 85 μV) and representations of this subject's evidence accumulation process on 6 low noise trials (Middle). Evidence for a correct response in one example trial (denoted by the red line) first remains neutral during an initial period of visual preprocessing time τ_{pre}. Then evidence is accumulated with an instantaneous evidence accumulation rate of mean δ (the drift rate) and standard deviation ς (the diffusion coefficient) via a Wiener process. The subject acquires either $\alpha = 1$ evidence unit or 0 evidence units to make a correct or incorrect decision respectively. After enough evidence is reached for either decision, motor response time τ_{motor} explains the remainder of that trial's observed reaction time. The 85^{th} and 15^{th} percentiles of Subject 10's single-trial drift rates $\delta_{i,10,1}$ in the low noise condition are shown as orange and green vectors, such that it would take 253 and 299 ms respectively to accumulate the .5 evidence accumulation units need to make a correct decision if there was no variance in the accumulation process. The larger drift rate is a linear function of the larger single-trial N200 amplitude (**), while the smaller drift rate is a linear function of the smaller N200 amplitude (*). The scalp activation (SVD weights multiplied by one trial's N200 amplitude) of this subject's response to the visual signal ranges from -13 to 13 μV on both trials. The two dark blue and red evidence time courses were randomly generated trials with the larger drift rate. The three dotted, light blue evidence time courses were randomly generated trials with the smaller drift rate. True Wiener processes with drifts $\delta_{i,10,1}$ and diffusion coefficient $\varsigma_{10,1}$ were estimated using a simple numerical technique discussed in Brown et al. (2006).

To understand the degree of influence of EEG on model parameters, approximate condition level evidence accumulation rates and non-decision times were calculated and then compared to the intercept of the respective parameter. Taking the mean peak positive and peak negative amplitudes and latencies across all subjects and trials in each noise condition and multiplying by the median posterior of the effects, it was found that evidence accumulate rate in low noise was 1.90 evidence units per second, 1.65 evidence units per second in medium noise, and 1.35 evidence units per second in high noise. It was also found that non-decision time was 393, 400, and 425 ms in the low, medium, and high noise conditions. The intercepts of non-decision time thus described approximately 86%, 94%, and 96% of the true condition means in low, medium, high noise conditions respectively. However, the intercepts of the drift rates only described approximately 77%, 79%, and 63% of the true condition means in low, medium, and high noise conditions respectively. While this gives the reader an idea of the strength of the influence of single-trial EEG measures on the parameters, better evaluations of the degree of effects are presented below.

3.3.2 Effects of attention on non-decision time in low-noise conditions

Strong evidence was found to suggest that in low noise conditions single-trial non-decision times τ_{ijk} are positively linearly related to delays in the EEG response to the visual signal as measured by the latency of the negative peak (N200) following stimulus onset. A probability greater than 99% of the condition-level effect being greater than zero in all subjects was found. This relationship to an EEG signature 150-275 ms post stimulus onset suggests an effect on preprocessing time rather than motor-response time. By exploring the posterior distribution of the mean effect across-participants $\mu_{(\gamma_\tau)j}$, it is inferred that non-decision time increases 12 ms (the posterior median) when there is a 40 ms increase in the latency of the single-trial negative peak (where 40 ms was the standard deviation across all trials and

subjects) in the low noise condition, with a 99% credible interval of 3 to 21 ms. **Figure 3.6** shows the per-subject effects of signal N200 latency on non-decision time in the low noise condition. No evidence was found to suggest that the signal N200 latency affected non-decision time in medium nor high noise conditions. 95% credible intervals for these increases in the subject mean non-decision time for 40 ms increased N200 delays were -9 to 4 ms and -8 to 3 ms respectively. No evidence was found to suggest that attentional delay to the noise, the noise P200 latency, affected non-decision time.

Weak evidence was found to suggest that magnitude of the response to the stimulus affects non-decision time in the low noise condition. The posterior median suggests that a 26.83 μV (the standard deviation) decrease in magnitude of the negative peak (i.e. moves the negative peak towards zero) leads to a 11 ms increase in non-decision time. The 95% credible interval of this effect of N200 signal magnitude on non-decision time was a 2 to 21 ms. No evidence was found to suggest that the magnitude in the medium and high noise conditions affected non-decision time.

3.3.3 Effects of attention on evidence accumulation

Evidence was found to suggest that per-trial response to the visual signal (measured by the negative peak, N200, amplitude) is positively correlated with per-trial evidence accumulation rates δ_{ijk} in each condition. In the low noise condition, $\mu_{(\gamma_{(\delta)})j}$, which describes the across-subject mean of the effect of negative peak on drift rate, was found to have a 95% credible interval of .02 to .34 evidence units per second increase (where it takes $\alpha = 1$ evidence unit to make a decision) and a posterior median of .17 evidence units per second increase for each magnitude increase (i.e. away from zero) of 26.83 μV, the standard deviation of the negative peak. Given the same magnitude increase, the posterior median of the effects in the medium and high noise conditions were .13 and .14 evidence units per second respectively with 95%

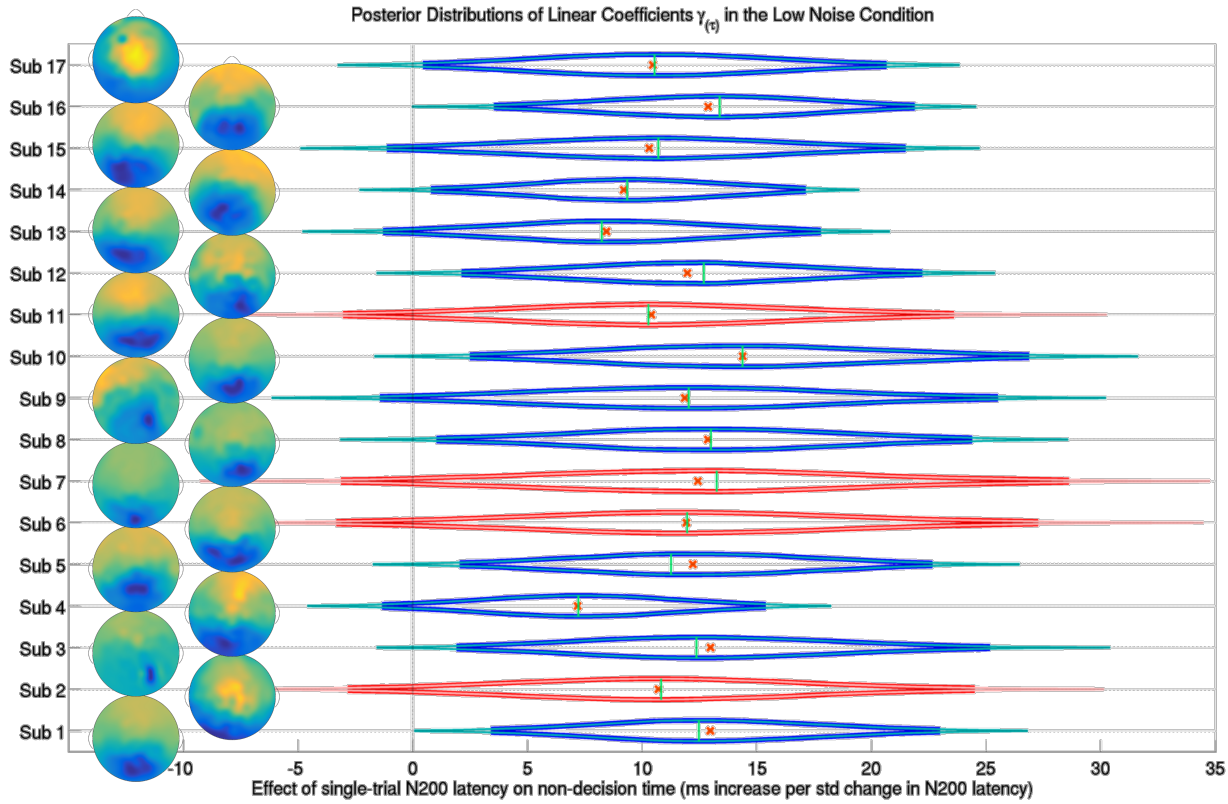

Figure 3.6: The posterior distributions of the effect of a trial's N200 latency during the response interval (onset attention latency to the signal stimulus) on trial-specific non-decision times τ_{ijk} for each subject j in the low noise condition $k = 1$. Subjects 2, 6, 7 and 11 were left out of the training set and their predicted posterior distributions are shown in red. Thick lines forming the distribution functions represent 95% credible intervals while thin lines represent 99% credible intervals. Crosses and vertical lines represent posterior means and modes respectively. Also shown are the topographic representations of the channel weights of the first SVD component of each subject's response interval ERP, indicating the location of single-trial N200s over occipital and parietal electrodes. Evidence suggests that longer attentional latencies to the signal, N200 latencies, are linearly correlated with longer non-decision times in the low noise condition.

credible intervals $-.01$ to .28 and 0 to .28 respectively.

Strong evidence was found to suggest that the magnitude of the positive peak of the response to the *visual noise* during the cue interval affected the future evidence accumulation rate in the medium noise and possibly high noise conditions. The median of the posterior distribution of the condition-level effect was .20 evidence units per second when there was a 27.67 μV increase, the standard deviation of the peak magnitude. A 99% credible interval of this effect was .04 to .32 evidence units per second. **Figure 3.7** shows the effects of the noise P200 amplitudes on specific subjects' single-trial drift rates in the medium noise condition. The probability of there being an effect of this P200 amplitude during the cue interval in the *high* noise condition was 94.6% (i.e. the amount of the posterior density of the condition-level effect above zero). The median of the posterior distribution of this effect was .09 evidence units per second with a 95% credible interval of $-.02$ to .22 evidence units per second when there is a 27.67 μV increase in a high noise trial.

3.3.4 Cross-validation

In-sample and out-of-sample posterior predictive coverage plots of correct-RT distributions for each condition and subject are provided in the Supplementary Materials. All three models perform well at predicting correct-RT distributions and overall accuracy of training data (i.e. *in-sample* prediction; see **Table 3.3** and Discussion section). However by cross-validation we found that the addition of single-trial EEG measures of attentional onset improved *out-of-sample* prediction of accuracy and correct reaction time distributions of *known* subjects (i.e. those subjects who had 2/3 of their trials used in the training set). R^2_{pred} indicates the percentage of variance explained by prediction in the given statistic across subjects. **Table 3.1** contains R^2_{pred} values for accuracy as well as summary statistics of correct-RT distributions of known subjects. Prediction was improved when using Model 3 for these

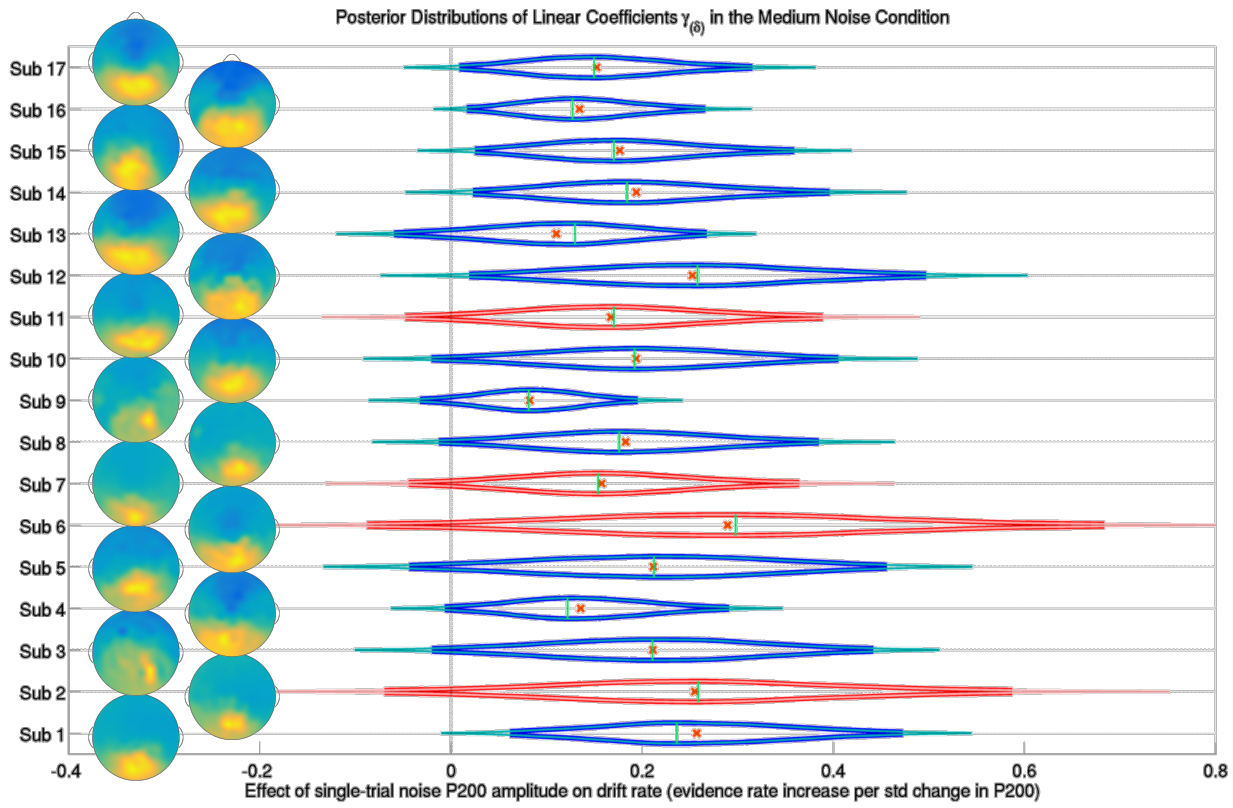

Figure 3.7: The posterior distributions of the effect of a trial's P200 amplitude during the cue interval (onset of attention to the noise stimulus) on trial-specific evidence accumulation rates δ_{ijk} for each subject j in the medium noise condition $k = 2$. Subjects 2, 6, 7 and 11 were left out of the training set, their predicted posterior distributions are shown in red. Thick lines forming the distribution functions represent 95% credible intervals while thin lines represent 99% credible intervals. Crosses and vertical lines represent posterior means and modes respectively. Also shown are the topographic representation of the channel weights of the first SVD component of each subject's cue interval ERP, indicating the location of single-trial P200s over occipital and parietal electrodes. Evidence suggests that the effect of the attention to the noise, reflected in P200 amplitudes, positively influenced the drift rate of each subject in each trial, in the medium and high noise conditions.

subjects with at least 77.3% of variance in correct-RT medians being explained by out-of-sample prediction, but Model 2 performed almost as well in comparison to Model 1, the model without single-trial EEG inputs. Model 2 was able to predict at least 76.3% of the variance in correct-RT medians while Model 1 was able to predict at least 74.5% of the variance in correct-RT medians. In the low noise condition, Model 2 did not improve upon Model 1's explanation of variance in subject-level accuracy but better predicted accuracy in the other two conditions.

Larger gains in out-of-sample prediction were found for *unknown* subjects (i.e. those subjects who were not used in the training set). These improvements were particularly pronounced in the low noise condition. Model 2 outperformed Model 3, which outperformed Model 1 in turn, as shown in **Table 3.2**. From these results it is clear that Model 2 was the best model for out-of-sample prediction overall, especially for new subjects. In the low noise condition, Model 2 was able to explain 22.1% of between-subject variance in correct-RT 25^{th} percentiles in the low-noise condition while Model 1 was not able to predict any between-subject variance in this measure. While the included single-trial EEG measures in these type of models do not perform as well as new subject prediction when *subject-average* EEG measures of attention are included (Nunez et al., 2015), single-trial EEG does improve prediction. The improvements in R^2_{pred} across models suggest that it is possible that similar models with more single-trial measures of EEG could explain new subjects' accuracy and correct-RT distributions.

3.3.5 P200 and N200 localizations

Both the subject-average P200 and N200 components were localized in time and space. All *single-trial* P200 and N200 amplitudes and latencies (i.e. single-trial peak positive and negative amplitudes 150 to 275 ms in the noise and response interval respectively) were

Prediction of new data from *known* subjects

		Model 1 Comparison	Model 2 EEG-δ,τ	Model 3 EEG-δ,τ,ς
Low	25^{th} $\mathbf{t_1}$ Percentile	81.6%	85.2%	85.6%
	$\mathbf{t_1}$ Median	74.5%	76.3%	77.3%
	75^{th} $\mathbf{t_1}$ Percentile	60.2%	63.3%	63.8%
	Accuracy	24.5%	20.8%	27.7%
Medium	25^{th} $\mathbf{t_1}$ Percentile	84.1%	85.1%	86.1%
	$\mathbf{t_1}$ Median	86.8%	87.6%	88.9%
	75^{th} $\mathbf{t_1}$ Percentile	63.1%	68.2%	69.3%
	Accuracy	58.1%	63.5%	63.5%
High	25^{th} $\mathbf{t_1}$ Percentile	73.0%	76.4%	76.6%
	$\mathbf{t_1}$ Median	77.4%	76.8%	77.8%
	75^{th} $\mathbf{t_1}$ Percentile	71.0%	74.2%	74.2%
	Accuracy	46.3%	48.9%	51.3%

Table 3.1: Percentage of across-subject variance explained by out-of-sample prediction (R^2_{pred}) for accuracy and summary statistics of correct-RT distributions of those subjects' that were included in the training set. 13 of the subjects' data were split into 2/3 training and 1/3 test sets. Posterior predictive distributions that predicted test set behavior were generated for 13 of the subjects by drawing from posterior distributions generated by the training set. In the Low, Medium, and High noise conditions, the 25^{th}, 50^{th} (the median), and 75^{th} percentiles and means were predicted reasonably well by the model without single-trial measures of EEG, Model 1. However including single-trial measures of EEG improved prediction of correct-RT distributions, especially in the Low noise condition, with Model 3 (which assumes evidence accumulation rate, non-decision time, and evidence accumulation variance vary with EEG per-trial) only slightly outperforming Model 2 (which assumes evidence accumulation rate and non-decision time vary with EEG per-trial).

Prediction of new data from *new* subjects

			Model 1 Comparison	Model 2 EEG-δ,τ	Model 3 EEG-δ,τ,ς
Low	25^{th}	$\mathbf{t_1}$ Percentile	−9.8%	22.1%	17.9%
		$\mathbf{t_1}$ Median	−37.6%	−11.6%	−14.3%
	75^{th}	$\mathbf{t_1}$ Percentile	−67.3%	−46.7%	−47.6%
		Accuracy	−8.1%	−22.8%	−65.9%
Medium	25^{th}	$\mathbf{t_1}$ Percentile	−5.3%	8.4%	5.3%
		$\mathbf{t_1}$ Median	−30.7%	−22.5%	−25.7%
	75^{th}	$\mathbf{t_1}$ Percentile	−62.6%	−53.9%	−56.2%
		Accuracy	−13.1%	−35.5%	−67.7%
High	25^{th}	$\mathbf{t_1}$ Percentile	−1.7%	8.2%	7.7%
		$\mathbf{t_1}$ Median	−12.0%	−4.3%	−2.9%
	75^{th}	$\mathbf{t_1}$ Percentile	−37.6%	−25.5%	−23.7%
		Accuracy	−4.3%	−14.1%	−29.1%

Table 3.2: Percentage of across-subject variance explained by out-of-sample prediction (R^2_{pred}) for accuracy and summary statistics of new subjects' and correct-RT distributions. Posterior predictive distributions were generated for 4 new subjects by drawing from condition level posterior distributions. Most R^2_{pred} measures are negative because the amount of variance in prediction was greater than the variance of the measure across subjects; however the relative values from one model to the next are still informative about the improvement in prediction ability. The model without single-trial EEG measures, Model 1, does not predict new subjects' correct-RT distributions. Models *with* single-trial EEG measures of onset attention, Model 2 and Model 3, can predict some variance of the new subjects' 25^{th} percentiles, with Model 2 outperforming Model 3.

averaged across trials for each subject. Localization in time and on the brain are based on these across-trial averages. The subject mean and standard deviation of the P200 latency during the cue interval was 220 ± 12 ms while the mean and standard deviation the of N200 latency during the response interval was 217 ± 6 ms. Although these latencies differ slightly from traditional P2 and N1 findings (Luck et al., 2000), when viewing the event-related waveforms over all trials it is clear that the P2 and N1 are influenced by the single-trial measures. As an example, every single-trial evoked response of Subject 12 to the noise and signal are shown in **Figure 3.4**, sorted by peak P200 amplitude in the cue interval and sorted by peak N200 amplitude in the response interval. The P200 and N200 latencies of this subject correspond to traditional P2 and N1 components.

While EEG localization is an inexact process that is unsolvable without additional assumptions, the surface Laplacian has been shown to match closely to simulated cortical activity using forward models (i.e the mapping of cortical activity to scalp potentials) and have shown consistent results when used with real EEG data (Nunez and Srinivasan, 2006). Unlike 3D solutions, projections to the surface of the cortex (more accurately, the dura) are theoretically solvable, and have been used with success in past studies (see Nunez et al., 1994, for an example).

In this study we have found surface spline-Laplacians (Nunez and Pilgreen, 1991) on the realistic MNI average scalp (Deng et al., 2012) of the mean positive peak during the cue interval (the P200) and the mean negative peak during the response interval (the N200) by averaging over trials and subjects. The surface Laplacians were then projected onto one subject's cortical surface using Tikhonov (L2) regularization and a Finite Element (FE; Pommier and Renard, 2005) forward model to the MNI 151 average head, maintaining similar distributions of activity of the surface Laplacians on the cortical surface. The subject's brain was then labeled using the Destrieux cortical atlas (Fischl et al., 2004). Cortical topographic maps of both peaks are given in **Figure 3.8**. Because we expect the majority of the Laplacian

to originate from superficial gyri (Nunez and Srinivasan, 2006), we have localized both the positive and negative peaks only to maximally active gyri. This localization suggested that both the P200 and N200 were in the following extrastriate and parietal cortical locations: right and left middle occipital gyri, right and left superior parietal gyri, right and left angular gyri, the left occipital inferior gyrus, the right occipital superior gyrus, and the left temporal superior gyrus. Although we should note that the exact localization must have some errors due to between-subject variance in cortex and head shape and between-subject variance in tissue properties.

Brain regions found using this cortical-Laplacian method point to activity in early dorsal and ventral pathway regions associated with visual attention (Desimone and Duncan, 1995; Corbetta and Shulman, 2002; Buschman and Miller, 2007) and decision making (Mulder et al., 2014). Corroborating our findings, White et al. (2014) found that blood-oxygen-level dependent (BOLD) activity in the right temporal superior gyrus, right angular gyrus, and areas in the right lateral occipital cortex (e.g. the right middle occipital gyrus) correlated with non-decision time during a simple visual and auditory decision making task. It was hypothesized that this activity was due to motor preparation time instead of visual preprocessing time (White et al., 2014); however the time-scale of BOLD signals does not provide additional knowledge to separate visual preprocessing time from motor preparation time. Informed by EEG, BOLD signals associated with evidence accumulation rates have been previously localized to right and left superior temporal gyri and lateral occipital cortical areas, thought to correspond to early bottom-up and late top-down decision making processes respectively during a visual face/car discrimination task (Philiastides and Sajda, 2007). The right and left middle occipital gyri have also previously been shown to contribute to evidence accumulation rates during a random dot motion task (Turner et al., 2015).

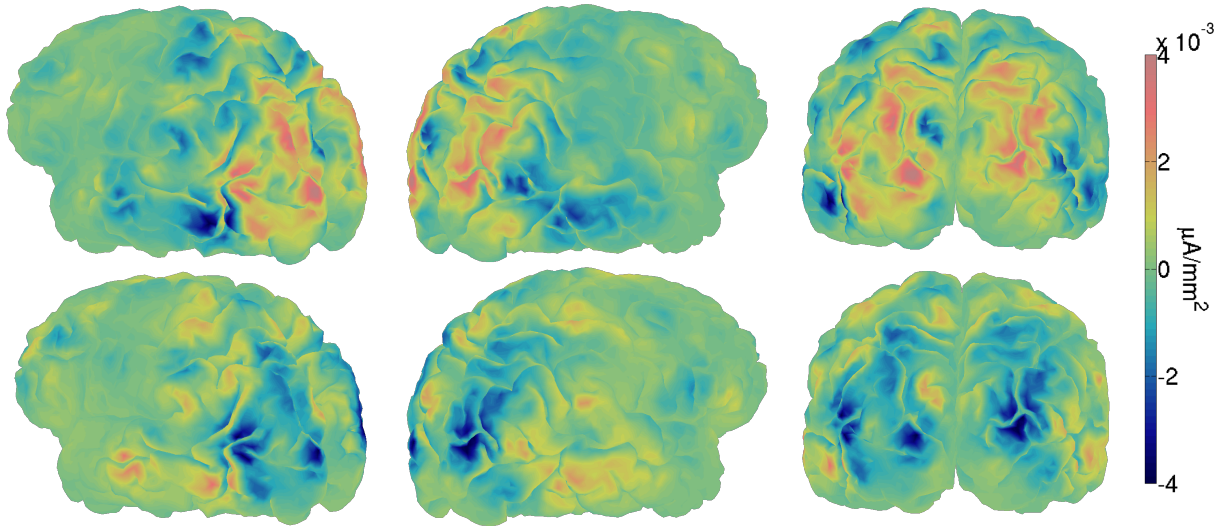

Figure 3.8: Right and left sagittal and posterior views of localized single-trial P200 evoked potentials during during the cue interval (Top) and localized single-trial N200 evoked potentials during the response interval (Bottom) averaged across trials and subjects. The cortical maps were obtained by projecting MNI-scalp spline-Laplacians (Nunez and Pilgreen, 1991; Deng et al., 2012) onto a subject's anatomical fMRI image via Tikhonov (L2) regularization, maintaining similar distributions of activity of the surface Laplacians on the cortical surface. Blue and orange regions in microamperes per mm^2 correspond to cortical areas estimated to produce negative and positive potentials observed on the scalp respectively. These two particular projections of the Laplacians suggest that P200 and N200 activity occurs in extrastriate cortices and areas in the parietal lobe.

3.4 Discussion

3.4.1 Attention influences perceptual decision making on each trial

The results of this study suggest that fluctuations in attention to a visual signal accounts for some of the trial-to-trial variability in the brain's speed of evidence accumulation on each trial in each condition. There is also evidence to suggest that increased response to the competing visual noise increases the brain's speed of evidence accumulation, but only in medium and high noise conditions. Although a simple explanation of this effect would be differences in trial-to-trial arousal, we note that the effect only occurs in medium-noise and high-noise conditions. We have previously found evidence that noise suppression during the cue interval predicts enhanced drift rates based on the subject-average SSVEP responses in these data (Nunez et al., 2015). Thus, our effect may reflect the attention the subject places on the cue, which determines whether to engage mechanisms of noise suppression, but we could not directly assess this possibility as new methods must be developed to measure SSVEPs on single trials. This would reflect a hypothesis based on the Perceptual Template Model that predicts that subjects will suppress attention to visual distractors during tasks of high visual noise (Lu and Dosher, 1998).

We assume that the effect of N200 latency on non-decision time during the response interval is on preprocessing time instead of motor response time. There was no a priori reason to believe that an attention effect that takes place 150-275 ms after stimulus onset would affect the speed of efferent signals to the muscle, given that the response times were at least 500 ms. These findings lead us to conclude that the effect of attention in response to the signal in low noise conditions is to reduce preprocessing delay time. This appears consistent with predictions of signal enhancement in low-noise conditions in the Perceptual Template Model.

However, we note there is not a perfect equality between N200 latency and preprocessing time (the condition and subject-level coefficient posteriors are not centered on 1). And the identification of both preprocessing time and motor response time is not possible with a drift-diffusion model fit without additional assumptions or external inputs such as EEG.

Reaction time (RT) and choice behavior during visual decision making tasks are well characterized by models that assume a continuous stochastic accumulation of evidence. And many observations of increasing spike-rates of single neuron action potentials lend support to this stochastic theory of evidence accumulation on a neural level (Shadlen and Newsome, 1996, 2001). Recently some macroscopic recordings of the cortex have shown that increasing EEG potentials ramping up to P300 amplitudes are correlates of the stochastic accumulation of evidence (O'Connell et al., 2012; Philiastides et al., 2014; Twomey et al., 2015). It has been hypothesized that this EEG data reflects the evidence accumulation process itself (or a mixture of this process with other decision-making correlates) and not a correlated measure such as top-down attention. This hypothesis leads to the natural prediction that single-trial drift rates are explained by single-trial P300 slopes. However within a small region of the cortex, neurons will have diverse firing patterns during the decision making process, only some of which are observed to have increasing spiking-rate behavior indicative of stochastic evidence accumulation (Meister et al., 2013). The properties of volume conduction through the cortex, skull, and skin only allow for synchronous post-synaptic potentials to be observed at the scalp (Nunez and Srinivasan, 2006; Buzsaki, 2006). Therefore an increasing spike-rate as observed on the single-neuron level is not likely to be observed as an increasing waveform in EEG recorded from the scalp. We also would expect the evidence accumulation process to terminate before the response time since a portion of the response time must be dedicated to the motor response after the decision is made. This may not be the case for the ramping P300 waveform on single-trials even though it is predictive of model parameters (Philiastides et al., 2014). If the stochastic evidence accumulation process was truly reflected in the ramping of EEG, a testable prediction would be that the variance around the mean rate of the

P300 ramp on each trial would be linearly related to the diffusion coefficient ς, in addition to single-trial P300 slopes being linearly related to the drift rate δ. There are many EEG measures thought to be related to *attention* such as event related potential (ERP) components, power in certain frequency bands, and steady-state visual evoked potential (SSVEP) responses. It is likely that EEG measures that share similar properties with stochastic evidence accumulation processes are in fact due to these correlates of attention or other forms of cortical processing that can influence the decision making process.

3.4.2 External predictors allow for trial level estimation of diffusion model parameters

The Wiener distribution (i.e. the diffusion model) used in this study does not incorporate trial-to-trial variability in drift rates *within the probability density function* as assumed by Ratcliff (1978). Instead we assume that each trial's drift rate is *exactly equal to* a linear function of EEG data and use an evidence accumulation likelihood function that does not assume drift rates vary trial-to-trial by any other means. Per-trial non-decision times and diffusion coefficients were also assumed to be *exactly equal to* linear functions of EEG data.

Per trial estimates of diffusion model parameters cannot be obtained without imposing constraints or including external inputs. In this study, we have shown that the single-trial P200 and N200 attention measures can be used to discover per trial estimates of all three free parameters, non-decision time, drift rate, and the diffusion coefficient. Non-decision time, the drift rate, and the *evidence boundary* could also be modeled as per trial estimates of external inputs, as would be useful in other speeded reaction time tasks where external per-trial physiological measurements are available. Other possible per-trial external inputs that could be used include: magnetoencephalographic (MEG) measures, functional magnetic resonance imaging (fMRI) measures, physiological measurements such as galvanic skin

response (GSR), and near infrared spectroscopy (NIRS), where each modality may have multiple external inputs (e.g. multiple linear EEG regressors of single-trial parameters, as in this study). The more external inputs correlated with single-trial parameters included in the model, the better the single-trial estimate of the parameters will be. This will allow decision model researchers to better explore the efficacy of the diffusion model by comparing single-trial estimates of parameters.

3.4.3 Behavior prediction and BCI applications

We have observed that single-trial measures of EEG in a hierarchical Bayesian approach to decision-making modeling improves overall accuracy and correct-RT distribution prediction for subjects with observed behavior. This paradigm also lead to significant improvement in overall correct-RT distribution prediction for those subjects whose behavior was *missing*. That is, we have shown that a new subject's correct-RT distributions can be predicted when only their single-trial EEG is collected, given that other subjects' EEG and behavior has been analyzed. If the goal of a future project is solely prediction (and no explanation of the cognitive or neural process is desired, as was in this paper), a whole host of single-trial EEG measures could be included in a hierarchical model of decision making, using perhaps a simpler model of decision making such as the linear ballistic accumulator model to ease analysis (e.g. Forstmann et al., 2008; Ho et al., 2009; van Maanen et al., 2011; Rodriguez et al., 2015) or a more complicated model to improve prediction. The set of single-trial measures could include: ERP-like components as we discussed in this paper, measures of evoked amplitudes in certain frequency bands, and measures of steady-state visual evoked potentials.

However, as observed in **Table 3.3**, we have not shown prediction of reaction times for when subjects committed errors because 1) few errors were committed by the subjects in the

presented data and 2) the type of model we used has been shown to explain incorrect-RT distributions well only when intrinsic trial-to-trial variability in evidence accumulation rates (as opposed to *extrinsic* due to neural regressors, for which it has not been shown) is included in the likelihood function (Ratcliff, 1978; Ratcliff and McKoon, 2008). In future work we plan to compare the predictive ability of more complicated models containing *both* intrinsic trial-to-trial variability and extrinsic trial-to-trial variability in evidence accumulation rate due to external neural measures.

The prediction ability of the presented models for accuracy and correct-RT distributions may have implications for Brain-Computer Interface (BCI) frameworks, especially in paradigms which attempt to enhance a participant's visual attention to particular task to improve reaction time. To maximize prediction, every EEG attention measure should be included, and the predictors that offer the best out-of-sample prediction within initial K-fold validation sets should be included. After collection of behavior and EEG from a few participants and a hierarchical Bayesian analysis of the data, later participants' single-trial preprocessing times or evidence accumulation rates could be predicted using only single-trial EEG measures. Conceivably this would allow for trial-by-trial intervention in order to enhance a participant's attention during the task, perhaps using neural feedback (e.g. via direct current stimulation or transcranial magnetic stimulation) during those trials in which a participant is predicted to be slow in their response because of a small evidence accumulation rate (e.g. when an N200 magnitude is small) or a slow preprocessing time (e.g. when an N200 latency is long). Although whether the participant could use such feedback in time to affect reaction time and accuracy remains to be tested.

Prediction of *training* data from *known* subjects

		Model 1 Comparison	Model 2 EEG-δ,τ	Model 3 EEG-δ,τ,ς
Low	25^{th} $\mathbf{t_1}$ Percentile	96.5%	97.1%	97.8%
	$\mathbf{t_1}$ Median	96.6%	96.5%	97.3%
	75^{th} $\mathbf{t_1}$ Percentile	91.4%	93.4%	94.5%
	Accuracy	95.1%	95.2%	97.3%
	$\mathbf{t_0}$ Median	−118.8%	−108.3%	−111.6%
Medium	25^{th} $\mathbf{t_1}$ Percentile	86.0%	87.5%	88.6%
	$\mathbf{t_1}$ Median	95.9%	95.6%	96.3%
	75^{th} $\mathbf{t_1}$ Percentile	84.7%	89.4%	90.1%
	Accuracy	90.7%	94.1%	95.3%
	$\mathbf{t_0}$ Median	−163.9%	−158.6%	−163.9%
High	25^{th} $\mathbf{t_1}$ Percentile	85.4%	87.3%	86.7%
	$\mathbf{t_1}$ Median	93.1%	92.5%	92.9%
	75^{th} $\mathbf{t_1}$ Percentile	79.1%	83.8%	84.0%
	Accuracy	95.9%	97.4%	95.9%
	$\mathbf{t_0}$ Median	−73.4%	−71.2%	−76.4%

Table 3.3: Percentage of variance across subjects explained by in-sample prediction (R^2_{pred}) for summary statistics of known subjects' accuracy-RT distributions. All three models fit accuracy and correct-RT $\mathbf{t_1}$ data very well, explaining over 92% of median correct-RT and over 90% of accuracy in each condition. However none of the models explain incorrect-RT $\mathbf{t_0}$ distributions well, a known problem for simple diffusion models that can be overcome by including variable drift rates directly in the likelihood function (Ratcliff, 1978; Ratcliff and McKoon, 2008).

3.4.4 Neurocognitive models

The term "neurocognitive" has been used to describe the recent trend of combining mathematical behavioral models and observations of brain behavior to explain and predict perceptual decision making (Palmeri et al., 2017). The usefulness of combining behavioral models and neural dynamics has been motivated on theoretical grounds. Behavioral models suggest links between subject behavior and cognition while laboratory observation and neuroimaging can suggest links between neural dynamics and cognition. The combination of these methods then provides a predictive chain of neural dynamics, cognition, and behavior. Another obvious benefit is the inference gain when predicting missing data. That is, we will be able to better predict behavior when brain activity is available. This is especially true when using hierarchical Bayesian models as they maintain uncertainty in estimates through different levels of the analysis (Vandekerckhove et al., 2011; Turner et al., 2013). While there are a variety of methods using cognitive models to find cognitive correlates in the brain dynamics (see Turner et al., 2017, for a review of these methods), some studies do not further constrain the cognitive models by informing those models with known neural links to specific cognitive processes. In this chapter, and our previous study of individual differences (Chapter 2; Nunez et al., 2015), we demonstrate another important use of neural data in cognitive models. Independent *neural* measures of cognitive processes, such as attention, can be used to better understand how cognition influences the mechanisms of behavior, furthering explanation and prediction of the cognitive process.

Data and code sharing

Pre-calculated EEG measures, raw behavioral data, MATLAB stimulus code, JAGS code, and an example single-trial EEG R script are available upon request and in the following repository (as of February 2016) if their use is properly cited.

https://github.com/mdnunez/mcntoolbox/

Chapter 4

EEG measures of neural processing speed reflect visual encoding time

4.1 Introduction

While research exists on the time course of primates' visual system in response to external visual inputs (Schmolesky et al., 1998), the time course of humans' visual response remains largely unexplored due to the invasiveness of prevailing techniques (e.g. single-cellular recordings). The electroencephalogram (EEG) is a noninvasive technique that records cortical synaptic activity that is synchronized across the cortex and is thought to represent higher-level function in electrophysiology (Buzsaki, 2006; Nunez and Srinivasan, 2006; Nunez et al., 2016). In an exploratory analysis, we propose a method of estimating and verifying individuals' visual encoding time using traditional EEG measures before decision processing of visual information. These techniques are based on joint estimation of traditional event-related potentials (ERPs) and cognitive models of perceptual decision making that describe human reaction time and choice distributions. Using these methods, both individual differences and

within-individual (trial-to-trial) variability in human visual encoding time can be found that predict observed individual differences and within-individual differences in reaction time.

A multitude of theoretical and experimental work studying micro-scale electrophysiological decision making, i.e. single neurons that track the cognitive process and neural computation of decision making (Shadlen and Newsome, 2001; Gold and Shadlen, 2007), while there exists only a small amount of theoretical work to inform experimental results on the *meso-* or *macro*-scale level in electrophysiology (e.g see Martin et al., 2010; O'Connell et al., 2012; Frank, 2015). In this Chapter I explore the veracity of a marco-level electrophysiological measure of visual encoding by drawing inference from a well-established cognitive model of human decision-making behavior. In particular I explore an early visual stimulus-locked negative peak latency in parietal and occipital electrodes, dubbed an "N200" ERP component to label its peak direction (negative) and its approximate latency in milliseconds (approximately 200 ms). By exploring the predictability of this stimulus-locked EEG measure embedded in cognitive models of perceptual decision making, inference about the neural process of decision making can be better understood.

4.1.1 Previous explorations of electrophysiological measures of visual encoding

The speed of human visual system has previously been measured with event-related potentials (Thorpe et al., 1996) using clever condition differences between in go/no-go tasks. Thorpe et al. (1996) estimated that visual processing in their particular task occurred at about 150 ms in humans. Other evidence has recently been found to suggest that N200 latencies influence the onset of decision making processes after initial visual encoding (Loughnane et al., 2016). In particular, Loughnane et al. (2016) found evidence that *contralateral* encoding time over parietal cortex are captured by the first negative peaks, "N1"s, in event-related

potentials (ERPs). In another study, single-trial EEG potentials thought to be generative of N1 peaks (computationally found using the algorithm from Anderson et al., 2016) were suspected to be related to a "preattentive" phase of decision-making, reflecting the time for attended visual information to reach extrastriate areas (Zhang et al., 2016). Further evidence in support of this theory are that single-neuron recordings of evidence accumulation from lateral intraparietal areas (LIP) in primates typically begin at similar time periods after an experimental stimulus is displayed (Shadlen and Kiani, 2013). Finally, Martin et al. (2010) found evidence that the same N1/N200 peak tracked non-decision time by looking at linear relationships between various ERP peak latency periods and residual reaction times.

Note that "N1" peaks are also referred to as "N200" peaks, referring to approximate time in milliseconds of occurrence (e.g. see Hong et al., 2009). Some literature also refers to similar peaks as the "N2" peak, unfortunately leading to confusion in the nomenclature (e.g. see Martin et al., 2010; Schubert et al., 2017). However the evidence does not seem to suggest that these potentials reflect different processes across visual studies and experiment types. In this chapter I will refer to "N200" when I refer to both the traditional trial-average ERP latencies and single-trial estimates of these latencies.

4.1.2 Separation of encoding, decision, and motor time during human decision making

Given the evidence in the literature, it is hypothesized that particular latencies of peaks in event related potentials (ERPs; EEG in response to stimuli at certain time points) predict decision making processes and, in particular, reflect encoding time of visual stimuli. Here *encoding time* is defined as the amount of time for visual processing to occur in the human brain before decision processes can begin. Behaviorally, trial-to-trial differences and subject-to-subject differences in encoding time are predicted to affect only trial-to-trial and

subject-to-subject differences in the onset of reaction time distributions, not the reaction time distribution shapes or accuracy, reflecting the fact that encoding does not affect the decision making process itself. Models of both trial-to-trial and session-to-session differences in encoding time will be evaluated with hierarchical diffusion models that include single-trial and session-level estimates of stimulus-locked EEG responses. A linear relationship between the predictor variable estimating speed of cortical response (either traditional trial-averaged ERP potential peaks or single-trial latency estimates as calculated from EEG) and estimates of encoding time from reaction time distributions is hypothesized that has a linear slope coefficient of 1, indicating that a 1 ms increase in cortical response time reflects a 1 ms increase in reaction time distributions.

4.2 Methods

4.2.1 Session-level observations

In order to evaluate the hypothesis, data was analyzed from two similar experiments. Data from Experiment 1 consisted of EEG recordings and behavioral observations from 12 unique subjects with 2 sessions of EEG each. Data from Experiment 2 consisted of EEG recordings and behavioral observations from 4 unique subjects with 7 sessions each. Sessions of EEG collection and task performance for each subject in Experiment 1 were separated by at least 24 hours. Sessions of EEG collection and task performance for each subject in Experiment 2 were separated by 1 week. Models 2 and 3 treat all EEG sessions as observations for the regression relationships, collapsing and disregarding subject differences. While subject-differences are important, subject differences were ignored in order to increase the observation count. The confirmatory study described in Chapter 5 contains a less heterogeneous data set with more unique subjects.

4.2.2 Experimental tasks

Subjects were tasked to perform a two-alternative forced choice task, specifically to determine whether a Gabor stimulus contained high or low spatial frequency content. Stimuli for Experiments 1 and 2 are built and displayed using the MATLAB Psychophysics toolbox (Psychtoolbox-3; www.psychtoolbox.org). Example stimuli for both Experiments 1 and 2 are given in **Figure 4.1**).

Subjects in the two experiments described performed visual-based tasks in a dark room on a $\sim$ 61 cm LCD monitor with a distance of 57 cm from retina to display. The monitor resolution was set as $1920 * 1280$ pixels with a 120 Hz refresh rate. Visual angle refers to the angular distance from a stationary subject's retina to an object in the visual field. A distance of 57 cm was chosen to ensure easy calculation of visual angles, with x cm equaling $\sim x^{\circ}$ visual angle. The subjects were told to maintain fixation throughout the task on a .25 cm (.25°) fixation spot while attending to to the stimulus.

Gabors are spatial grating patterns with a Gaussian falloff of contrast that produces maximal firing in certain neurons within the primary visual cortex (Webster et al., 1985). Subjects in both experiments were tasked with identifying the spatial frequency (either "high" or "low" spatial frequency) of large Gabors that are randomly rotated on each trial, matching observed Gabor templates (observed before the first session) with the spatial frequency observed on the current trial. In Experiment 2, the high and low spatial frequencies of the target Gabors were 2.4 and 2.6 cycles per degree visual angle (cpd) respectively. While in Experiment 1, the targets were randomly drawn from a $\mathcal{B}eta(13.5, 1.5)$ distribution or a $\mathcal{B}eta(1.5, 13.5)$ distribution with means 2.4 and 2.6 cpd respectively. Before and during the Gabor display on each trial, a disk of visual noise was displayed. In Experiment 1, checkerboard noise was displayed. While in Experiment 2, the noise was spatially bandpass filtered to include only masking noise centered around both 2 and 3 cpd was used, equally masking both the high

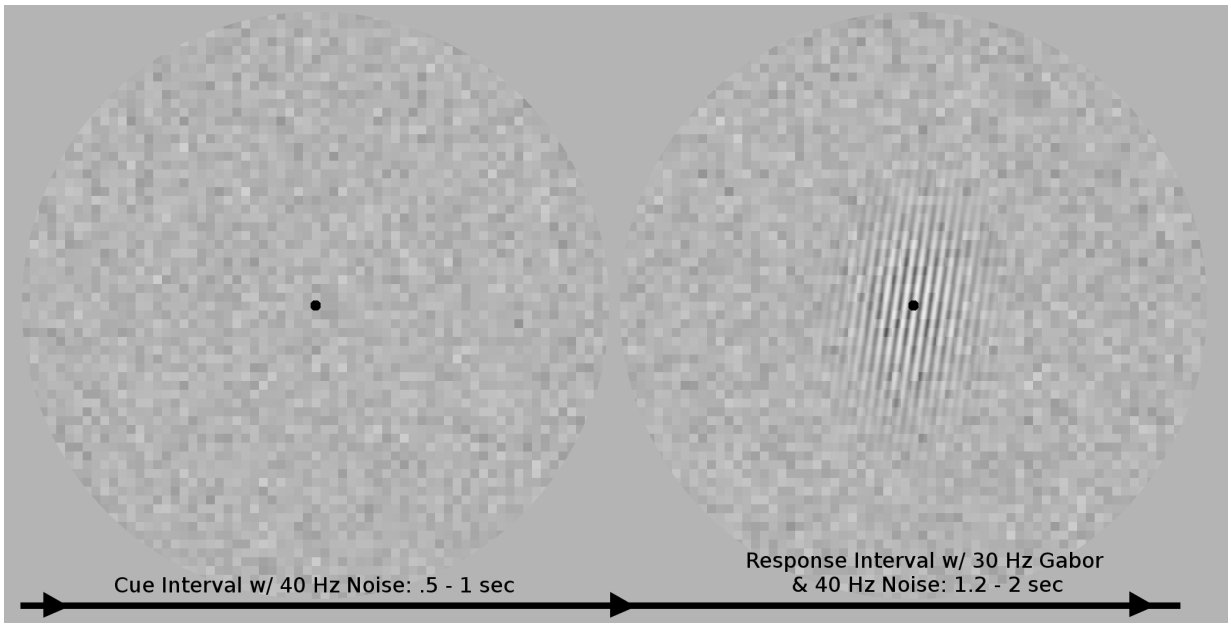

(a) Experiment 1

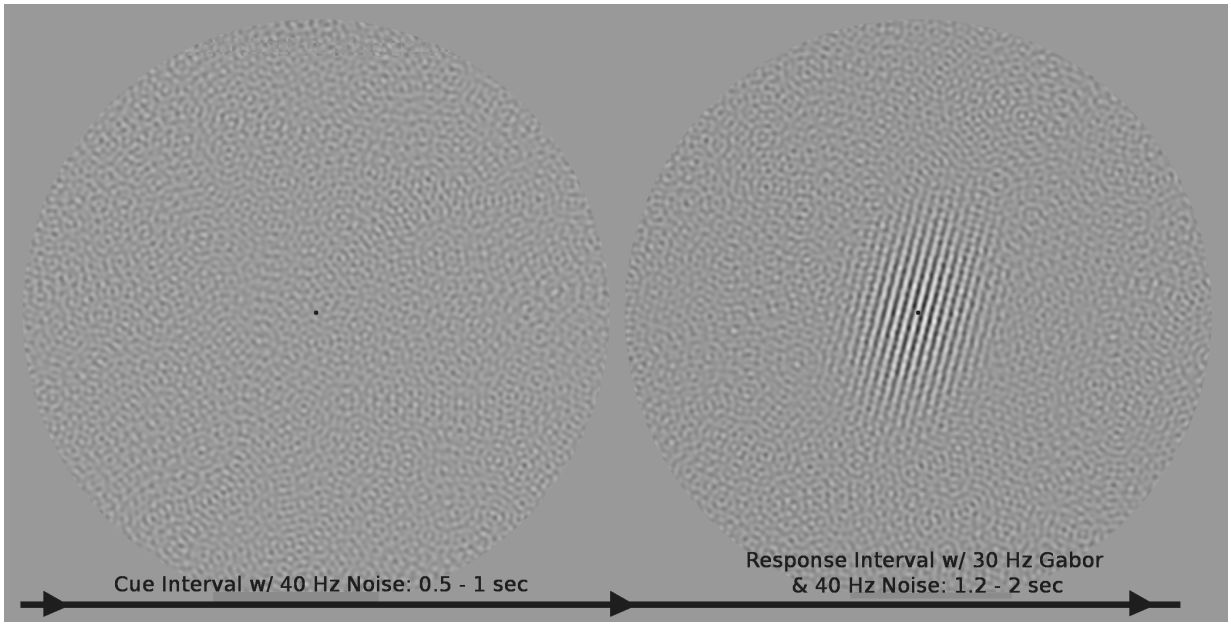

(b) Experiment 2

Figure 4.1: Example stimuli of the *cue* and *response* intervals of medium noise conditions from Experiments 1 (top) and 2 (bottom). In both experiments an SSVEP paradigm was be used in which the visual noise changed at 40 Hz and the Gabor signal flickered at 30 Hz to evoke 40 Hz and 30 Hz responses in electrocortical activity that track attention to the noise and signal stimuli respectively. Single-trial ERPs were measured both to the onset of visual signal in the *response* intervals.

and low spatial frequency targets. Both experiments contained three conditions of high, medium, and low visual noise contrast.

The time course of each trial was as follows: subjects were asked to fixate on a fixation spot (.25° visual angle) in the center of the screen throughout the experiment, visual noise changing at 40 Hz was displayed a variable length of time between 500 ms to 1000 ms during the *cue* interval, then the Gabor signal stimulus was flickered at 30 Hz embedded in the noise stimuli for 1200 ms to 2000 ms during the *response* interval. During the response interval, subjects are asked to respond as accurate as possible in the time allowed using a button box (using their left hand to respond for low spatial frequency targets and their right hand for high spatial frequency targets, although this will be reversed in some subjects for Experiment 2). Auditory feedback stimuli occurred after the response interval to indicate trial accuracy; auditory feedback was given in order to maintain subject vigilance. An entire session of task performance and EEG collect took subjects approximately 1 hour with breaks. Each session produced 8 blocks of 60 trials each for a total of 480 trials, with the noise contrast conditions intermixed within each block. Subjects were allowed to take breaks between blocks.

4.2.3 EEG recording

EEG was collected in both experiments using Electrical Geodesic, Inc.'s 128 electrode sensor net and a Net Amps 200 series amplifier. Electrical activity from the scalp was recorded at a sampling rate of 1000 samples per second and hardware band pass filtered to either a 1 to 50 Hz window (EEG data from Experiment 1) or a 1 to 100 Hz window (EEG data from Experiment 2). The hardware band pass filter in Experiment 2 was chosen purposely to maintain high frequencies so that broadband noise is submitted to an Independent Component Analysis (ICA) in order to aid artifact correction (i.e. removing electrical activity due to muscle, movement and environmental electrical influence) (Makeig et al., 1996; Nunez

et al., 2016). However both analyses yielded artifact-resistant measures after ICA that were used for further processing as detailed below.

4.2.4 Calculation of event related potentials

Event-related potentials (ERPs) to the onset of the Gabor patch during the *response* interval were analyzed using five key steps steps: 1) low pass filtering the EEG data to obtain cleaner estimates of evoked potentials, 2) shifting each trial's amplitude to rereference to a baseline (i.e. the average amplitude in a 100 ms window before the stimulus), 3) averaging across trials to obtain a traditional ERP estimate at 128 electrodes for each subject (minus those electrodes marked as artifactual), 4) taking a singular value decomposition (SVD) of the ERPs, and then 5) using the first component as an channel weighting function in order to obtain better estimate of the peaks and latencies for both session-specific trial-averaged ERPs and single-trial estimates of ERPs (although traditionally calculated ERPs at specific electrodes were also compared to this method yielding no difference in latency calculations).

The window to calculate minima in order to evaluate N200 latency responses was found empirically. Out of the following windows: 100 to 300 ms post-stimulus, 100 to 250 ms post-stimulus, 150 to 225 ms post-stimulus, and 150 to 275 ms post-stimulus, the 150 to 275 ms window post-stimulus was found to capture the N200 latency on the subject-level well in each condition because few edge-located minimums were detected in the trial-averaged N200 estimates (i.e. because only 3 estimates out of 147 estimates of trial-averaged N200 latency were detected at the boundary, it was thought that his window well captured the N200 latency). If either trial-averaged or single-trial N200 latencies were found with a minimum on the boundary, they were removed from the analysis. A minimum at either 150 or 275 ms was indicative that N200 latencies were not well estimated in that noise condition or on that trial.

The goal of the singular value decomposition (SVD) was to improve the signal-to-noise ratio of both traditional ERP estimates and single-trial estimates by weight-averaging adjacent electrodes as a spatial filter (Parra et al., 2005). SVD is the algorithm used my most principal component analysis (PCA) algorithms and produces non-stochastic, deterministic results. Better estimates of the traditional ERPs were obtained by finding the first principal component of the matrix of ERPs (samples T by channel C data, see the discussion in Chapter 3). The first principal component consists of both a weight vector that produces a time series over of the trial-average and an associated vector of weights representing the location of that component. Single-trial estimates of that ERP are then obtained by weight-summing over the vector of weights as described in **Chapter 3** and **Figure 3.2**.

Different methods for finding channel weights for single-trial estimation of N200 latencies were explored. The best method was quantified as the method that produced clear peaks at typical N200 latencies over all trials (and no where else) such that the peak latencies only differed by a range of about 50 ms across trials. The following methods were explored in depth: 1) calculating ERPs in each SNR condition and biasing single-trial EEG in each condition using an SVD component based on that conditions average, 2) calculating ERPs only in the low noise condition and biasing single-trial EEG in each condition using an SVD component based only on the low noise condition, or 3) calculating ERPs across all trials, disregarding condition, and biasing single-trial EEG by that SVD component. Method 1 was found to be the best calculation. An Independent Component Analysis (ICA) decomposition of ERP response matrix $T \times C$ was also explored. However independent components (ICs) did not produce consistent estimates of components across subjects. Thus the SVD was found to be a more stable and objective calculation such that the first component from each subject and condition was used.

4.2.5 Presence of fast errors in reaction time distributions

One possible confound in evaluating the relationship between N200 latencies and reaction time distributions is the presence of fast error reaction times. That is, fast reaction time observations that are not due to a decision making process, and occur due to another random process. Because these reaction times are not due to a decision process, the associated accuracy on those trials should be about 50% in a two-alternative forced choice task. These reaction times typically pollute the left tail of reaction time distributions as subjects are often able to become attentive to the task later if attention lapses at the beginning of the trial. Recent research has studied and developed theory on the notion of **mind wandering**, which is defined as cognition that departs from that which is useful to the task at hand (e.g. Hawkins et al., 2015; Lin et al., 2016).

In order to remove fast responses times that were not due to a decision process, an exponentially weighted moving average was used to calculate accuracy sorted by reaction time, from shortest reaction time to longest reaction time. A cutoff of 60% accuracy was used and the accuracy data and the associated short reaction times were removed from the data. However a fixed cutoff of 350 ms yielded no difference in the ultimate results, and therefore a fixed cutoff was maintained in the data. This procedure resulted in less than 7.6% of the response data being dropped for each visual noise condition across sessions of EEG, with the median being 0% removed in one visual noise condition of one session.

4.2.6 Simple analyses

A criticism of past analyses was that they depended upon complicated assumptions that exist within drift diffusion models (even though they are empirically validated, see Voss et al., 2004). Therefore it was apparent that simple analyses were needed in order to explore neural processing time and visual encoding time relationships, independently of cognitive

models. Simple linear regressions of both single-trial N200 latencies versus reaction times and traditional N200 latencies versus reaction time distribution statistics were performed.

Statistics reported are 95% confidence intervals related to the p values and t statistics in the Neyman and Pearson (1933) tradition as well as the natural logarithm of Bayes Factors which describe the amount of relative evidence (probability in the Bayesian definition) of a model which has a non-zero regression slope over a model which has a regression slope of zero (Kass and Raftery, 1995; Rouder and Morey, 2012). Adjusted R^2 is also reported that describes the fraction of variance of the dependent variable (RT statistics) explained by the regressor variable (measures of neural processing speed). R^2_{adj} ranges from 0 to 1. All statistics were generated by JASP, a open-source graphical software package for statistical analysis (JASP Team, 2017).

4.2.7 Integrated neurocognitive model fitting

As shown in previous research (Frank, 2015; Nunez et al., 2017) single-trial electrophysiological measures can be used to estimate single-trial non-decision times and decision parameters that are not identifiable on single-trials without experimental data from other modalities. This is performed using hierarchical Bayesian models with linear connectors between neural data and drift-diffusion model parameters. This technique can yield inference about the effect of single-trial electrocortical measures of attention on decision making parameters.

Parameter estimates from three hierarchical Bayesian models were found. These linear relationships were estimated in a single-step in a hierarchical Bayesian framework.

In Model 1, an exploratory analysis was performed where drift diffusion modeling was applied to reaction time and accuracy data from Experiment 1 that accounted for single-trial changes in non-decision time τ, within-trial evidence accumulation rate δ and within-trial evidence

accumulation standard deviation α that were linearly related to single-trial changes in single-trial N200 amplitudes, single-trial N200 latencies, single-trial P300 amplitudes, single-trial P300 latencies, and two steady-state visual evoked potential (30 Hz and 40 Hz stimulus frequency tagged EEG responses).

The relationship between single-trial estimates of neural processing speed as measured by N200 latencies responses and parameter estimates with cognitive interpretations was tested by assuming simple link functions between cognitive mechanisms in hierarchical models. That is, the diffusion model parameters on each trial were assumed to be equal to a simple linear combination of the vector of single-trial EEG inputs $\mathbf{x}_{ijk}$ on that trial i with η_{jk} and $\boldsymbol{\gamma}_{jk}$ as the intercept and slope parameters respectively:

$$
\begin{aligned}
\delta_{ijk} &= \eta_{(\delta)jk} + \boldsymbol{\gamma}_{(\delta)jk}^{\intercal}\mathbf{x}_{ijk} \\
\tau_{ijk} &= \eta_{(\tau)jk} + \boldsymbol{\gamma}_{(\tau)jk}^{\intercal}\mathbf{x}_{ijk} \\
\alpha_{ijk} &= \eta_{(\alpha)jk} + \boldsymbol{\gamma}_{(\alpha)jk}^{\intercal}\mathbf{x}_{ijk}
\end{aligned}
$$

Decision-diffusion modeling was also applied to reaction time and accuracy data from Experiments 1 & 2 jointly (data consisted of 12 unique subjects with 2 sessions of EEG each and 4 unique subjects with 7 sessions of EEG each respectively), containing between session-differences in non-decision time τ, within-trial evidence accumulation rate δ, and speed accuracy trade-off parameter α that were explained by session j and condition k differences in traditional N200 latencies z_{jk} (first negative peak latencies of ERPs; see Luck et al., 2000) of subject-level ERPs. All built models will were assumed to be hierarchical, describing intrinsic session j and condition k variability which 1) ensured model fits with small amounts of data (Lee, 2008; Lee and Newell, 2011; Vandekerckhove et al., 2011) and 2) provided the ability to predict data more easily for observed and unobserved subjects in future analyses (Wagenmakers, 2009; Nunez et al., 2017). Hierarchical parameters were also split by

experiment e. Bayesian hierarchical drift-diffusion models were fit using Wiener likelihood approximations in JAGS, a program that easily samples from marginal posterior distributions in hierarchical models using Markov Chain Monte Carlo (MCMC) (Plummer, 2003; Wabersich and Vandekerckhove, 2014).

In Models 2 and 3, wide (uninformative) priors were given the effects $\boldsymbol{\gamma}$ of EEG on the parameters of interest, centered at 1 in order to calculate Bayes Factors using the Savage-Dickey density ratio (see paragraph below; Verdinelli and Wasserman, 1995). In Model 3, wide (uninformative) priors were given to the modifier effects $\boldsymbol{\theta}$ of EEG on the parameters of interest, centered at 0 for the same reason (see paragraph below). Model 2 had the following prior and hyperprior structure:

$$
\begin{array}{llllll}
\tau_{jk} & \sim \mathcal{N}(\eta_{(\tau_{ek})} + \gamma_{1ek} z_{jk}, \sigma^2_{(\tau)}) \in (0,1) & , & \eta_{(\tau)k} & \sim \mathcal{N}(0.3, 0.25^2), & \sigma_{(\tau)} \sim \Gamma(0.2,1) \\
\delta_{jk} & \sim \mathcal{N}(\eta_{(\delta_{ek})} + \gamma_{2ek} z_{jk}, \sigma^2_{(\delta)}) \in (-9,9) & , & \eta_{(\delta_{ek})} & \sim \mathcal{N}(1, 2^2) & , \quad \sigma_{(\delta)} \sim \Gamma(1,1) \\
\alpha_{jk} & \sim \mathcal{N}(\eta_{(\alpha_{ek})} + \gamma_{3ek} z_{jk}, \sigma^2_{(\alpha)}) \in (0.1,3), & & \eta_{(\alpha_{ek})} & \sim \mathcal{N}(1, 0.5^2) & , \quad \sigma_{(\alpha)} \sim \Gamma(1,1) \\
z_{jk} & \sim \mathcal{N}(\mu_{(z_{ek})}, \sigma^2_{(z)}) & , & \mu_{(z_{ek})} & \sim \mathcal{N}(0.2, 0.1^2) & , \quad \sigma_{(z)} \sim \Gamma(0.2,1) \\
\gamma & \sim \mathcal{N}(\mu_{(\gamma)}, \sigma^2_{(\gamma)}) & , & \mu_{(\gamma)} & \sim \mathcal{N}(1, 3^2) & , \quad \sigma_{(\gamma)} \sim \Gamma(1,1)
\end{array}
$$

Model 3 was created in order to test the post-hoc hypothesis that noise condition and visual noise type (differing across the two experiments) modified the effect of N200 latency on non-decision time. In Models 3, the effect of trial-averaged N200 latency on non-decision time was assumed to be moderated by noise condition and experiment with modifier. Thus Model 3 had the following prior and hyperprior structure:

$$
\begin{array}{llllll}
\tau_{jk} & \sim \mathcal{N}(\eta_{(\tau_{ek})} + \gamma_{1ek} z_{jk} + \boldsymbol{\theta_1}^\intercal \mathbf{z}_{jk}, \sigma^2_{(\tau)}) \in (0,1) & , & \eta_{(\tau)k} & \sim \mathcal{N}(0.3, 0.25^2), & \sigma_{(\tau)} \sim \Gamma(0.2,1) \\
\delta_{jk} & \sim \mathcal{N}(\eta_{(\delta_{ek})} + \gamma_{2ek} z_{jk} + \boldsymbol{\theta_2}^\intercal \mathbf{z}_{jk}, \sigma^2_{(\delta)}) \in (-9,9) & , & \eta_{(\delta_{ek})} & \sim \mathcal{N}(1, 2^2) & , \quad \sigma_{(\delta)} \sim \Gamma(1,1) \\
\alpha_{jk} & \sim \mathcal{N}(\eta_{(\alpha_{ek})} + \gamma_{3ek} z_{jk} + \boldsymbol{\theta_3}^\intercal \mathbf{z}_{jk}, \sigma^2_{(\alpha)}) \in (0.1,3), & & \eta_{(\alpha_{ek})} & \sim \mathcal{N}(1, 0.5^2) & , \quad \sigma_{(\alpha)} \sim \Gamma(1,1) \\
z_{jk} & \sim \mathcal{N}(\mu_{(z_{ek})}, \sigma^2_{(z)}) & , & \mu_{(z_{ek})} & \sim \mathcal{N}(0.2, 0.1^2) & , \quad \sigma_{(z)} \sim \Gamma(0.2,1) \\
\gamma & \sim \mathcal{N}(1, 3^2), \quad \theta \sim \mathcal{N}(0, 1^2) & & & &
\end{array}
$$

In **Model 2** Bayes Factors for effect parameters in each experiment and condition were calculated with Python scripts using a Savage-Dickey density ratio (Verdinelli and Wasserman, 1995) of the posterior density over the prior distribution at $\gamma = 1$. These Bayes Factors provide the degree of evidence (defined as a probability ratio in Bayesian statistics) of a model where effect of trial-averaged N200 latency on non-decision time is 1-to-1 versus a model where the effect is unknown (the prior model where $\gamma \sim \mathcal{N}(1, 3^2)$). In **Model 3** Bayes Factors for *additional* effect parameters in each experiment and condition were calculated using a Savage-Dickey density ratio of the posterior density over the prior distribution at $\theta = 0$, i.e. the null model of no additional N200 effect with increasing visual noise contrast. These Bayes Factors provide the degree of evidence (defined as a probability ratio in Bayesian statistics) of a model where effect of trial-averaged N200 latency on non-decision time is 0 versus a model where the effect is unknown (the prior model where $\theta \sim \mathcal{N}(0, 3^2)$).

4.3 Results

4.3.1 Simple results: Scatter plots and linear regression

In order to explore the hypothesis that trial-averaged N200 latency predicts session-level visual encoding time, linear models between reaction time percentiles (both correct and incorrect reaction times combined) and N200 latencies were computed. If a drift diffusion model of quick decision making is to be believed, then the *shifts* of reaction time (RT) distributions (and thus the small RT percentiles) should correlate with non-decision time (i.e. encoding + motor response) while the shape of RT distributions (and thus the medium and larger RT percentiles) should correlate both with non-decision time *and* decision making

properties. Thus linear regressions can be explored without fitting diffusion models to data directly. A close to 1-to-1 relationship was found between in a simple linear regression between trial-averaged N200 latencies and 10th RT percentiles as shown in the left portion of **Figure 4.2** with a regression coefficient of $\beta_1 = 1.32$ ($N = 143$, 95% confidence interval: $[0.83, 1.81]$, $p < .001$, $t = 5.33$, $\log(\text{BF}) = 10.44$). Although it should be noted that only 16% of the variance in 10th RT percentiles were explained by the variance across trial-averaged N200 latencies ($R^2_{adj} = .162$).

The relationship is also expected to be observed in the reaction time distributions itself, albeit with more variance around the regression line. Single-trial estimates of N200 latencies were compared to reaction times across all data points. Again, a close to 1-to-1 relationship was found between in a simple linear regression between single-trial N200 latencies and reaction times as shown in the right portion of **Figure 4.2** with a regression coefficient of $\beta_1 = 0.95$ ($N = 13{,}868$, 95% confidence interval: $[0.84, 1.07]$, $p < .001$, $t = 15.85$, $\log(\text{BF}) = 120.42$). However only about 2% of the variance of raw single-trial reaction times were explained ($R^2_{adj} = .018$). This may be expected if other mechanisms besides visual encoding time contribute to the variance in reaction times, which was explored with the following cognitive modeling results.

4.3.2 Posterior distributions of Model 1

Each model was fit using JAGS with six Markov Chain Monte Carlo (MCMC) chains run in parallel of 52,000 samples each with 2,000 burn-in samples and a thinning parameter of either 100 (Model 1) or 10 (Models 2 and 3) resulting in either 500 (Model 1) or 5,000 (Models 2 and 3) posterior samples in each chain. The posterior samples from each chain were combined to form one posterior sample of 3,000 (Model 1) or 30,000 (Models 2 and 3) samples for each parameter. All three models converged as judged by $\hat{R}$ (see Gelman and

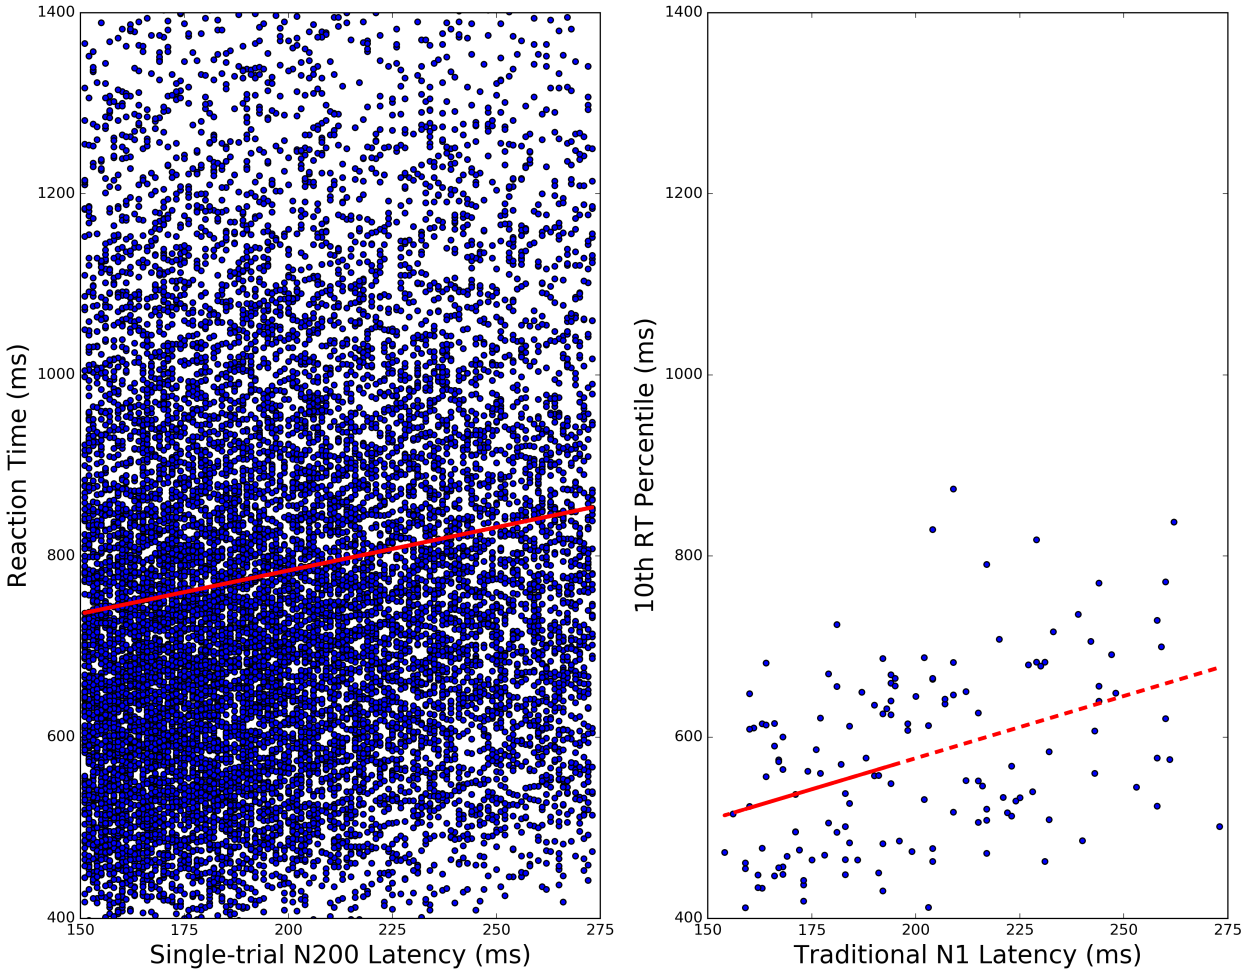

Figure 4.2: Left: A scatter plot of single-trial N200 latencies versus reaction times and the best fit linear regression line. Observations were generated per trial ($N = 15{,}680$) Right: A scatter plot of trial-averaged N200 latencies (i.e. traditional N1 latencies) and 10th reaction time percentiles. Observations were generated per noise condition and per EEG collection session ($N = 141$).

Rubin, 1992) being less than 1.02 for all parameters in each model.

Posterior distributions of model parameters for 7 subjects' single-trial N200 linear effects on single-trial non-decision time (the sum of encoding and motor response time as estimated by a hierarchical Bayesian account of a decision-diffusion model) support the hypothesis of a 1-to-1 correspondence (see **Figure 4.3**). However, this relationship between N200 latency and non-decision time was weaker and non-existent in medium and high noise conditions. In medium and high noise conditions, further processing may be required to estimate encoding time from single-trial N200 latency measures by introducing other sources of variance. In low noise condition, the evidence suggests that the EEG response to the onset of the stimulus reflects visual encoding time.

4.3.3 Posterior distributions of Models 2 and 3

Model 2 was used to evaluate the effect of trial-averaged N200 latency (i.e. "traditional" N1 latency) on non-decision time across EEG sessions and noise conditions. The mean effect across all conditions of each experiment μ_γ provides a sense of how N200 processing time tracks non-decision time in general. The posterior distribution of this parameter was found to be around 1 in each condition and experiment as shown in **Figure 4.4**, indicating that 1 ms increase in N200 processing corresponded to 1 ms increase in non-decision time. The posterior distribution of the overall hierarchical effect of N200 peak latency on non-decision time across sessions is given in **Figure 4.5**, indicating that the overall evidence for a 1-to-1 relationship between N1 latency and non-decision time is 5.86, indicating moderate evidence.

Model 3 was fit in order to evaluate the observed modifying effect of noise condition on effect of N200 peak latency (see increasing trend across conditions in **Figure 4.4**). Some evidence was found for the null model as indicated by the Bayes Factors for the null effect $BF0$ as shown in **Figure 4.6**. The base effect posterior distribution (evidence for the effect of N200

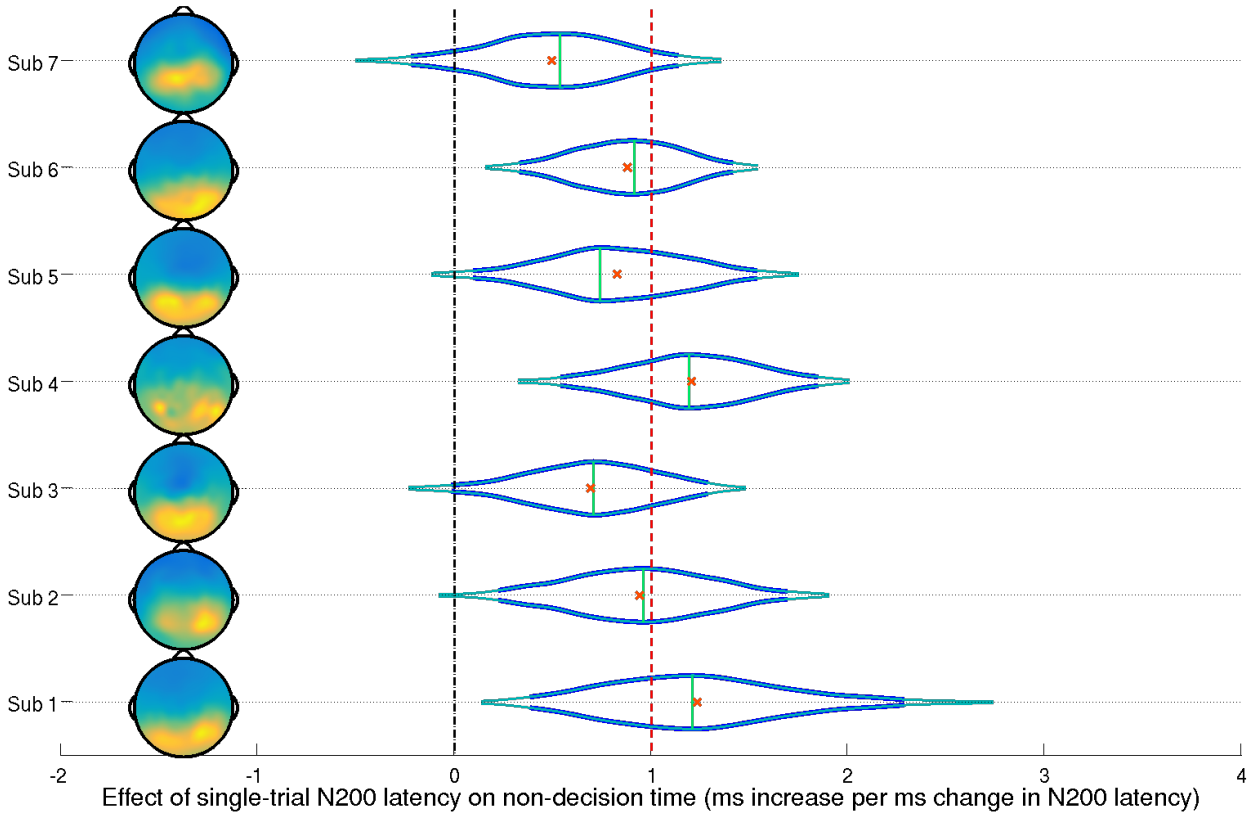

Figure 4.3: The posterior distributions of the effect of a trial's N200 latency (a visual processing component of the signal stimulus) on trial-specific non-decision times for each subject in a low noise condition. Thick lines forming the distribution functions represent 95% credible intervals while thin lines represent 99% credible intervals. Crosses and vertical lines represent posterior means and modes respectively. Also shown are the topographic representations of the channel weights of the first SVD component of each subject's ERP, indicating the location of single-trial N200s over occipital and parietal electrodes.

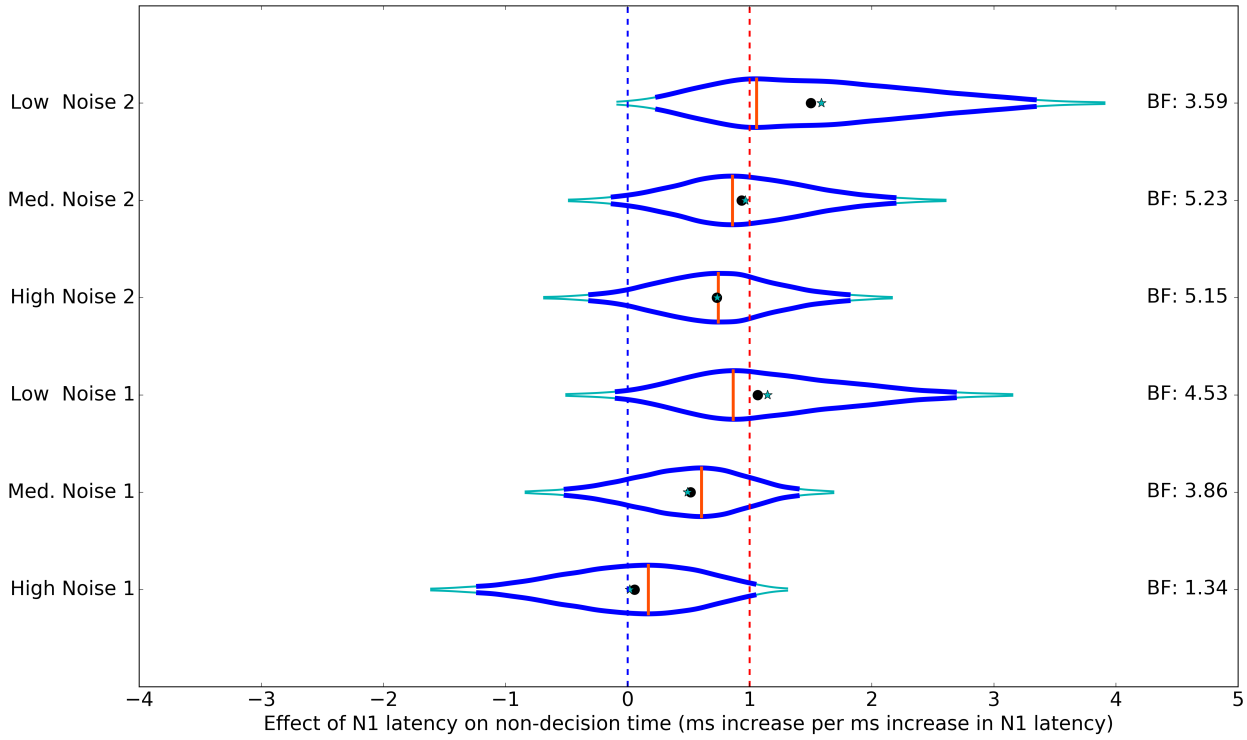

Figure 4.4: Session-level effects of N200 latency on non-decision time during the training experiment. Data was collected from 4 subjects over 7 training sessions each, resulting in 28 observations used in the linear model of trial-averaged N200 peak latency on non-decision time embedded in a hierarchical linear model. Some evidence exists for the effects of N200 peak latency on non-decision time to be 1-to-1 (one millisecond increase in N200 latency corresponds to a millisecond increase in non-decision time) as indicated by the Bayes Factors calculated with a Savage-Dickey density ratio (Verdinelli and Wasserman, 1995) of the posterior density over the prior distribution at $\gamma = 1$.

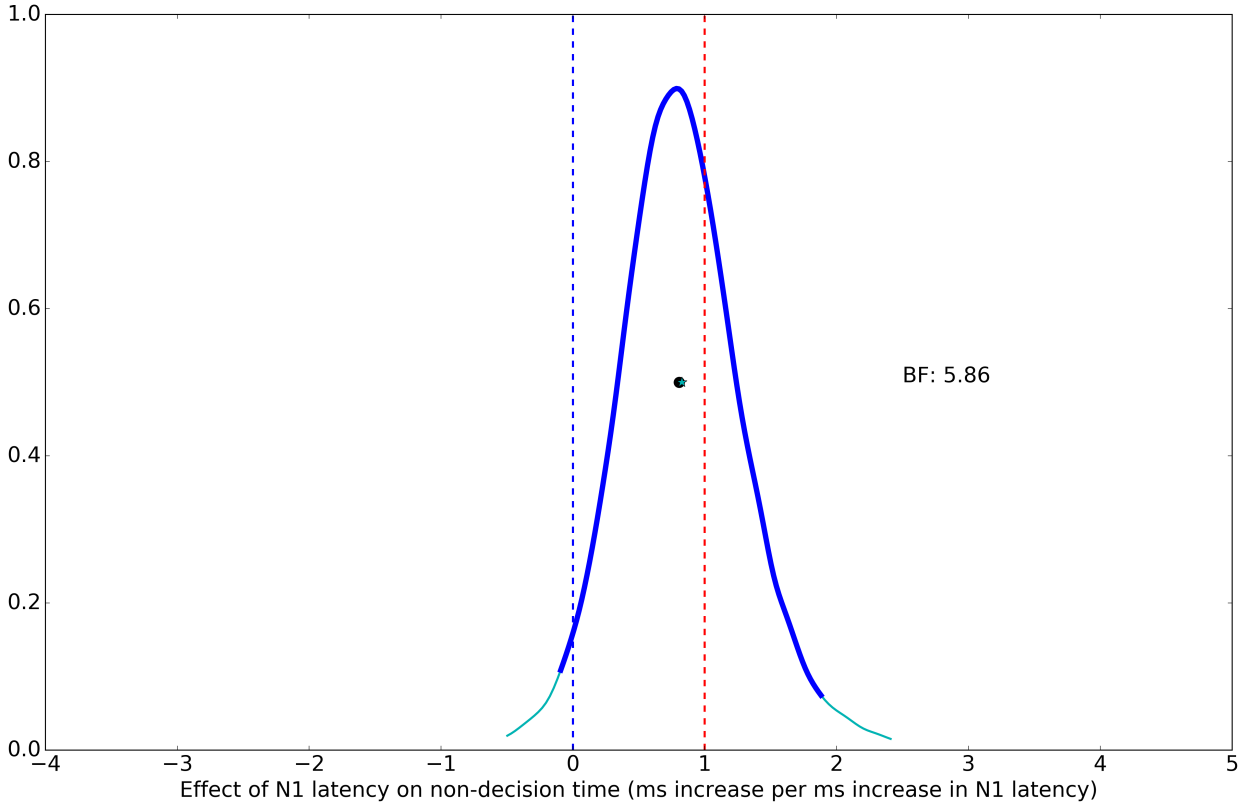

Figure 4.5: Overall hierarchical parameter of the effect of trial-averaged N200 latency on non-decision time during the training experiment. Evidence exists for the effects of N200 peak latency on non-decision time to be 1-to-1 (one millisecond increase in N200 latency corresponds to a millisecond increase in non-decision time) as indicated by the Bayes Factor of 5.86 of the posterior density over the prior distribution at $\mu_\gamma = 1$.

latency on non-decision time in Experiment 1 in the low noise condition) was close to that of the overall effect posterior in Model 2. Therefore not much evidence exists for moderator effects θ of condition and experiment.

4.4 Discussion

Evidence is presented both for and against the N200-visual encoding hypothesis. Evidence against is often left out of manuscripts presented in exploratory analyses in both the cognitive psychology and cognitive neuroscience fields, often presented as confirmatory results. Such "censoring" of data can lead to biases statistical analyses (Guan and Vandekerckhove, 2016). In this Chapter I present exploratory evidence both for and against this hypothesis.

4.4.1 Evidence for N200 latencies tracking visual encoding

Both simple regression analyses yielded evidence for event-related potential measures of neural processing speed reflecting visual encoding times due to the simple linear regression slope coefficients being close to 1. These regression fits indicate 1) that there was approximately a 1 ms increase in reaction times when there was a 1 ms increase in single-trial N200 latency and 2) that there was approximately a 1 ms increase in 10th percentile reaction times, a non-parametric estimate of non-decision times, when there was a 1 ms increase in trial-averaged N200 (i.e. traditional N1) latencies across both experiments and all experimental sessions and all visual noise conditions of EEG. Furthermore, moderate evidence was found of this 1-to-1 relationship of both single-trial N200 latencies and trial-average N200 latency on non-decision time in hierarchical drift-diffusion Models 1, 2, and 3.

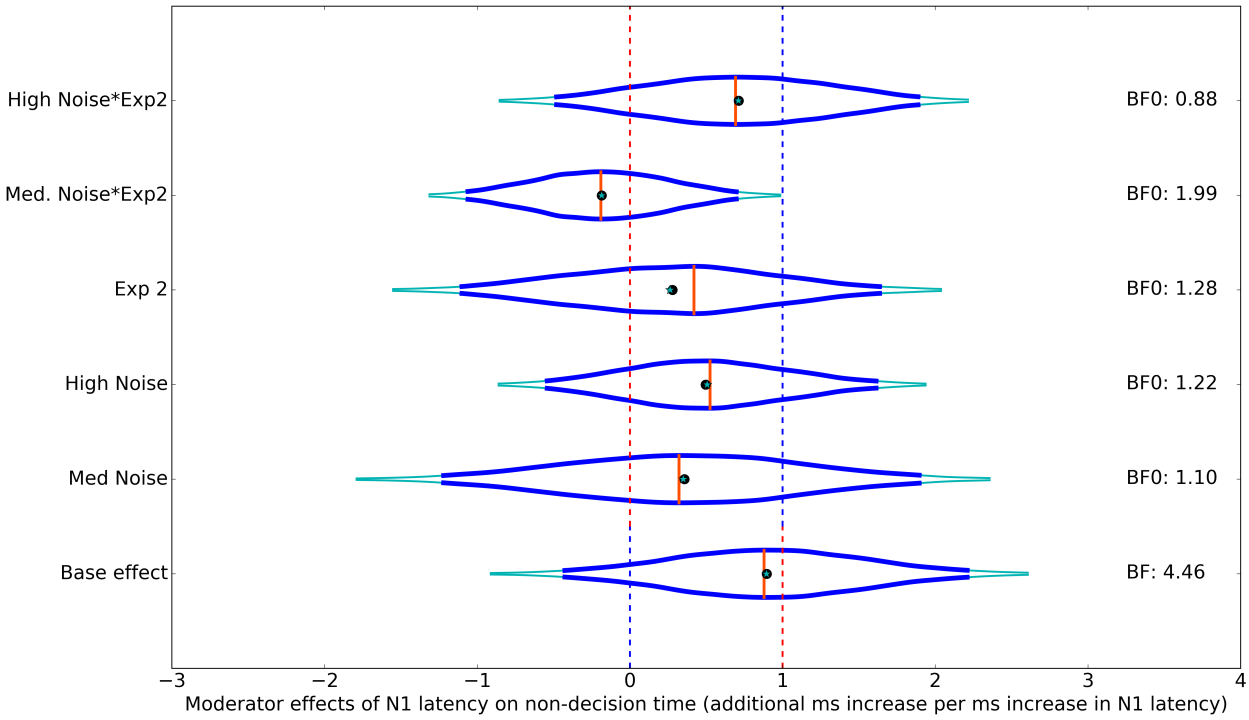

Figure 4.6: Session-level moderator effects of N200 latency on non-decision time during the training experiment. Data was collected from 4 subjects over 7 training sessions each, resulting in 28 observations used in the linear model of trial-averaged N200 peak latency on non-decision time embedded in a hierarchical linear model. Not much evidence exists for moderator effects of condition as indicated by the Bayes Factor for the null effect $BF0$ (that is evidence for $\theta = 0$ or no effect of condition on the N200-non-decision time relationship)

4.4.2 Evidence against N200 latencies tracking visual encoding

When exploring the relationship between cortical processing time (as measured by early ERP latencies) and subject-to-subject differences in non-decision time, the ability of ERP latencies to reflect visual encoding time is more clearly dependent upon the quality of visual stimuli. In two experiments with different noise types, the effects of session-level N1 latency differences on session-level non-decision time differences (assumed by the model to be the sum of visual encoding and residual motor response time) seemed to be mediated by the contrast condition of the visual distractors. This suggest that how well the peak neural signal latency tracks visual encoding time is dependent upon the quality of the external signal itself. However not much evidence was found for this qualitative finding as indicated by the Bayes Factors of the null model in **Figure 4.6**.

Single-trial N200 latencies may not be related to non-decision time in high noise conditions. However, further work is required to find better, testable estimates of single-trial neural processing speed as measured by ongoing EEG. It is also possible that the simple decision-diffusion model does not provide a satisfactory account of visual encoding. Or that the beginning of the negative deflection of the single-trial and subject-level EEG is a better indicator of encoding time. Initial evidence for this latter hypothesis is that early negative deflections better match visual encoding time as estimated by other modalities (e.g. see Schmolesky et al., 1998).

4.4.3 A new neurocognitive theory

A neurocognitive model of encoding, decision (via a diffusion process), and motor response is proposed that is identified with separable measures of encoding and motor response via EEG and decision processes that are estimated from observed decision and reaction time distributions. Fitting this type of model to data would allow inference about trial-to-trial

differences and individual differences in observed cortical signals as well as human behavior during quick decision making.

While some parameters of the model presented in **Equation 4.1** are classical unidentifiable with behavior alone, both unidentifiability issues can resolved by including electrophysiological measures as correlates of cognition directly as observed data in hierarchical models. In this Chapter I have presented initial evidence that non-decision time τ is related to the N200 latency almost 1-to-1 in response to stimuli. In the context of previous research (e.g. Martin et al., 2010; Loughnane et al., 2016; Zhang et al., 2016), this is further evidence that single-trial encoding times τ_e (i.e. visual processing time before the decision begins) are reflected in N200 latencies in the parietal cortex.

Neurocognitive hierarchical models of encoding, decision, and motor response time thus become identifiable when treating EEG measures (N200 latencies) as measures of encoding time. A quantitative neurocognitive theory of visual encoding, rapid decision-making, and motor preparation is therefore proposed that explains EEG measures of encoding, EEG measures of motor preparation, behavioral accuracy in two alternative forced choice (2AFC) tasks, and reaction time distributions. The theory is an extension of a class of Decision-Diffusion Models (e.g. "drift-diffusion" models, see Ratcliff and McKoon, 2008, for a review) which predicts that reaction time and accuracy are explained by a continuous accumulation of evidence towards certain pre-decided evidence thresholds (see **Figure 4.7**).

The base likelihood of all hierarchical evidence accumulation models will be assumed to be the same. A drift-diffusion model of accuracy ($w = 0$ or $w = 1$) and reaction time t on one trial is predicted to be a function of encoding time τ_e, motor response time τ_m, the drift rate – the average rate of evidence accumulation within one trial δ, the diffusion coefficient – the variance around the average rate of evidence accumulation ς, and the amount of evidence required to make a decision α_t, which may be assumed to decrease with time. The bias towards correct or incorrect responses is assumed to be equal to $z_0 = 0.5\alpha_t$. The joint

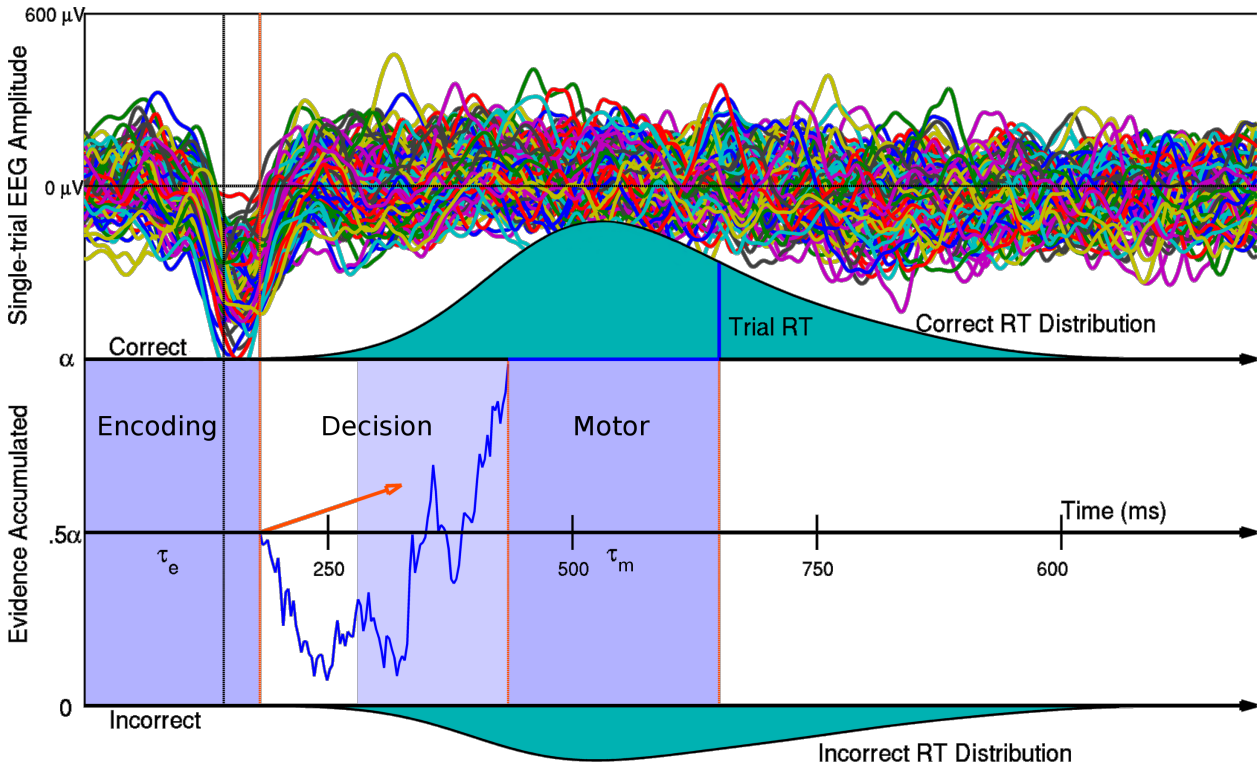

Figure 4.7: A graphical illustration of a Neural Decision Diffusion model in which the encoding time τ_e on single-trials describes the latency of the negative peaks of the EEG on 146 single-trials in occipital and parietal locations. Single-trial observations of the N200 latency are found by using a decomposition of the average ERP response at each electrode and then biasing the raw EEG by the resulting channel weights (this algorithm detailed in Chapter 3). Total non-decision time τ reflects both stimulus encoding time τ_e as well as residual motor response τ_m (i.e. motor preparation time after the decision is made) and can be estimated from reaction time distributions.

density f of RT t and accuracy w of this simplified diffusion model is given in **Equation 4.1**. The density is derived from the solution given by (Ratcliff, 1978; Tuerlinckx, 2004).

$$\begin{cases} f(t, w = 0 \mid \alpha, \varsigma^2, \tau_e, \tau_m, \delta) = \frac{\pi \varsigma^2}{\alpha_t^2} e^{-\frac{1}{2\varsigma^2}\left[\delta\alpha_t + \delta^2(t-\tau_e-\tau_m)\right]} \sum_{k=1}^{\infty} \left[k \sin\left(\frac{1}{2}\pi k\right) e^{-\frac{1}{2\alpha_t^2} k^2 \pi^2 \varsigma^2 (t-\tau_e-\tau_m)} \right] \\ \\ f(t, w = 1 \mid \alpha_t, \varsigma^2, \tau_e, \tau_m, \delta) = f(t, w = 0 \mid \alpha_t, \varsigma^2, \tau_e, \tau_m, -\delta) \end{cases} \tag{4.1}$$

Note that the density in **Equation 4.1** is classically unidentifiable when estimating the model with only reaction time and accuracy observations for two reasons. 1) Encoding time τ_e and motor response time τ_m both contribute to residual response time $\tau = \tau_e + \tau_m$ which approximately equals the smallest observed reaction time that is not due to fast guesses. 2) Only two of the three parameters related to evidence accumulation (i.e. drift rate δ, the diffusion coefficient ς, and the boundary separation α) can be found with behavioral data alone because, for example, if α_t is constant with t, multiplying ς by two and dividing both α and δ by two would result in the same fit of choice-RT (Wabersich and Vandekerckhove, 2014).

The goal of future work is to further develop theory that reconciles psychological phenomena with observations about the brain found in the cognitive neuroscience literature. This work would lead to the ability to fit the density given in **Equation 4.1** with *identified* parameters that explain both psychological phenomena and observed brain electrophysiology.

Chapter 5

Do EEG markers of motor processing reflect motor preparation time?

5.1 Introduction

Simple theories of speeded decision-making often assume that the cognitive processing required occurs in three sequential time periods: 1) a period of visual encoding, 2) a period of decision-making, and then 3) a period of motor response time (e.g. Ratcliff and McKoon, 2008). In a new preregistered study, I propose to test the common sequential assumption of cognitive processing during quick decision making by using EEG measures and hierarchical drift-diffusion models, analyses similar to those presented in the previous chapter. If evidence of sequential processing is found, the second goal is to find estimates of these three time periods in milliseconds for each subject using a mixture of human behavior observations (choice and reaction time) and evoked electroencephalographic (EEG) measures.

To further understanding of the time course of quick decision making, EEG measures of motor response time will be evaluated in relation to the predicted model of decision making.

The mu rhythm is a high alpha / beta band (often estimated at $\sim 9-15$ Hz) response in EEG over the motor cortex (central electrodes) that typically decreases in magnitude during motor preparation (Pfurtscheller et al., 2006). However the beta rhythm ($\sim 18-25$ Hz) can also be important markers of sensorimotor preparedness (McFarland et al., 2000). This neural motor measure has even become a reliable predictor in Brain-Computer Interfaces (e.g. McFarland et al., 2005). For this reason it is thought that the time course of "desynchronization" (decrease in EEG power) in the beta-band will be indicative of motor response time. This hypothesis will be evaluated using analysis of diffusion models with beta-band power decreases as measures of motor response time. Furthermore, it is hypothesized that non-decision time for each subject, as estimated by reaction time distributions, will total N200 negative peak latency time (an ERP possibly reflecting visual encoding time as explored in Chapter 4) and evoked motor response time as estimated by the peak negative amplitude of the beta band-desynchronization. The estimation of encoding time from the stimulus-locked EEG, the decision processes from the accuracy and response time distributions, and the motor response time from the mu-band desynchronization will refine knowledge of both individual differences and between experimental condition differences in quick decision making.

5.2 Methods

5.2.1 Preregistration

The confirmatory analysis of encoding time and motor response time separation was preregistered at the Open Science Framework (`https://osf.io`). Some data was collected before the preregistration was completed. However no data had been analyzed nor cleaned before the preregistration. Pilot data had been viewed and analyzed. However the pilot data will

not be used in future analysis. At the time of this writing, data was still being collected and therefore the final results are expected to be published in late 2017.

5.2.2 Subjects

Participants were recruited by word-of-mouth and email on the main campus of the University of California, Irvine. Participants must have been at least 18 years of age and have no prior history or family history of epilepsy. Participants were paid $60 for two sessions of EEG data collection during a visual decision making experiment, multiple trials of a simple game built using Psychtoolbox 3 in MATLAB. Evoked-EEG measures, reaction times, and choices were collected using MATLAB scripts.

5.2.3 Dataset

The target sample size is 32 subjects, each with 2 sessions of recordings. One session will yield 1) 480 trials of EEG data and behavior, 2) a 2 minute record of EEG while the subject is at rest with eyes closed, and 3) a 2 minute record of EEG while the subject is at rest with eyes open, fixated on a computer screen. We will recruit up to 50 participants, assuming that not all participants will complete two sessions.

The sample size for both trials and number of participants was constrained by time. We did not wish to keep participants in the experiment room for longer than 2 hours because participants would have become fatigued, resulting in poor experimental performance. Empirically it was found that 480 trials of data could be collected (80 trials per experimental cell) per experimental session, resulting in a total of 960 trials of EEG and behavioral observations per participant across sessions. We will terminate our data collection when we have both EEG and behavioral data from 32 participants with two complete sessions each.

5.2.4 Experimental tasks

Subjects were tasked to perform a two-alternative forced choice task, specifically to determine whether a Gabor stimulus contained high or low spatial frequency content. Stimuli were built and displayed using the MATLAB Psychophysics toolbox (Psychtoolbox-3; www.psychtoolbox.org). A graphical representation of one trial of the experiment is given in **Figure 5.1**.

Subjects in the two experiments described performed visual-based tasks in a dark room on a $\sim$ 61 cm LCD monitor with a distance of 57 cm from retina to display. The monitor resolution was set as $1920 * 1280$ pixels with a 120 Hz refresh rate. A distance of 57 cm was chosen to ensure easy calculation of visual angles, with x cm equaling $\sim x^{\circ}$ visual angle.

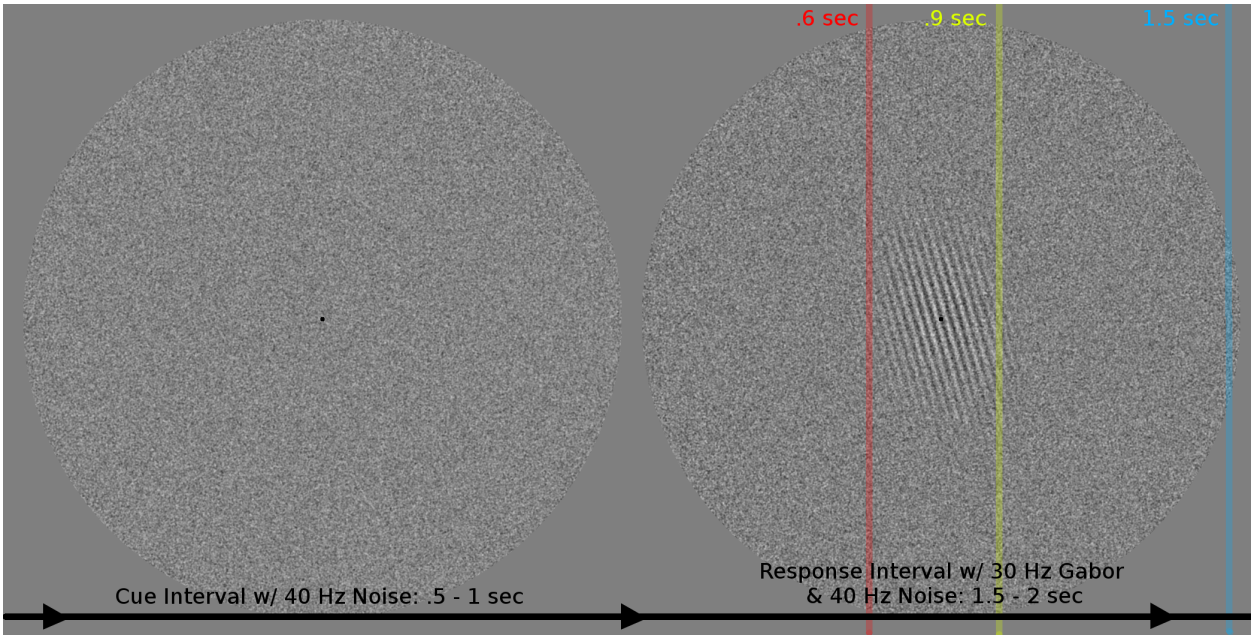

Figure 5.1: Example stimuli of the *cue* and *response* intervals of the experiment with three conditions of different times to respond (.6, .9, and 1.5 seconds). An SSVEP paradigm was be used in which the visual noise changed at 40 Hz and the Gabor signal flickered at 30 Hz to evoke 40 Hz and 30 Hz responses in electrocortical activity that track attention to the noise and signal stimuli respectively. For the first half of an experimental sessions, subjects had to respond with either only their right or left hands, and then use the other hand for the second half of the experiment.

The subjects were told to maintain fixation throughout the task on a .25 cm (.25°) fixation spot while attending to to the stimulus. The high and low spatial frequencies of the target Gabors were 2.4 and 2.6 cycles per degree visual angle (cpd) respectively. The visual noise was broadband noise. The experiment contained only one condition of contrast noise, found to make the experiment somewhat difficult for subjects.

This experiment had a within-subject design with 2 factors ("time boundary" and "hand of response"). The "time boundary" factor had thee levels: 1.5 seconds to respond, .9 seconds to respond, and .6 seconds to respond. The "hand of response" factor has two levels: left hand and right hand. Such that subjects responded with only their left hand or with only their right hand during particular trials (using their ring and index fingers for each choice button on a response-collection device). The trial types were blocked, therefore there were $2 * 3 = 6$ different types of blocks. Subjects started each session with either their right or left hands and then were displayed the blocks with either increasing or decreasing time-to-respond in either hand. Block order was counterbalanced across subjects. Subjects were thus assigned randomly to one of eight block orders (where each session of the experiment has 8 blocks). Subjects participated in two sessions of EEG data collection on separate days, less than a week apart. The block order for the second session reversed the "hand of response" order and reversed the block order of the "time boundary" within each hand condition.

The time course of each trial is as follows: subjects are asked to fixate on a fixation spot (.25° visual angle) in the center of the screen throughout the experiment, visual noise changing at 40 Hz is displayed for 500 ms to 1000 ms during the *cue* interval, then the Gabor signal stimulus is flickered at 30 Hz embedded in the noise stimuli for 1500 ms to 2000 ms during the *response* interval. During the response interval, subjects are asked to respond as accurate as possible in the time allowed during that block (either .6 seconds, .9 seconds, or 1.5 seconds). using a button box. Each session produced 6 blocks of 80 trials each for a total of 480 trials. Subjects were allowed to take breaks between blocks.

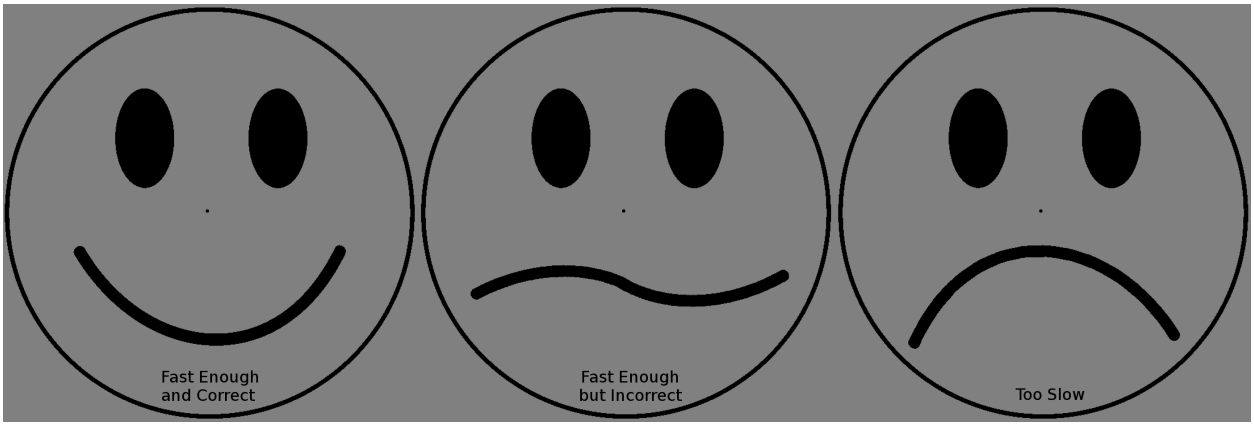

Figure 5.2: Example feedback images given to subjects during the experiment. Faces were displayed directly after completion of each trial. (Left) Happy faces were displayed when a subject answered within the allowed time period (either .6, .9, or 1.5 seconds into the response interval, depending upon the condition of that block of trials) and chose the correct spatial frequency. (Middle) "Unsure" faces were displayed when a subject answered within the allowed time period but chose the incorrect spatial frequency. (Right) Sad faces were displayed when a subject *did not* answer in the allowed time period for the block of trials. Sad faces were used to encourage subjects to answer within the time boundary.

For trial-to-trial feedback, cartoon face stimuli will be used, as humans are thought to use emotion and facial expressions as a mechanism of feedback and simple face stimuli have even been used in robotic learning experiments as proof-of-concept (Andry et al., 2011). "Unsure" faces were given for trials answered incorrectly but in the correct amount of time, happy faces were given for trials answered correctly in the correct amount of time, and sad faces were given for trials not answered in the correct amount of time for the block condition. Sad faces were given to encourage subjects to answer within the time boundary. However for all subsequent analyses, correct and incorrect trials will still be calculated based on the full response interval time period, so that the time boundary does not directly affect accuracy calculations. Example feedback stimuli are displayed in **Figure 5.2**.

5.2.5 EEG recording

Electroencephalograms (EEG) were recorded using Electrical Geodesics, Inc.'s high density 128-channel Geodesic Sensor Net and Advanced Neuro Technology's amplifier. Electrical activity from the scalp was recorded at a sampling rate of 1024 samples per second with OpenVibe recording software and Lab Streaming Layer to output the data for offline analysis using MATLAB and Python scripts.

5.2.6 Calculation of event related potentials

Following the results of Chapter 4, the ERP was windowed 150 to 275 ms post-stimulus in order to calculate trial-averaged N200 latency (the time of the minimum latency within the aforementioned window). However if N200 latencies were found with a minimum on the boundary, they were removed from the analysis. A minimum at either 150 or 275 ms was indicative that N200 latencies were not well estimated in that session. Singular value decomposition (SVD) was used to find single-trial latency estimates and reduced latency estimates of trial-averaged N200 latencies (see Nunez et al., 2017, and Methods of Chapter 4)

To find single-value estimates of motor response time, beta desynchronization was calculated by using the Fast Fourier Transform (FFT) function in MATLAB with a moving Gaussian window over the time course of the response-locked data. Then left and right central electrodes (C3 and C4 in the 10-20 electrode layout, see Jasper, 1958) were used to find beta band (14-18 Hz) desynchronization time courses for left-hand and right-hand responses (found using the contralateral electrode on the opposite hemisphere to the used hand in that block). The peak negative magnitude of the beta band desynchronization before response was used as an estimate of motor preparation time, to be compared to reaction time distributions and non-decision time estimates.

5.2.7 Data exclusion

In order to remove "fast error" trials that do not contain decision-making processes, reaction times less than 350 ms will be removed. Trial records deemed to contain too much EEG *artifact* by visual inspection (high amplitude or high frequency electrical activity that is impossible for cortically-generated activity, see Nunez et al., 2016) will be removed. Single-trial N200 latencies that are found to be on the boundary (either 150 ms or 275 ms post-stimulus; indicating a failure of the simple windowing algorithm to find a stimulus-related peak) will be removed. Trial-averaged N200 latencies that are found to be on the same boundaries will also be removed for the same reasons. Beta desynchronization peaks that are found to be on the -600 ms or -200 ms window boundaries will be removed for the same reasons. Test data will be removed from each block of the data such that 25% of trials will be marked as test data for use in out-of-sample posterior predictive distribution evaluations in exploratory analyses.

5.2.8 Simple analyses

In the same analysis as Chapter 4, simple linear regressions of both single-trial N200 latencies versus reaction times and trial-averaged N200 latencies versus reaction time 10th percentiles will be performed. The trial-averaged N200 latencies observations will be calculated per session and condition of the experiment (and therefore per block of the experiment, each block being a particular hand and response-cutoff condition). The response-locked peak beta desynchronization times will also be used as regressors in a simple linear model with 10th reaction time percentiles as the dependent variable. In addition, a linear regression will be performed with two regressors. The two independent variables will be trial-averaged N200 latencies and beta-desynchronization peak latency times, with reaction time 10th percentiles as the dependent variable.

In the simple linear regression analyses, statistics reported will be 95% confidence intervals of the effect parameters as well as the natural logarithm of Bayes Factors which describe the amount of relative evidence (probability in the Bayesian definition) of a model which has a non-zero regression slope over a model which has a regression slope of zero. Adjusted R^2 will also be reported that describes the fraction of variance of the dependent variable (RT statistics) explained by the regressor variables (N200 latencies and peak beta desynchronization times). All statistics will be generated by JASP (JASP Team, 2017), a open-source graphical software package for statistical analysis.

5.2.9 Integrated neurocognitive model fitting

Both Models 2 and 3 from Chapter 4 will be repeated, with two observations of trial-averaged N200 latency per hand condition and three observations per response-cutoff condition. If the Bayes Factors using the Savage-Dickey method indicate some evidence that there is a 1-to-1 linear effect of N200 latency on non-decision in the first hierarchical Bayesian model ($BF > 3$ for the hierarchical effect parameter), then an alternate hierarchical Bayesian model will be used to test whether there were additive effects of N200 latency for the right hand and longer time boundaries. If evidence is found for both simple hypotheses (evidence of both a 1-to-1 effect of N200 latency and 1-to-1 effect of peak beta desynchronization time on non-decision time) then an alternate hierarchical with both EEG regressors on non-decision time simultaneously will be fit.

The posterior distributions of the linear effect parameters in the hierarchical Bayesian models were evaluated using Bayes Factors calculated using the Savage-Dickey ratio to 1) indicate evidence of a 1-to-1 linear effect of N200 latency on non-decision time, 2) indicate evidence of a 1-to-1 linear effect of peak beta desynchronization time on non-decision time, and 3) indicate evidence of the null effect (a slope parameter of zero indicating no effect) of additive

effect parameters in the follow-up analysis.

5.3 Future directions

5.3.1 EEG markers of decision-making

In **Section 4.4.3** I describe an extended theory of decision making that is not identifiable with behavioral data alone. However the three evidence accumulation variables may be separable with the addition of macro-scale cortical behavior. For instance, the diffusion coefficient ς describes the variance in the relative evidence accumulated between two choices on every trial. Assuming neural population representations of evidence (Shadlen and Kiani, 2013), the representation of evidence is expected to vary within a trial at every time step based on the amount of variance in the neural population response due to oscillatory local field potentials (Mitchell et al., 2007, 2009). Because EEG is thought to be generated from synchronous slow-wave extracellular potentials (Buzsaki, 2006; Nunez and Srinivasan, 2006) and alpha power (8 to 12 Hz) decreases are a well known phenomenon that predicts accuracy (Ergenoglu et al., 2004), the reduction of humans' resting-state alpha rhythm over parietal and occipital cortex may reflect a mechanism of internal noise reduction in the neural representation of the evidence accumulation process. It is therefore predicted that trial-to-trial alpha power influences trial-to-trial evidence accumulation variances in the form of internal noise reduction. This hypothesis will be evaluated by first using 1) single-trial and subject-level connectors between alpha power and parameter estimates of the diffusion coefficient ς in hierarchical models and 2) in-sample and out-of-sample prediction for model validation. If enough evidence for this noise suppression mechanism is found, then neural representations of estimates of ς will be evaluated in models with the all three evidence accumulation parameters (e.g. α, ς, and δ) estimated.

The neural response to visual face feedback on the previous trial is expected to affect only the evidence boundary (i.e. affecting the subject's speed-accuracy tradeoff) and not other evidence parameters, so the effects of neural response to visual face feedback will be explored. Slow-wave motor EEG potentials and their relationship to decision-time will also be explored in hierarchical models.

5.3.2 In support of open science

Data and code will be released online in support of the greater scientific community and in an effort to promote "open science" (Eich, 2014). Specifically, raw EEG data from a few example subjects and sessions, behavioral data and preprocessed EEG measures for all subjects and sessions, generated visual stimuli written in MATLAB with Psychtoolbox 3, simple MATLAB or Python scripts to perform computations of any newly discovered analysis techniques, and MATLAB or Python code needed to generate results presented in published work will be released online.

Bibliography

Anderson, J. R., Zhang, Q., Borst, J. P., and Walsh, M. M. (2016). The discovery of processing stages: Extension of sternbergs method. *Psychological Review*, 123(5):481.

Andry, P., Blanchard, A., and Gaussier, P. (2011). Using the rhythm of nonverbal human–robot interaction as a signal for learning. *IEEE Transactions on IEEE Transactions on Autonomous Mental Development*, 3(1):30–42.

Bridwell, D. A., Hecker, E. A., Serences, J. T., and Srinivasan, R. (2013). Individual differences in attention strategies during detection, fine discrimination, and coarse discrimination. *Journal of Neurophysiology*, 110(3):784–794.

Bridwell, D. A. and Srinivasan, R. (2012). Distinct attention networks for feature enhancement and suppression in vision. *Psychological Science*, 23(10):1151–1158.

Brown, S. D. and Heathcote, A. (2008). The simplest complete model of choice response time: linear ballistic accumulation. *Cognitive psychology*, 57(3):153–178.

Brown, S. D., Ratcliff, R., and Smith, P. L. (2006). Evaluating methods for approximating stochastic differential equations. *Journal of Mathematical Psychology*, 50(4):402–410.

Buschman, T. J. and Miller, E. K. (2007). Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science*, 315(5820):1860–1862.

Buzsaki, G. (2006). *Rhythms of the Brain*. Oxford University Press. 02625.

Calhoun, V., Adali, T., Pearlson, G., and Pekar, J. (2001). A method for making group inferences from functional MRI data using independent component analysis. *Human Brain Mapping*, 14(3):140–151.

Callaway, E. and Halliday, R. (1982). The effect of attentional effort on visual evoked potential n1 latency. *Psychiatry Research*, 7(3):299–308.

Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M. A., Guo, J., Li, P., and Riddell, A. (2016). Stan: A probabilistic programming language. *Journal of Statistical Software*, 20:1–37.

Cassey, P., Gaut, G., Steyvers, M., and Brown, S. (2014). A simple accumulator model that simultaneously accounts for behaviour and spike-train data. Paper presented at the 55th Annual Meeting of the Psychonomic Society, Long Beach, CA.

Cavanagh, J. F., Wiecki, T. V., Cohen, M. X., Figueroa, C. M., Samanta, J., Sherman, S. J., and Frank, M. J. (2011). Subthalamic nucleus stimulation reverses mediofrontal influence over decision threshold. *Nature Neuroscience*, 14(11):1462–1467.

Corbetta, M. and Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3(3):201–215.

Davis, E. T. and Graham, N. (1981). Spatial frequency uncertainty effects in the detection of sinusoidal gratings. *Vision Research*, 21(5):705–712.

Delorme, A., Sejnowski, T., and Makeig, S. (2007). Enhanced detection of artifacts in EEG data using higher-order statistics and independent component analysis. *Neuroimage*, 34(4):1443–1449.

Deng, S., Winter, W., Thorpe, S., and Srinivasan, R. (2012). Improved surface laplacian estimates of cortical potential using realistic models of head geometry. *IEEE Transactions on Biomedical Engineering*, 59(11):2979–2985.

Desimone, R. and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18(1):193–222.

Ding, J., Sperling, G., and Srinivasan, R. (2006). Attentional modulation of SSVEP power depends on the network tagged by the flicker frequency. *Cerebral Cortex*, 16(7):1016–1029.

Dosher, B. A. and Lu, Z.-L. (2000a). Mechanisms of perceptual attention in precuing of location. *Vision Research*, 40(10):1269–1292.

Dosher, B. A. and Lu, Z.-L. (2000b). Noise exclusion in spatial attention. *Psychological Science*, 11(2):139–146.

Dutilh, G., Vandekerckhove, J., Forstmann, B. U., Keuleers, E., Brysbaert, M., and Wagenmakers, E.-J. (2012). Testing theories of post-error slowing. *Attention, Perception, & Psychophysics*, 74(2):454–465.

Eich, E. (2014). Business not as usual. *Psychological Science*, 25(1):3–6.

Ergenoglu, T., Demiralp, T., Bayraktaroglu, Z., Ergen, M., Beydagi, H., and Uresin, Y. (2004). Alpha rhythm of the EEG modulates visual detection performance in humans. *Cognitive Brain Research*, 20(3):376–383.

Eriksen, C. W. and Hoffman, J. E. (1972). Temporal and spatial characteristics of selective encoding from visual displays. *Perception & Psychophysics*, 12(2):201–204.

Etz, A. and Vandekerckhove, J. (2017). Introduction to bayesian inference for psychology. *Psychonomic Bulletin & Review*, pages 1–30.

Fischl, B., van der Kouwe, A., Destrieux, C., Halgren, E., Ségonne, F., Salat, D. H., Busa, E., Seidman, L. J., Goldstein, J., Kennedy, D., et al. (2004). Automatically parcellating the human cerebral cortex. *Cerebral Cortex*, 14(1):11–22.

Forstmann, B. U., Anwander, A., Schafer, A., Neumann, J., Brown, S., Wagenmakers, E.-J., Bogacz, R., and Turner, R. (2010). Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. *Proceedings of the National Academy of Sciences*, 107(36):15916–15920.

Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., Von Cramon, D. Y., Ridderinkhof, K. R., and Wagenmakers, E.-J. (2008). Striatum and pre-sma facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences*, 105(45):17538–17542.

Frank, M. J. (2015). Linking across levels of computation in model-based cognitive neuroscience. In Forstmann, B. U. and Wagenmakers, E.-J., editors, *An Introduction to Model-Based Cognitive Neuroscience*, pages 159–177. Springer New York.

Frank, M. J., Gagne, C., Nyhus, E., Masters, S., Wiecki, T. V., Cavanagh, J. F., and Badre, D. (2015). fMRI and EEG predictors of dynamic decision parameters during human reinforcement learning. *Journal of Neuroscience*, 35(2):485–494.

Garcia, J. O., Srinivasan, R., and Serences, J. T. (2013). Near-real-time feature-selective modulations in human cortex. *Current Biology*, 23(6):515–522.

Gelman, A. and Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, pages 457–472.

Gold, J. I. and Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30:535–574.

Guan, M. and Vandekerckhove, J. (2016). A bayesian approach to mitigation of publication bias. *Psychonomic Bulletin & Review*, 23(1):74–86.

Harter, M. R. and Aine, C. J. (1984). Brain mechanisms of visual selective attention. *Varieties of Attention*, pages 293–321.

Hawkins, G., Mittner, M., Boekel, W., Heathcote, A., and Forstmann, B. (2015). Toward a model-based cognitive neuroscience of mind wandering. *Neuroscience*, 310:290–305.

Hawkins, G. E., Mittner, M., Forstmann, B. U., and Heathcote, A. (2017). On the efficiency of neurally-informed cognitive models to identify latent cognitive states. *Journal of Mathematical Psychology*, 76:142–155.

Herz, D. M., Tan, H., Brittain, J.-S., Fischer, P., Cheeran, B., Green, A. L., FitzGerald, J., Aziz, T. Z., Ashkan, K., Little, S., et al. (2017). Distinct mechanisms mediate speed-accuracy adjustments in cortico-subthalamic networks. *Elife*, 6.

Ho, T. C., Brown, S., and Serences, J. T. (2009). Domain general mechanisms of perceptual decision making in human cortex. *The Journal of Neuroscience*, 29(27):8675–8687.

Hong, B., Guo, F., Liu, T., Gao, X., and Gao, S. (2009). N200-speller using motion-onset visual response. *Clinical Neurophysiology*, 120(9):1658–1666.

JASP Team, T. (2017). JASP (Version 0.8.1.2)[Computer software].

Jasper, H. H. (1958). Report of the committee on methods of clinical examination in electroencephalography. *Electroencephalography and Clinical Neurophysiology*, 10(2):370 – 375.

Jolliffe, I. T. (1982). A note on the use of principal components in regression. *Applied Statistics*, pages 300–303.

Jung, T.-P., Makeig, S., Humphries, C., Lee, T.-W., Mckeown, M. J., Iragui, V., and Sejnowski, T. J. (2000). Removing electroencephalographic artifacts by blind source separation. *Psychophysiology*, 37(02):163–178.

Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795.

Krishnan, L., Kang, A., Sperling, G., and Srinivasan, R. (2013). Neural strategies for selective attention distinguish fast-action video game players. *Brain Topography*, 26(1):83–97.

Lee, M. D. (2008). Three case studies in the Bayesian analysis of cognitive models. *Psychonomic Bulletin & Review*, 15(1):1–15.

Lee, M. D. (2011). How cognitive modeling can benefit from hierarchical bayesian models. *Journal of Mathematical Psychology*, 55(1):1–7.

Lee, M. D. and Newell, B. J. (2011). Using hierarchical Bayesian methods to examine the tools of decision-making. *Judgment & Decision Making*, 6(8).

Lee, M. D. and Wagenmakers, E.-J. (2014). *Bayesian cognitive modeling: A practical course*. Cambridge University Press.

Lee, T.-W., Girolami, M., and Sejnowski, T. J. (1999). Independent component analysis using an extended infomax algorithm for mixed subgaussian and supergaussian sources. *Neural Computation*, 11(2):417–441.

Lin, C.-T., Chuang, C.-H., Kerick, S., Mullen, T., Jung, T.-P., Ko, L.-W., Chen, S.-A., King, J.-T., and McDowell, K. (2016). Mind-wandering tends to occur under low perceptual demands during driving. *Scientific Reports*, 6.

Link, S. and Heath, R. (1975). A sequential theory of psychological discrimination. *Psychometrika*, 40(1):77–105.

Loughnane, G. M., Newman, D. P., Bellgrove, M. A., Lalor, E. C., Kelly, S. P., and OConnell, R. G. (2016). Target selection signals influence perceptual decisions by modulating the onset and rate of evidence accumulation. *Current Biology*, 26(4):496–502.

Lu, Z.-L. and Dosher, B. A. (1998). External noise distinguishes attention mechanisms. *Vision Research*, 38(9):1183–1198.

Luck, S. J., Woodman, G. F., and Vogel, E. K. (2000). Event-related potential studies of attention. *Trends in Cognitive Sciences*, 4(11):432–440.

Makeig, S., Bell, A. J., Jung, T.-P., and Sejnowski, T. J. (1996). Independent component analysis of electroencephalographic data. *Advances in Neural Information Processing Systems*, 8:145–151.

Martin, T., Huxlin, K. R., and Kavcic, V. (2010). Motion-onset visual evoked potentials predict performance during a global direction discrimination task. *Neuropsychologia*, 48(12):3563–3572.

McFarland, D. J., Miner, L. A., Vaughan, T. M., and Wolpaw, J. R. (2000). Mu and beta rhythm topographies during motor imagery and actual movements. *Brain Topography*, 12(3):177–186.

McFarland, D. J., Sarnacki, W. A., Vaughan, T. M., and Wolpaw, J. R. (2005). Brain-computer interface (BCI) operation: signal and noise during early training sessions. *Clinical Neurophysiology*, 116(1):56–62.

Meister, M. L., Hennig, J. A., and Huk, A. C. (2013). Signal multiplexing and single-neuron computations in lateral intraparietal area during decision-making. *The Journal of Neuroscience*, 33(6):2254–2267.

Mitchell, J. F., Sundberg, K. A., and Reynolds, J. H. (2007). Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron*, 55(1):131–141.

Mitchell, J. F., Sundberg, K. A., and Reynolds, J. H. (2009). Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron*, 63(6):879–888.

Morgan, S., Hansen, J., and Hillyard, S. (1996). Selective attention to stimulus location modulates the steady-state visual evoked potential. *Proceedings of the National Academy of Sciences*, 93(10):4770–4774.

Mulder, M., van Maanen, L., and Forstmann, B. (2014). Perceptual decision neurosciences a model-based review. *Neuroscience*, 277:872–884.

Müller, M. M., Picton, T. W., Valdes-Sosa, P., Riera, J., Teder-Sälejärvi, W. A., and Hillyard, S. A. (1998). Effects of spatial selective attention on the steady-state visual evoked potential in the 20–28 hz range. *Cognitive Brain Research*, 6(4):249–261.

Neyman, J. and Pearson, E. S. (1933). On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231:289–337.

Nunez, M. D., Nunez, P. L., and Srinivasan, R. (2016). Electroencephalography (EEG): neurophysics, experimental methods, and signal processing. In Ombao, H., Linquist, M., Thompson, W., and Aston, J., editors, *Handbook of Neuroimaging Data Analysis*, pages 175–197. Chapman & Hall/CRC.

Nunez, M. D., Srinivasan, R., and Vandekerckhove, J. (2015). Individual differences in attention influence perceptual decision making. *Frontiers in Psychology*, 6:18.

Nunez, M. D., Vandekerckhove, J., and Srinivasan, R. (2017). How attention influences perceptual decision making: Single-trial EEG correlates of drift-diffusion model parameters. *Journal of Mathematical Psychology*, 76, Part B:117–130.

Nunez, P., Silberstein, R., Cadusch, P., Wijesinghe, R., Westdorp, A., and Srinivasan, R. (1994). A theoretical and experimental study of high resolution EEG based on surface laplacians and cortical imaging. *Electroencephalography and Clinical Neurophysiology*, 90(1):40–57.

Nunez, P. L. and Pilgreen, K. L. (1991). The spline-laplacian in clinical neurophysiology: a method to improve EEG spatial resolution. *Journal of Clinical Neurophysiology*, 8(4):397–413.

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

O'Connell, R. G., Dockree, P. M., and Kelly, S. P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nature Neuroscience*, 15(12):1729–1735.

Olitsky, S., Lee, H., and Young, E. (2013). IVAC—interactive visual acuity chart. `http://www.smbs.buffalo.edu/oph/ped/IVAC/IVAC.html`. Accessed: 2013-12-13.

Palmeri, T., Love, B., and Turner, B. (2017). Model-based cognitive neuroscience. *Journal of Mathematical Psychology*, 76, Part B:59 – 64.

Palmeri, T. J., Schall, J. D., and Logan, G. D. (2015). *Neurocognitive modeling of perceptual decision making*. Oxford handbook of computational and mathematical psychology.

Parra, L. C., Spence, C. D., Gerson, A. D., and Sajda, P. (2005). Recipes for the linear analysis of EEG. *Neuroimage*, 28(2):326–341.

Pfurtscheller, G., Brunner, C., Schlögl, A., and Da Silva, F. L. (2006). Mu rhythm (de) synchronization and EEG single-trial classification of different motor imagery tasks. *Neuroimage*, 31(1):153–159.

Philiastides, M. G., Heekeren, H. R., and Sajda, P. (2014). Human scalp potentials reflect a mixture of decision-related signals during perceptual choices. *The Journal of Neuroscience*, 34(50):16877–16889.

Philiastides, M. G., Ratcliff, R., and Sajda, P. (2006). Neural representation of task difficulty and decision making during perceptual categorization: a timing diagram. *The Journal of Neuroscience*, 26(35):8965–8975.

Philiastides, M. G. and Sajda, P. (2007). EEG-informed fMRI reveals spatiotemporal characteristics of perceptual decision making. *Journal of Neuroscience*, 27(48):13082–13091.

Pisauro, A., Fouragnan, E., Retzler, C., and Philiastides, M. (2017). Neural correlates of evidence accumulation during value-based decisions revealed via simultaneous EEG-fMRI. *Nature Communications*.

Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd International Workshop on Distributed Statistical Computing (DSC 2003), Vienna, Austria.*

Pommier, J. and Renard, Y. (2005). Getfem++, an open source generic c++ library for finite element methods.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2):59.

Ratcliff, R. and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, 20(4):873–922.

Ratcliff, R., Philiastides, M. G., and 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 Sciences*, 106(16):6539–6544.

Ratcliff, R., Thapar, A., and McKoon, G. (2001). The effects of aging on reaction time in a signal detection task. *Psychology and Aging*, 16(2):323.

Regan, D. (1977). Steady-state evoked potentials. *JOSA*, 67(11):1475–1489.

Rodriguez, C. A., Turner, B. M., Van Zandt, T., and McClure, S. M. (2015). The neural basis of value accumulation in intertemporal choice. *European Journal of Neuroscience*, 42(5):2179–2189.

Rosenblum, M. G., Pikovsky, A. S., and Kurths, J. (1996). Phase synchronization of chaotic oscillators. *Physical Review Letters*, 76(11):1804.

Ross, S. M. (2014). *Introduction to probability models*. Academic press.

Rouder, J. N., Lu, J., Speckman, P., Sun, D., and Jiang, Y. (2005). A hierarchical model for estimating response time distributions. *Psychonomic Bulletin & Review*, 12(2):195–223.

Rouder, J. N. and Morey, R. D. (2012). Default bayes factors for model selection in regression. *Multivariate Behavioral Research*, 47(6):877–903.

Salvatier, J., Wiecki, T. V., and Fonnesbeck, C. (2016). Probabilistic programming in python using PyMC3. *PeerJ Computer Science*, 2:e55.

Sazonov, A. V., Ho, C. K., Bergmans, J. W., Arends, J. B., Griep, P. A., Verbitskiy, E. A., Cluitmans, P. J., and Boon, P. A. (2009). An investigation of the phase locking index for measuring of interdependency of cortical source signals recorded in the EEG. *Biological cybernetics*, 100(2):129–146.

Schmolesky, M. T., Wang, Y., Hanes, D. P., Thompson, K. G., Leutgeb, S., Schall, J. D., and Leventhal, A. G. (1998). Signal timing across the macaque visual system. *Journal of neurophysiology*, 79(6):3272–3278.

Schubert, A.-L., Hagemann, D., and Frischkorn, G. (2017). Is intelligence little more than the speed of higher-order processing? *Journal of Experimental Psychology: General*.

Shadlen, M. N. and Kiani, R. (2013). Decision making as a window on cognition. 80(3):791–806.

Shadlen, M. N. and Newsome, W. T. (1996). Motion perception: seeing and deciding. *Proceedings of the National Academy of Sciences*, 93(2):628–633.

Shadlen, M. N. and Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *Journal of Neurophysiology*, 86(4):1916–1936.

Shaw, M. L. and Shaw, P. (1977). Optimal allocation of cognitive resources to spatial locations. *Journal of Experimental Psychology: Human Perception and Performance*, 3(2):201.

Smith, P. L. and Ratcliff, R. (2009). An integrated theory of attention and decision making in visual signal detection. *Psychological Review*, 116(2):283–317.

Spieler, D. H., Balota, D. A., and Faust, M. E. (2000). Levels of selective attention revealed through analyses of response time distributions. *Journal of Experimental Psychology: Human Perception and Performance*, 26(2):506.

Tallon-Baudry, C., Bertrand, O., Delpuech, C., and Pernier, J. (1996). Stimulus specificity of phase-locked and non-phase-locked 40 hz visual responses in human. *The Journal of Neuroscience*, 16(13):4240–4249.

Tange, O. (2011). GNU parallel - the command-line power tool. *The USENIX Magazine*, 36(1):42–47.

Thorpe, S., Fize, D., and Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, 381(6582):520.

Tuerlinckx, F. (2004). The efficient computation of the cumulative distribution and probability density functions in the diffusion model. *Behavior Research Methods, Instruments, & Computers*, 36(4):702–716.

Tuerlinckx, F., Maris, E., Ratcliff, R., and De Boeck, P. (2001). A comparison of four methods for simulating the diffusion process. *Behavior Research Methods, Instruments, & Computers*, 33:443–456.

Turner, B. M., Forstmann, B. U., Love, B. C., Palmeri, T. J., and Van Maanen, L. (2017). Approaches to analysis in model-based cognitive neuroscience. *Journal of Mathematical Psychology*, 76:65–79.

Turner, B. M., Forstmann, B. U., Wagenmakers, E.-J., Brown, S. D., Sederberg, P. B., and Steyvers, M. (2013). A bayesian framework for simultaneously modeling neural and behavioral data. *Neuroimage*, 72:193–206.

Turner, B. M., Rodriguez, C. A., Norcia, T. M., McClure, S. M., and Steyvers, M. (2016). Why more is better: Simultaneous modeling of EEG, fMRI, and behavioral data. *Neuroimage*, 128:96–115.

Turner, B. M., van Maanen, L., and Forstmann, B. U. (2015). Informing cognitive abstractions through neuroimaging: The neural drift diffusion model. *Psychological Review*, 122(2):312–336.

Twomey, D. M., Murphy, P. R., Kelly, S. P., and O'Connell, R. G. (2015). The classic p300 encodes a build-to-threshold decision variable. *European Journal of Neuroscience*.

van Maanen, L., Brown, S. D., Eichele, T., Wagenmakers, E.-J., Ho, T., Serences, J., and Forstmann, B. U. (2011). Neural correlates of trial-to-trial fluctuations in response caution. *The Journal of Neuroscience*, 31(48):17488–17495.

van Ravenzwaaij, D., Provost, A., and Brown, S. D. (2017). A confirmatory approach for integrating neural and behavioral data into a single model. *Journal of Mathematical Psychology*, 76:131–141.

van Vugt, M. K., Simen, P., Nystrom, L. E., Holmes, P., and Cohen, J. D. (2012). EEG oscillations reveal neural correlates of evidence accumulation. *Decision Neuroscience*, 6:106.

Vandekerckhove, J. (2014). A cognitive latent variable model for the simultaneous analysis of behavioral and personality data. *Journal of Mathematical Psychology*, 60:58–71.

Vandekerckhove, J. and Tuerlinckx, F. (2007). Fitting the ratcliff diffusion model to experimental data. *Psychonomic Bulletin & Review*, 14(6):1011–1026.

Vandekerckhove, J. and Tuerlinckx, F. (2008). Diffusion model analysis with MATLAB: A DMAT primer. *Behavior Research Methods*, 40(1):61–72.

Vandekerckhove, J., Tuerlinckx, F., and Lee, M. D. (2011). Hierarchical diffusion models for two-choice response times. *Psychological Methods*, 16(1):44.

Verdinelli, I. and Wasserman, L. (1995). Computing Bayes factors using a generalization of the savage-dickey density ratio. *Journal of the American Statistical Association*, 90(430):614–618. 00422.

Voss, A., Rothermund, K., and Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory & Cognition*, 32(7):1206–1220.

Wabersich, D. and Vandekerckhove, J. (2014). Extending JAGS: A tutorial on adding custom distributions to JAGS (with a diffusion model example). *Behavior Research Methods*, 46(1):15–28.

Wagenmakers, E.-J. (2009). Methodological and empirical developments for the ratcliff diffusion model of response times and accuracy. *European Journal of Cognitive Psychology*, 21(5):641–671.

Wagenmakers, E.-J., Ratcliff, R., Gomez, P., and McKoon, G. (2008). A diffusion model account of criterion shifts in the lexical decision task. *Journal of Memory and Language*, 58(1):140–159.

Webster, M. A., De Valois, R. L., et al. (1985). Relationship between spatial-frequency and orientation tuning of striate-cortex cells. *JOSA A*, 2(7):1124–1132.

White, C. N., Congdon, E., Mumford, J. A., Karlsgodt, K. H., Sabb, F. W., Freimer, N. B., London, E. D., Cannon, T. D., Bilder, R. M., and Poldrack, R. A. (2014). Decomposing decision components in the stop-signal task: A model-based approach to individual differences in inhibitory control. *Journal of Cognitive Neuroscience*, 26(8):1601–1614.

Wiecki, T. V., Sofer, I., and Frank, M. J. (2013). HDDM: Hierarchical bayesian estimation of the drift-diffusion model in python. *Frontiers in Neuroinformatics*, 7.

Zhang, Q., Walsh, M. M., and Anderson, J. R. (2016). The effects of probe similarity on retrieval and comparison processes in associative recognition. *Journal of Cognitive Neuroscience*.

Appendix A

Supplementary materials Chapter 2

A.1 Supplementary Equations

A.1.1 Predictive power as measured by R^2_{pred}

We define R^2_{pred} as a measure of the percentage of total between-subject variance of a statistic T (e.g. the correct-RT median) explained by in-sample or out-of-sample prediction. It is a function of the mean squared error of prediction (MSEP) and the sample variance of the statistic T based on a sample size of $n = 17$ subjects. R^2_{pred} is defined in Equation A.1.

$$R^2_{\text{pred}} = 1 - \frac{\sum_{j=1}^{17}(T_j - T_{(\text{pred})j})^2/16}{\sum_{j=1}^{17}(T_j - \bar{T})^2/16} = 1 - \frac{\text{MSEP}_T}{\widehat{\text{Var}[T]}} \tag{A.1}$$

A.1.2 Supplementary Figures and Tables

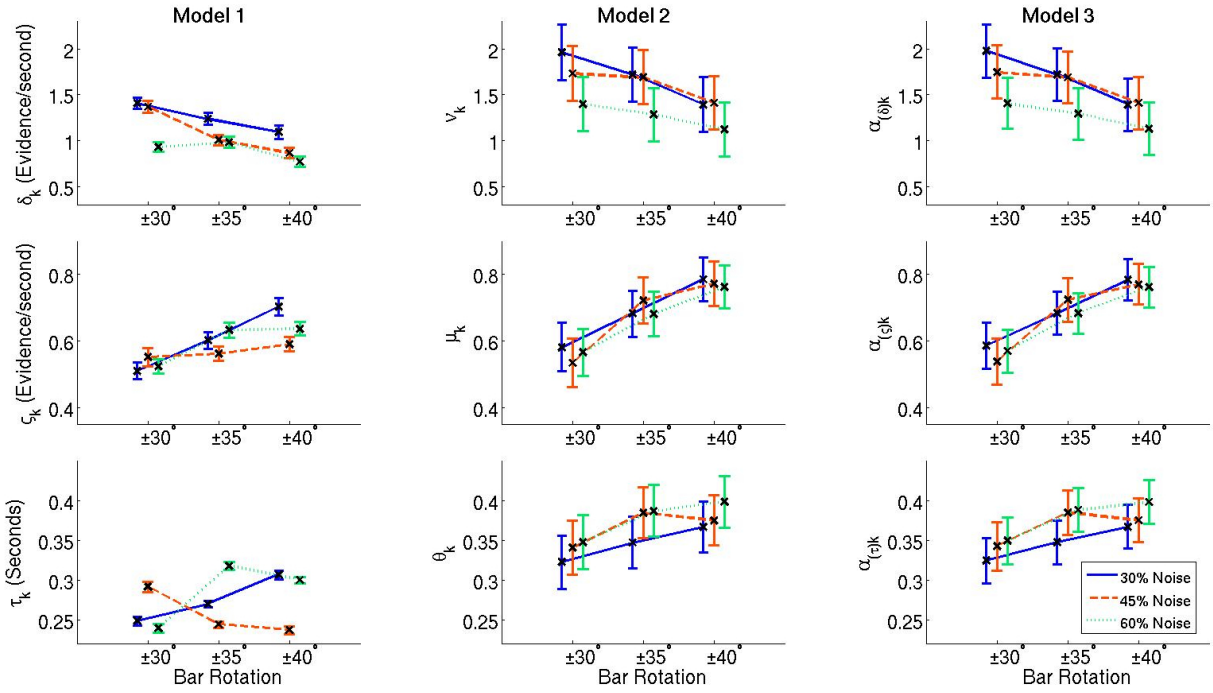

Figure A.1: The posterior means as represented by the black crosses and 95% credible intervals of the condition level parameter posterior distributions. From top to bottom, the rows correspond to the condition level drift rate, diffusion coefficient, and non-decision time posterior distributions for each model. The left column displays posterior samples of the parameters in **Model 1**, the model without assumed individual differences. The middle column shows condition level posterior distributions from **Model 2**. The larger variances in the posterior samples are due to explicitly modeled individual differences. The right column shows the credible intervals of the condition effects on each variable of **Model 3**. As expected the model with exogenous neural data as predictors of the subject level parameters has no effect on the condition level parameters.

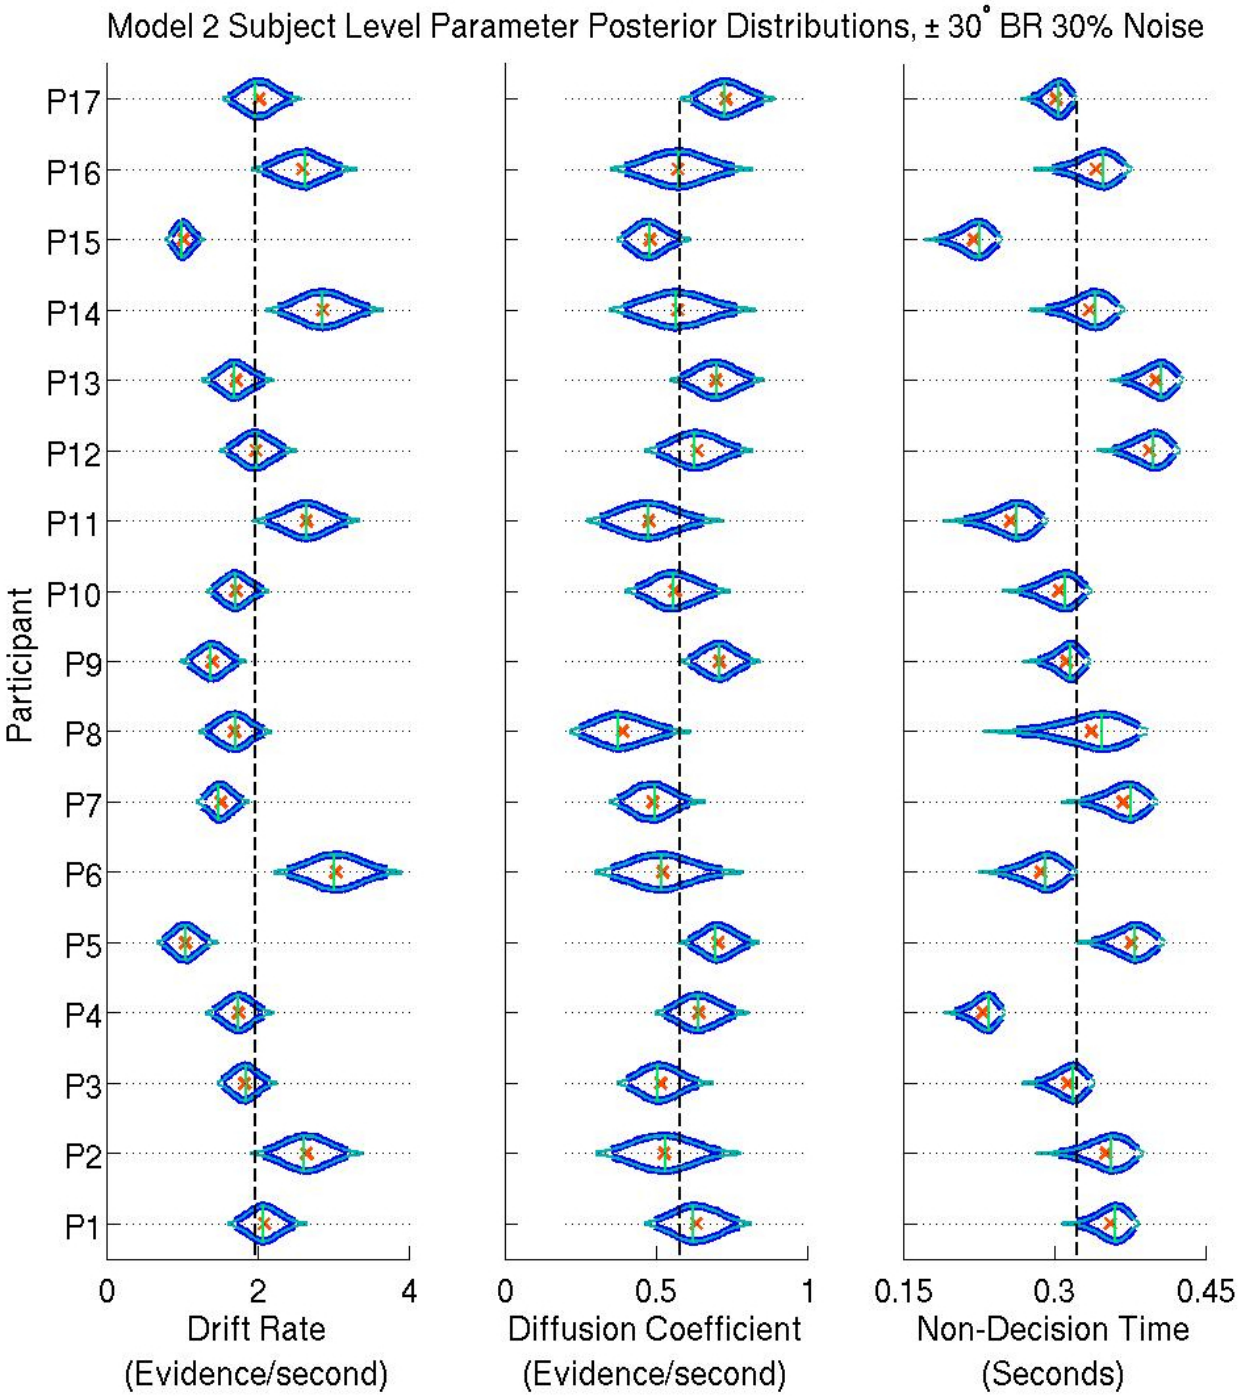

Figure A.2: Posterior distributions for the subject level parameters of the ±30° bar rotation and 30% noise condition. The thick blue line of the posterior density indicates 95% coverage, the density between the 2.5th and 97.5th percentiles. The thin teal line of the posterior density indicates 99% coverage, the density between the .5th and 99.5th percentiles. The vertical black line is the mean value of the condition level parameters. Nine subject level drift rates δ_{j1} deviate significantly at the 95% level from the mean posterior condition level drift rate ν_k. Three subject level diffusion coefficients ς_{j1} deviate significantly at the 95% level from the mean posterior condition level diffusion coefficient μ_k. And nine subject level non-decision times τ_{j1} deviate significantly at the 95% level from the mean posterior condition level non-decision time θ_k. It is clear from these results that there were differences between participants' cognition in the easiest condition.

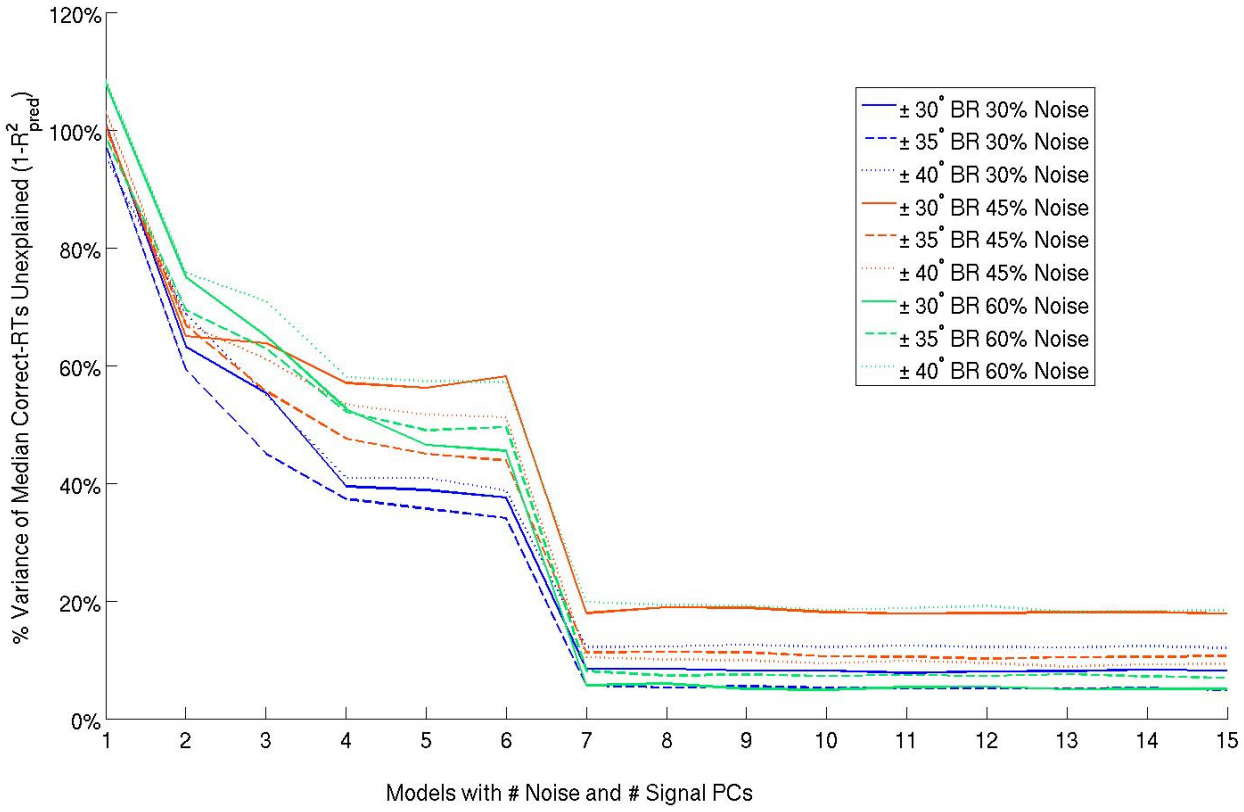

Figure A.3: Two principal components analyses (PCAs) on the noise and signal PLI variables were performed to obtain **Model 3**'s regressors. We generated in-sample posterior predictive distributions using samples from condition level parameter posterior distributions and EEG variables (to find subject level parameter predictive distributions) in order to find principal components that best predicted correct-RT distributions. Predictive power was measured as R^2_{pred} of the subject correct-RT medians of each condition. Each model type is plotted versus $1 - R^2_{\text{pred}}$ which is a measure of the percent variance of a statistic (e.g. median correct-RT) unexplained by *in-sample* prediction. We iteratively added one principal component per variable set (signal variables and noise variables) to the model. Based on this analysis principal components 2, 4, and 7 of both the noise and signal sets were tested further to find the model that best predicted *out-of-sample* reaction time of correct responses.

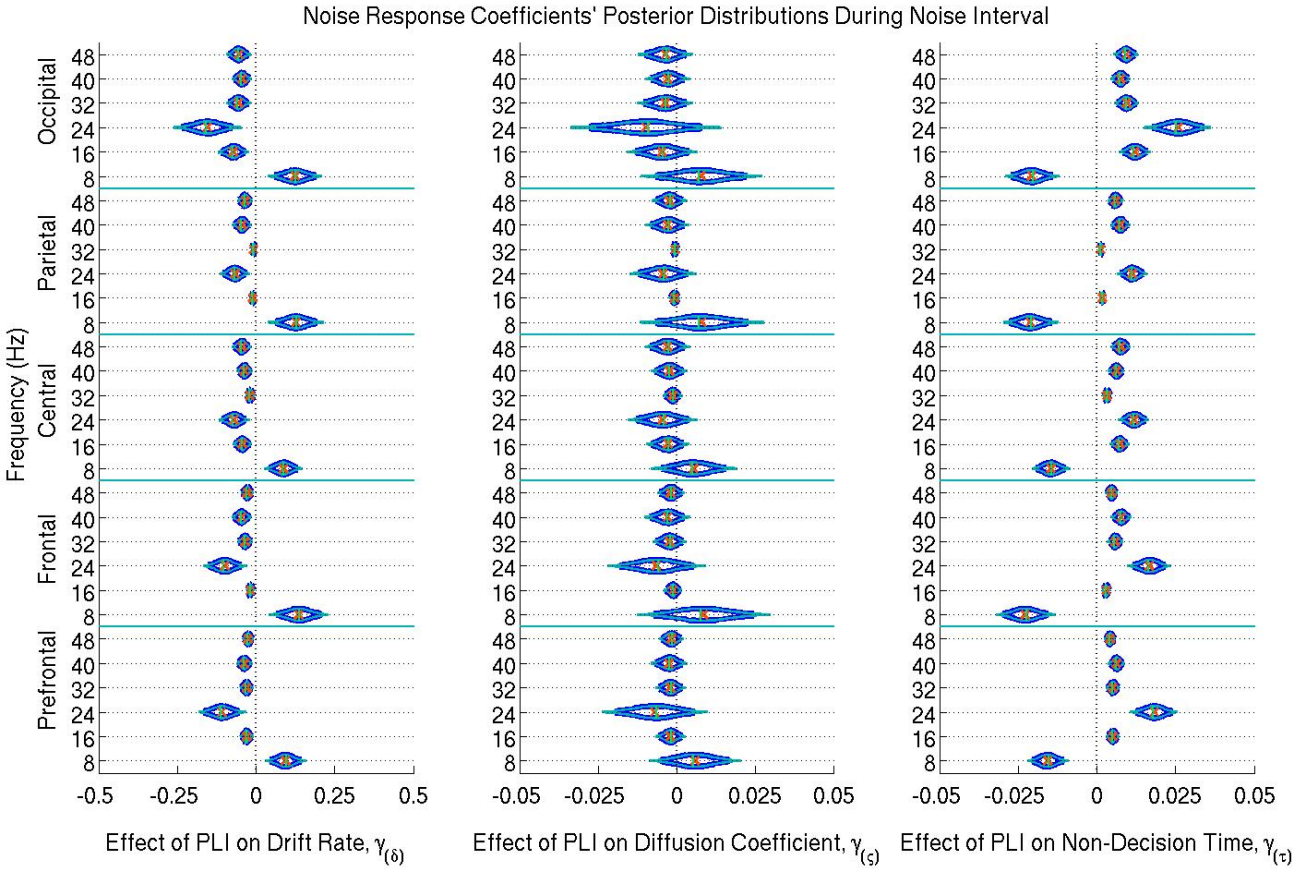

Figure A.4: The marginal posterior distributions of the noise PLI coefficients on the drift rate, diffusion coefficient, and non-decision time during the noise interval. Dark blue lines indicate 95% credible intervals, smaller teal lines indicate 99% credible intervals, horizontal green lines indicate posterior medians, and the orange exes indicate posterior means. We consider noise PLI coefficients to be significant if the 95% credible intervals do not include 0. At all noise harmonic frequencies (16, 24, 32, 40 and 48 Hz) during the noise interval, those subjects who suppressed noise had faster evidence accumulation rates (drift rates); this effect was found at all electrode groups. However noise enhancement at 8 Hz was associated with slower evidence accumulation. Those subjects who better suppressed noise at harmonic frequencies also had faster non-decision times. E.g., a participant whose PLI responses were suppressed .2 units more than another participant's responses at all locations and frequencies during the noise interval is expected to accumulate 0.131 evidence units per second faster and have 22 ms faster non-decision times, leading to faster and more correct responses.

Correct-RT Mean	**Out-of-sample Prediction**		
Condition	M1	M2	M3
±30° BR 30% Noise	−25.2%	−23.2%	−4.0%
±35° BR 30% Noise	−38.1%	−21.8%	6.5%
±40° BR 30% Noise	−26.3%	−17.1%	18.9%
±30° BR 45% Noise	−21.4%	−23.0%	−0.8%
±35° BR 45% Noise	−43.1%	−18.1%	16.0%
±40° BR 45% Noise	−73.8%	−17.0%	25.3%
±30° BR 60% Noise	−47.5%	−28.6%	−3.3%
±35° BR 60% Noise	−44.0%	−29.4%	−11.7%
±40° BR 60% Noise	−48.8%	−13.6%	36.8%

Table A.1: Percentage of between-subject variance in correct-RT means explained by out-of-sample prediction (R^2_{pred}) for each experimental condition.

Correct-RT 25th Percentile	**Out-of-sample Prediction**		
Condition	M1	M2	M3
±30° BR 30% Noise	−50.2%	−24.0%	23.6%
±35° BR 30% Noise	−70.6%	−31.5%	13.0%
±40° BR 30% Noise	−51.8%	−26.5%	15.8%
±30° BR 45% Noise	−47.4%	−25.4%	19.2%
±35° BR 45% Noise	−97.6%	−30.9%	16.6%
±40° BR 45% Noise	−120.3%	−43.5%	−0.1%
±30° BR 60% Noise	−110.2%	−49.0%	−5.3%
±35° BR 60% Noise	−70.8%	−35.6%	7.3%
±40° BR 60% Noise	−79.2%	−45.7%	1.7%

Table A.2: Percentage of between-subject variance in correct-RT 25th percentiles explained by out-of-sample prediction (R^2_{pred}) for each experimental condition.

Correct-RT 75th Percentile	**Out-of-sample Prediction**		
Condition	M1	M2	M3
±30° BR 30% Noise	−62.5%	−41.3%	−15.7%
±35° BR 30% Noise	−61.3%	−21.8%	0.4%
±40° BR 30% Noise	−50.0%	−22.9%	9.9%
±30° BR 45% Noise	−39.1%	−27.3%	−20.7%
±35° BR 45% Noise	−88.2%	−18.4%	12.4%
±40° BR 45% Noise	−193.2%	−24.0%	13.2%
±30° BR 60% Noise	−86.3%	−24.8%	−6.5%
±35° BR 60% Noise	−82.4%	−33.6%	−19.5%
±40° BR 60% Noise	−71.6%	−11.8%	27.8%

Table A.3: Percentage of between-subject variance in correct-RT 75th percentiles explained by out-of-sample prediction (R^2_{pred}) for each experimental condition.

Appendix B

Supplementary materials Chapter 3

B.0.1 Posterior predictive plots

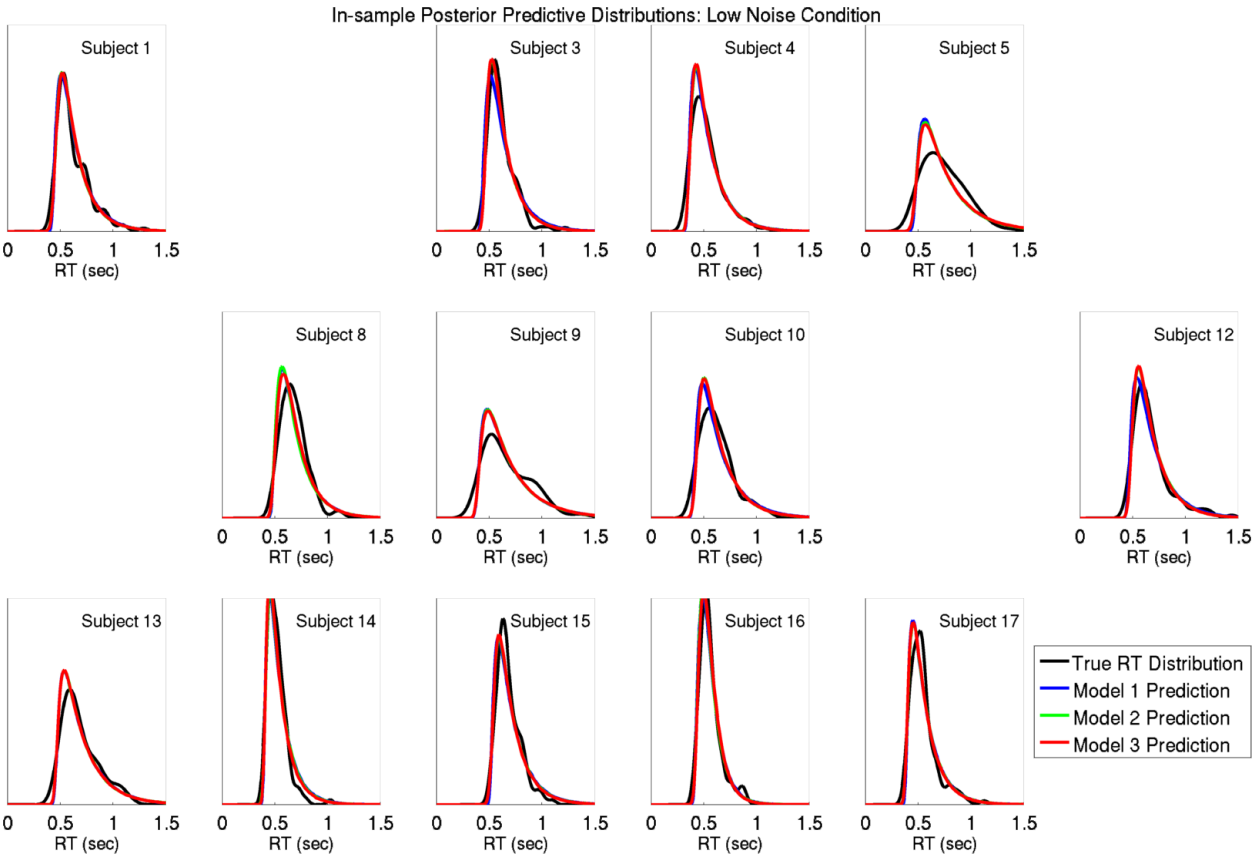

Figure B.1: For each subject, posterior predictive distributions of correct-RTs during trials from the low noise condition were compared to true correct-RT distributions from training (i.e. "in-sample") data. Note that subjects 2, 6, 7, and 11 were randomly chosen to be left out of the training set in order to test the prediction ability of each model for unknown subjects. Each model performs well at predicting training correct-RT data from known subjects in the low noise condition. More comprehensive evaluations of in-sample prediction are provided in **Table 3** in the paper.

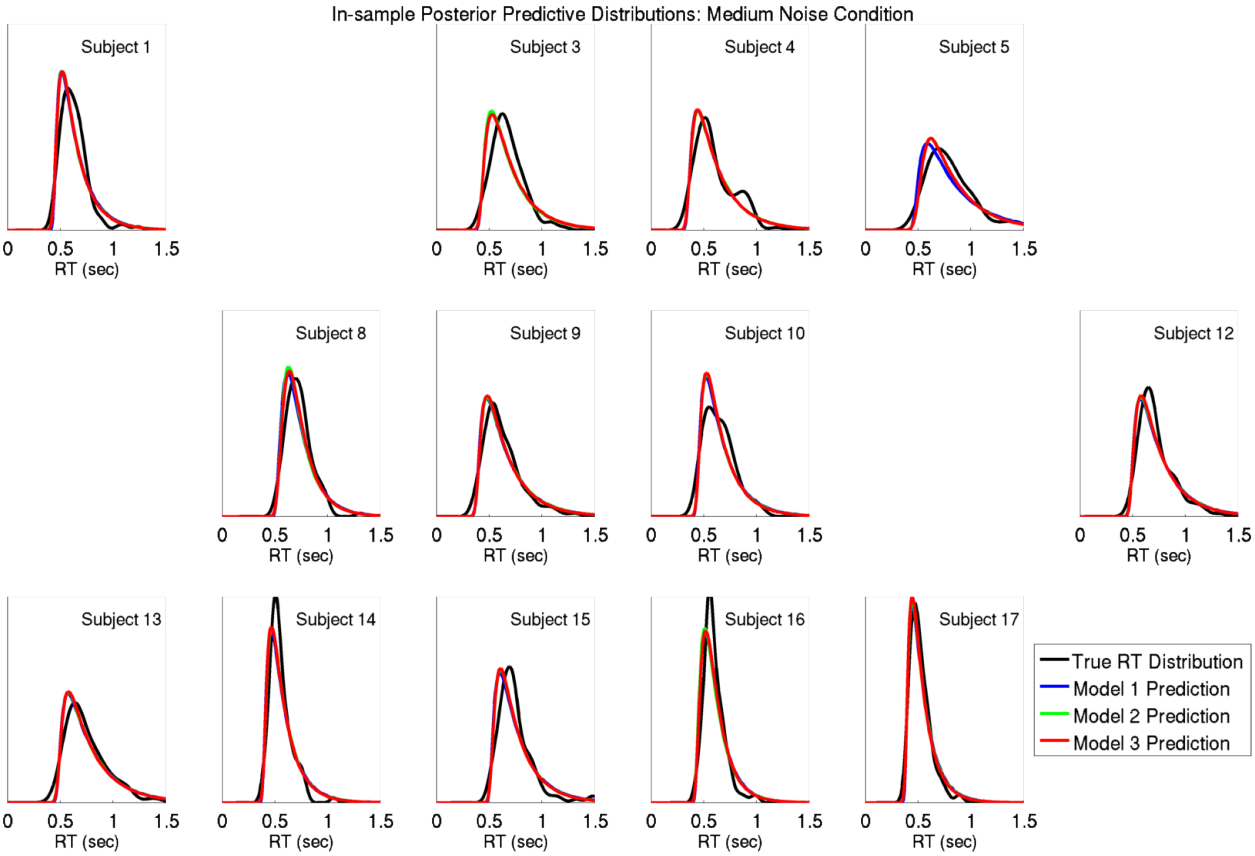

Figure B.2: For each subject, posterior predictive distributions of correct-RTs during trials from the medium noise condition were compared to true correct-RT distributions from training data. Each model performs well at predicting training correct-RT data from known subjects in the medium noise condition. More comprehensive evaluations of in-sample prediction are provided in **Table 3** in the paper.

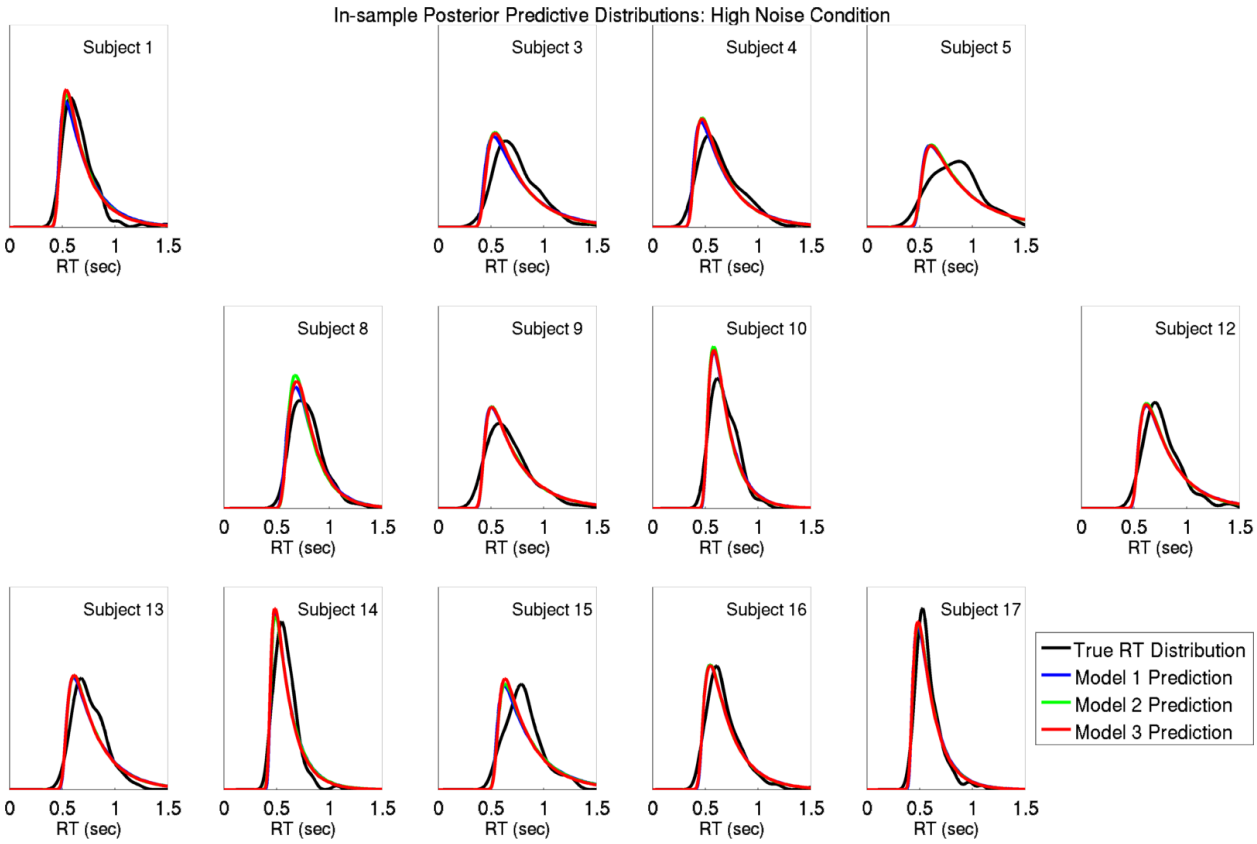

Figure B.3: For each subject, posterior predictive distributions of correct-RTs during trials from the high noise condition were compared to true correct-RT distributions from training data. Each model performs well at predicting training correct-RT data from known subjects in the high noise condition. More comprehensive evaluations of in-sample prediction are provided in **Table 3** in the paper.

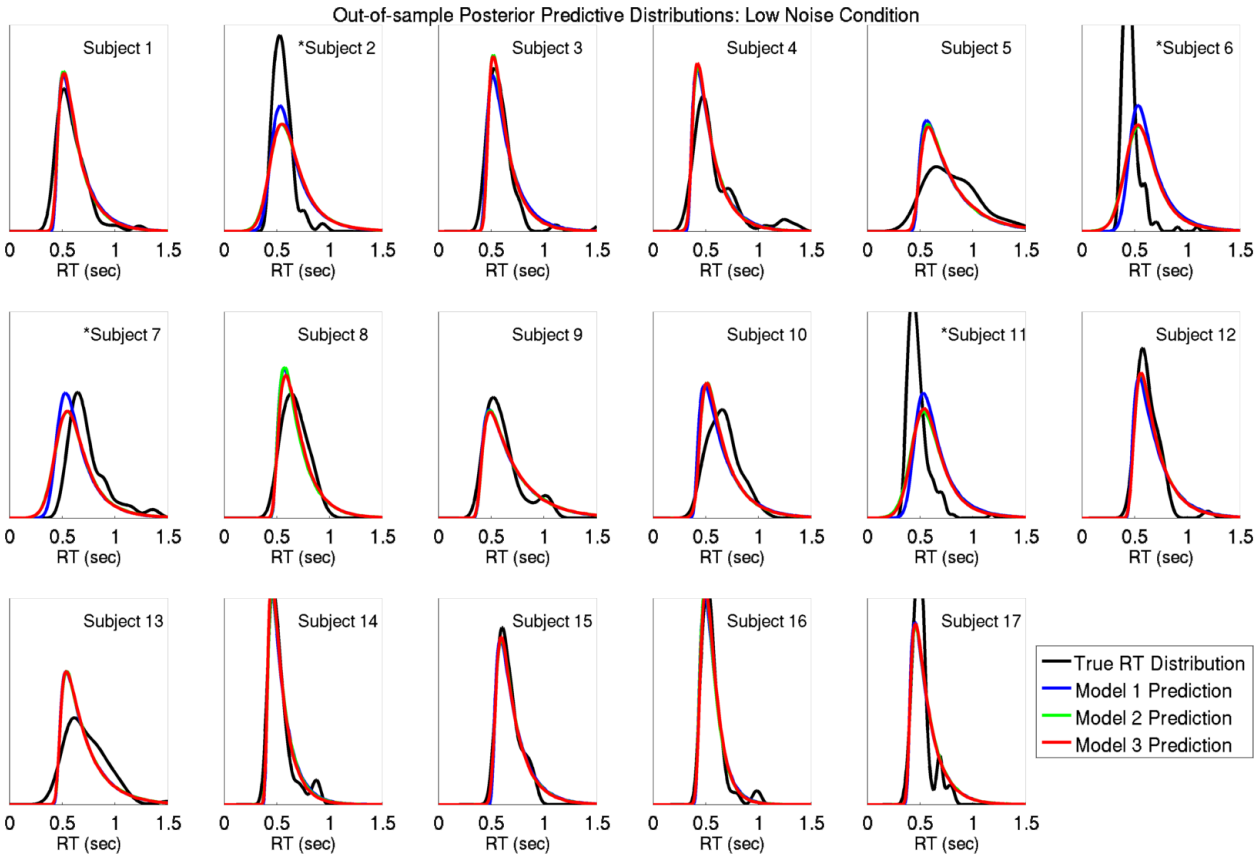

Figure B.4: For each subject, posterior predictive distributions of correct-RTs during trials from the low noise condition were compared to true correct-RT distributions from test (i.e. "out-of-sample") data. The predictive ability of Models 2 and 3 were influenced by observed single-trial EEG data during the randomly-assigned test trials. Note that the predictive ability of each model for subjects 2, 6, 7, and 11 is decreased in comparison to the other subjects because these subjects were left out of the training data. More comprehensive evaluations of out-of-sample prediction for both "known" and "unknown" subjects are provided in **Tables 1** and **2** in the paper.

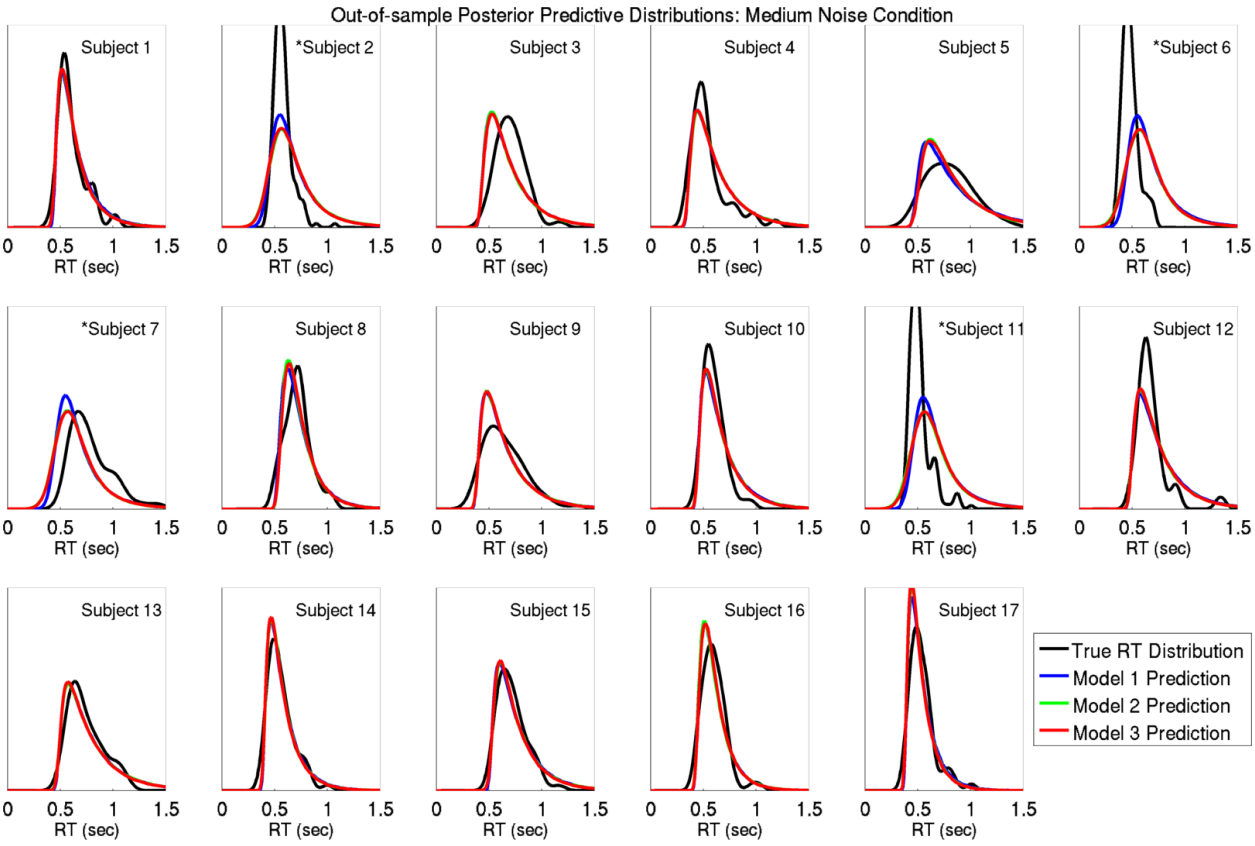

Figure B.5: For each subject, posterior predictive distributions of correct-RTs during trials from the medium noise condition were compared to true correct-RT distributions from test data. More comprehensive evaluations of out-of-sample prediction for both "known" and "unknown" subjects are provided in **Tables 1** and **2** in the paper.

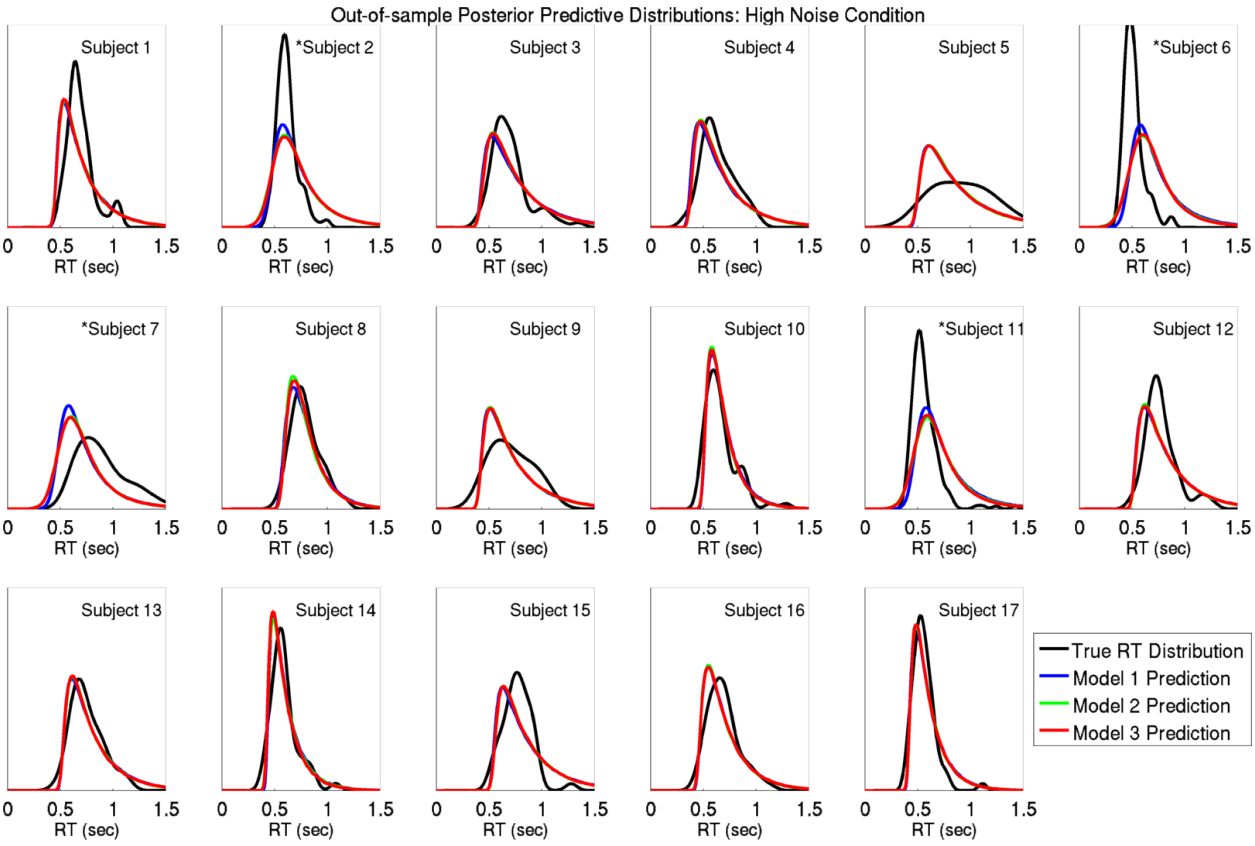

Figure B.6: For each subject, posterior predictive distributions of correct-RTs during trials from the high noise condition were compared to true correct-RT distributions from test data. More comprehensive evaluations of out-of-sample prediction for both "known" and "unknown" subjects are provided in **Tables 1** and **2** in the paper.