2019

Cooperative Success Under Shared Cognitive States and Valuations

Shannon Marie Klotz
Scripps College

Cooperative Success Under Shared Cognitive States and Valuations

A Thesis Presented

by

Shannon Marie Klotz

To the Keck Science Department

Of Claremont McKenna, Pitzer, and Scripps Colleges

In partial fulfillment of

The degree of Bachelor of Arts

Senior Thesis in Neuroscience

April 29th, 2019

Table of Contents

Abstract 3

Historical Background 4

Challenges 6

Proposed Solutions 21

Methods 30

Results 37

Discussion 40

Acknowledgements 42

References 43

Appendix A 52

Appendix B 53

Abstract

A mental model of the another person's state of mind including their thoughts, feelings, and beliefs, otherwise known as Theory of Mind (ToM), can be created to better predict their behavior and optimize our own decisions. These representations can be explicitly modeled during both the development and presence of stable cooperation via communication outcomes, allowing us to understand the sophistication or depth of mental coordination, involved in an individual's social perception and reasoning. Almost all current scientific studies of ToM take a spectatorial approach, relying on observation followed by evaluation (e.g., the Sally-Anne Task). However given evidence that social cognition fundamentally shifts during valuationally significant social encounters with others, this study adopts a second-person approach. Each participant's actions under dynamic uncertainty influence the joint reward probabilities of both, favoring cooperation and coordination. Only Teachers have knowledge of the correct action-reward contingencies, while Learners must ascertain the Teacher's directive and correctly adjust their actions to obtain the optimal reward. The complexity of cooperative behaviors cannot be captured with simple reinforcement learning models, however a similarity in valuation exists, probing further investigation.

Keywords: Decision Making, Theory of Mind, Neuroscience

Historical Background

Cooperative success in a dynamically updating world requires transmission, acquisition, and application of expertise and knowledge in social settings. By both learning another's thoughts, feelings, and beliefs and being able to distinguish them from one's own, otherwise known as developing Theory of Mind (ToM) or mentalizing, we can create a mental model of the Other's state of mind and ultimately better predict their behavior to optimize our own decisions. These mental representations can be explicitly modeled during both the development and presence of stable cooperation via communication outcomes, allowing us to understand the sophistication, or degree of mental coordination, involved in each individual's social perception and reasoning. We commonly find ourselves in states of asymmetric knowledge distribution when we or others have better access to useful information, and in situations in which cooperation is advantageous.

There exists a range of approaches throughout psychology and neuroscience that are taken to understand cooperation, Theory of Mind, and the dynamics of updating values based on changing signals from the environment. Typically, the approaches typically are used independently of one another but could be combined to yield insight beyond the reach of any one of them in isolation. Cooperation is often studied using joint action approaches which rely on coordination in goal-oriented motor tasks (Sebanz, Bekkering, and Knoblich 2006), however joint action approaches do not require modeling another person's beliefs (cite). False belief tasks are commonly used to capture this presence of Theory of Mind in the form of a play or cartoon by giving an individual more information than another and establishing whether they can then differentiate between their own knowledge and the other's (e.g. Sally-Anne task) (Baron-Cohen, Leslie, and Frith 1985; Wimmer 1983). Only answers consistent with the unknowing perspective, not the participant's own, demonstrate the necessary insight. These tests rarely incorporate either competitive or cooperative valuationally consequential actions and instead impose a spectatorial point of view on participants. Probabilistic reinforcement learning tasks permit a vested interest as well as cognitive modeling of behavior and learning over time (Rolls, McCabe, and Redoute 2008). The implicit weight, or expected value, assigned to a reward can be behaviorally modeled using reinforcement learning methods. Inherent values are updated over time via incentive salience and errors in prediction of reward.

By identifying underlying neural processes and networks involved in the computation of the self and the other, we can better comprehend the components comprising and perhaps necessary for successful cooperation. Recruitment and integration of broad neural regions has been found during tasks explicitly requiring ToM (Adolphs 2009; Carrington and Bailey 2009), notably including the temporoparietal junction (TPJ), medial prefrontal cortex (mPFC), superior temporal sulcus (STS), and precuneus (PCUN) (Fig.1). As the mPFC has been shown to have greater activation while moderating one’s social behavior and executive functioning and the TPJ is believed to have a role in self-other distinction processes, they are commonly referred to as areas of a mentalizing network, leading to our focus on their role and neural connectivity in integrating Theory of Mind and expected value in cooperative learning.

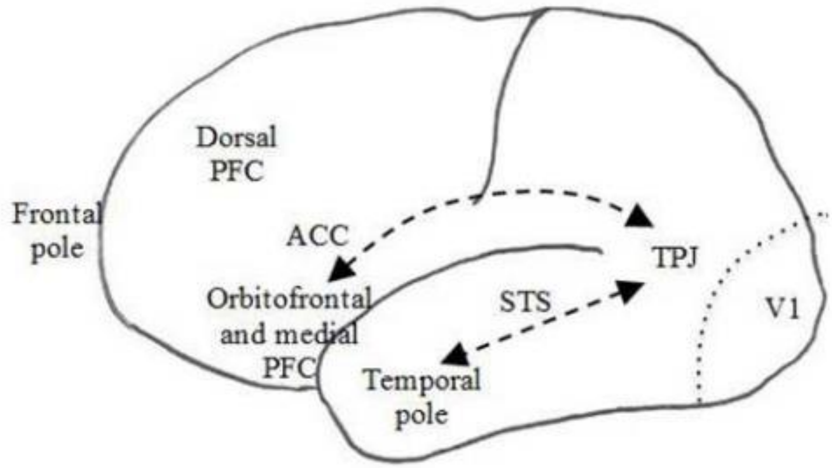

Figure 1. Theory of mind regions of interest (ROIs) in the brain.

This work aims to examine individual and shared valuations in relation to a neural network of cooperation and Theory of Mind integrating evidence from cognitive psychology, social neuroscience, and computer science to apply game theory and reinforcement learning research. By investigating the relationship and underlying neural mechanisms, we may better our understanding of cooperative dynamics in teamwork settings and the requirements for cooperative success.

Challenges

ToM and False Belief Tasks

Social cognition evolved via ongoing interactions and communication between individuals in dyads or small groups, producing a collaborative knowledge network (Dunbar, 1998). In 1981 Stephan, Frahm, and Baron presented a correlation between neocortex size and mean social group size, leading to the creation of social brain hypothesis. It suggested that primates' large brain to body size ratio reflects the computational demands of complex structures in their social systems. Then came a proposal by Dunbar about a limiting volume value for visual areas, meaning that once an organism passed a crucial threshold, additional space could be dedicated to nonvisual areas of the neocortex and cognitive processing, including language. Individuals with more processing capacity were found to have proportionately larger social networks (Lewis et al., 2011; Powell et al., 2012, 2014; Kanai et al., 2012), meaning that their hypothesis applies not just at the level of the species but also at the level of the individual. Communication involves second-order representations in which we represent someone else's representation of our own mental state. Ongoing accumulation and contribution of knowledge enables us to make inferences about the beliefs and intentions of those we are communicating with. These inferences guide our predictions of appropriate behaviors and decisions, leading to more successful interactions. These predictions rely on an intentional stance, also known as having a Theory of Mind or mentalizing (the initial estimate of the mental state of the person we are interacting with) using long-term dispositions, short-term emotional states, desires, beliefs. If the brain can both follow the mental state of the person we are interacting with and make predictions about future behavior on the basis of that mental state using the same mechanism as Bayesian approaches state, there is much more to learn about the brain's organization. The brain is thought to be comprised of social regions (Frith 2007) and this is supported by imaging during mentalizing tasks handling communicative intentions and perspective taking, as well as lesion studies. As social information processing is modular, there are areas that stand out in these studies including the posterior superior temporal sulcus (pSTS), temporoparietal junction (TPJ), and anterior rostral medial prefrontal cortex (armPFC), amygdala, orbitofrontal cortex (OFC), and insula. The first two deal with success during mentalizing tasks, predicting the trajectory of movements where greater activity is associated with prediction errors and understanding that someone has a false belief and therefore different spatial perspectives, respectively. It is not clear whether the region is needed to be effective in the tasks, however, as PFC is elicited in these situations more research is necessary to understand it's role. The amygdala, which is not specific to social cognition, regulates social behavior and recognizing emotional facial expressions, while the orbitofrontal cortex is involved in reward processing, and the insula deals with representing states of our own body, feeling empathy for others (Adolphs 2009; Dodell-Feder et al. 2011; Koster-Hale et al. 2017).

It is believed that theory of mind must have preceded any ability to use language in the communicative way in which it is used today due to the functional limitations it would otherwise experience due to lack of a Broca's area. Some hypothesize as far back as 6 million years (Mithen 1996), but recently evolution by degrees as close as 40,000 years ago has been argued to be the origin (Baron-Cohen 2012), citing the "social intelligence" of monkeys and apes today not being fundamentally the same as the clearly recognize goal states found in humans with a full theory of mind.

Recent advances in technology have spurred unprecedented increases in both the speed of information transfer and size of informational social networks. Unfortunately, the same pitfall of all oral, written, and behavioral communication plagues this expansion of social knowledge networks, unreliable transfer. Whether these networks are between two individuals, such as with a teacher and student or a consumer and a customer service representative, or on a larger scale like a company, it is necessary for any social system with an imbalance of information to quickly and clearly identify the individual with the most current and accurate information, better known as an expert. As knowledge is constantly changing, an answer at one point in time may no longer be the "right" answer later. Similarly, the title of expert is shifting.

Current approaches to understanding knowledge transmission and ToM in these social networks take a spectatorial approach, relying on observation followed by evaluation, such as the case with the classic Sally-Anne Task (Baron-Cohen et al. 1985). In this test a cartoon or play is passively presented to a participant where Sally and Anne are in a room with a ball, a basket, and a box. Sally puts the ball in the basket and proceeds to leave the room. While she is gone, Anne moves the ball to the box. Sally then re-enters the room. The participant is then asked where Sally is going to look for the ball. For a participant to pass this test and thereby possess ToM capabilities, otherwise known as the ability to attribute mental states, they must answer that Sally believes that the doll is in the basket. The participant, holding two contrasting belief states in their own mind must temporarily adopt Sally's limited perspective and disregard their own omnipresent knowledge. There are two separable components of the Theory of Mind mechanism, the Self and the Other. These can be distinguished through distinct reaction times and neural signatures. Humans constantly process our own perspective regardless of ultimate interaction demands, whereas the Other perspective is only processed when explicitly required (Bradford, Jentzsch, and Gomez 2015). Behaviorally, these 'perspective-shift' requirements increase the cognitive load of executive functioning processes when attributing beliefs to others and lengthen reaction times. In terms of neuroanatomy, medial regions of the prefrontal cortex and precuneus respond more when tasks require ToM reasoning in relation to oneself and lateral TPJ regions when another person's behavior violates received information about her thoughts and feelings (Thye et al. 2018).

False-belief task paradigms like that of the Sally-Anne Task investigate actual, but limited domains of social cognition, committed to spectator theories of knowledge using "isolation paradigms" (Becchio et al. 2012) to theorize about social interaction and emotional engagement. In isolation paradigms participants observe others or think about their mental states, but are not required to or are restricted from interacting with the actual individual. The social cognition that drives social exchanges and is responsible for social learning fundamentally shifts during social encounters with others (Schilbach et al. 2013). When participants take an all-knowing, but inconsequential perspective which is not fully emerged in the social interaction, such as with spectatorial tasks, they are not actively participating in social interactions (Barresi and Moore 1996) as there is no valuational reward or consequence for their actions.

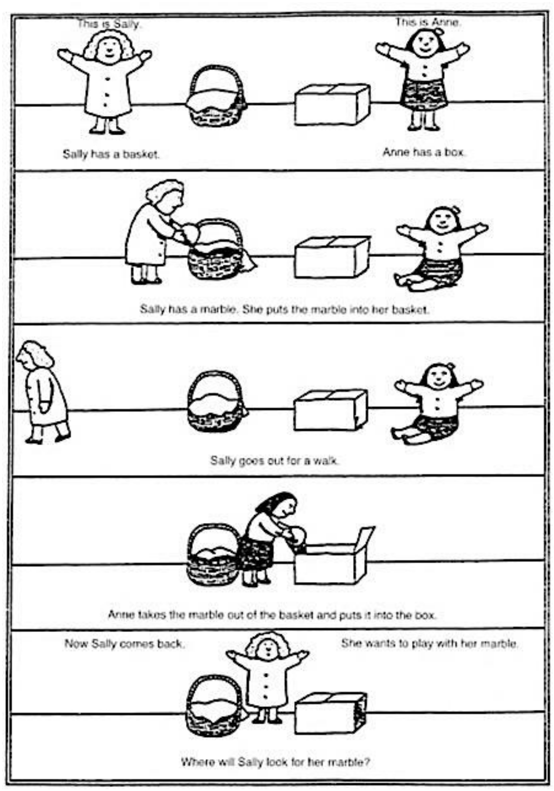

Figure 2. Sally-Anne False Belief Test.

ToM and Competitive vs Cooperative Cognition: Evolutionary and Developmental Considerations

As humans are a uniquely cooperative species, researchers have sought to understand how social interactions may have triggered our cognitive and social evolution. Cooperation is known to be difficult due to tempting immediate self-interest seen through challenges in sustaining relationships. Constructs of trust, prosociality, reciprocity, culture, norm violation, fairness, reputation, learning, adaptation, and reward and punishment all play roles in cooperation, making it a difficult area of study (van Dijk, Parks, and van Lange 2013).

Any in-group or social organization requires some level of cooperation, however we do not always choose to coordinate with members of our community. Competitive behavior spans conventional aggression including territoriality, dominance, and threats, to physical violence. During competitive decision-making, the ACC stores information on whether a course of action demands competition, what the intensity of the competition will be, and, whether or not the competition is 'worth it' to achieve an end reward (Hillman and Bilkey 2012).

Evolutionarily, as social groups became larger, more complex, and mutually dependent, it was crucial to be able to identify others who cooperated and those who did not. Cooperation was previously thought to be either an anomaly or a default mode, however, current research suggest that it is best understood as dependent upon traits and states (De Dreu 2013). Rather than identifying the mental process to build up cooperation, the necessary factors for cooperation are better understood through the moment that cooperation breaks down (De Dreu 2013).

Unitary ToM? Considerations From Cognitive Science and Neuroscience

In *The Architecture of Cognition* a rigorous general theory of mind was proposed that mental life has a common set of principles, or essential unity, despite the diversity of its component functions (Anderson 1996; Holyoak 1983). Through revision of the ACT (Adaptive Control of Thought) system, processing and representation of material in our system is determined via an "if, then" matching network. Social interactions can take different forms, leading to different requirements of Theory of Mind capabilities. Four types have been proposed: shared imaginative play, discourse about inner states, collaborative narratives, arguments in disputes, exposure to and engagement in deception (Dunn 1999a). Anderson's theory is strongly founded on goals, meaning the functionality or type of interaction can vary the degree of activation in the network.

Computational Cognitive Modeling

To understand a system as complex and as diverse as the human mind, we need more than just simulation tools or programming languages. Proper exploration of the various cognitive functionalities underlying human behavior demands development of detailed, process-based understanding by identifying computational models of the representations, mechanisms, and processes involved (Sun 2008).

When modeling social exchanges, researchers must keep in mind that the brain must both represent the available choices and calculate the differential value of each. This is done by developing expectations about the present and future concerning the magnitude of reward or punishment resulting from a choice. One approach to finding the neural and psychological mechanisms involved in valuation has been through implementation of reward stimuli to map out brain responses by varying the dimension or importance of the rewards. Unfortunately, this can only give partial information about neural responses related to valuation. To best capture what problems the brain is trying to solve and what mathematical formalisms and computational algorithms solution the brain employs under reward, we implement reinforcement learning models. These models assume that the seen mechanisms represent only one valuation and decision-making system within the brain, and the neuroimaging findings regarding valuation in social exchange, decision-making, and learning similarly show striking consistency in the set of neural structures that respond to rewards across a broad range of domains, conforming to predictions made by formal models of reinforcement learning (Montague, King-Casas, and Cohen 2006).

Localizing Neural Substrates

In order to map the unobservable dynamics of mental processes to the physical dynamics in the brain we use both formal cognitive models (discussed below) and measurable signals from the brain. Conventional electroencephalography (EEG) continuously and passively records electrophysiological activity of neuronal assemblies, or local field potentials (LFPs), using electrodes placed on the scalp. The electrodes are held to the scalp via a cap, while the data is transferred to an amplifier and ultimately a data acquisition computer (Fig. 1). Inferring brain areas responsible for underlying neural activity of interest, in this case using scalp potential data, is also known as solving the forward problem (Hallez et al. 2007).

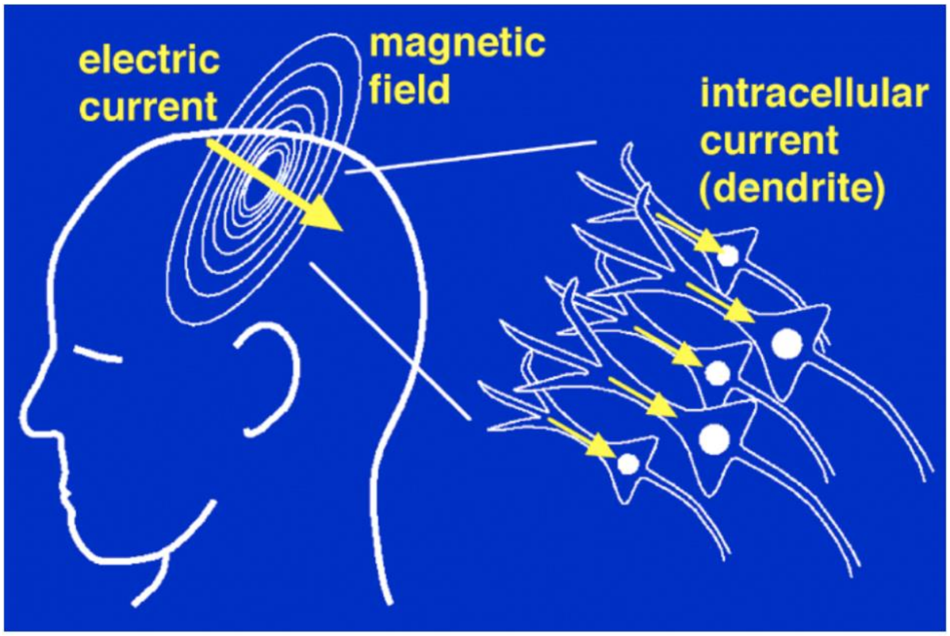

Figure 3. Field potentials elicited via dendritic ensembles.

EEG is known to exhibit exceptional temporal resolution of around 1 millisecond, however it has difficulty localizing neural activity in space. Compared to microelectrode recordings that can reflect the sum of action potentials of cells in a region of 50-350 μm (Gray et al. 1995; Legatt, Arezzo, and Vaughan 1980), early EEG could only limit its range to approximately 6-9 cm (Nunez et al., 1994; Babiloni et al., 2001). As neuroimaging technology improves, standard guidelines as to how researchers approach data, it's results, and what this means for the association are necessary. Cacioppo et al. propose four: 1) we already know that brain activation is corresponded to some aspect of social cognition, emotion, or behavior, so shallow studies that claim discoveries from covariation of a psychological measure, process, or representation and an event in the brain bring little to the table; 2) while focusing on regions of interest can be of use and finding localizations of information-processing components would provide much simpler answers to our questions, the brain needs to be considered a network that requires its total mass to carry out functions;

3) erroneous interpretations arise from belief that changes in activation during a certain information-processing operation do not necessarily mean that the region is the neural substrate for that operation; 4) although neural results can be stunning and colorful, this does not mean that they are of psychological significance (Cacioppo et al. 2003). These practices help us to move past shallow research questions and practices and look instead at processes, consequences, and conditions using all of the power available to us.

While improvements in EEG acquisition and head modeling are enhancing the spatial resolution of scalp EEG (Ferree, Clay, and Tucker 2001), the skull continues to be an impediment due to variance in conductivity. Due to physical distance from the neurons and potential distortion from the skull and scalp, even the best spatial resolution remains around 1-2 cm (Ryynanen, Hyttinen, and Malmivuo 2006). Although this approach does not allow for single cell recordings, we can measure signal from neurons working in concert. Due to similar changes in electrical polarity of cortical pyramidal cells' postsynaptic dendritic currents, the summed electric current flowing from multiple nearby neurons within a small volume of nervous tissue can be investigated.

Figure 4. BioSemi ActiveTwo EEG cap with amplifier.

Localization of these assemblies, or source reconstruction, requires realistic head models of the participants. We can use standardized models of head geometry and brain tissues spatially warped to a template brain to reliably identify brain structures and compare participants.

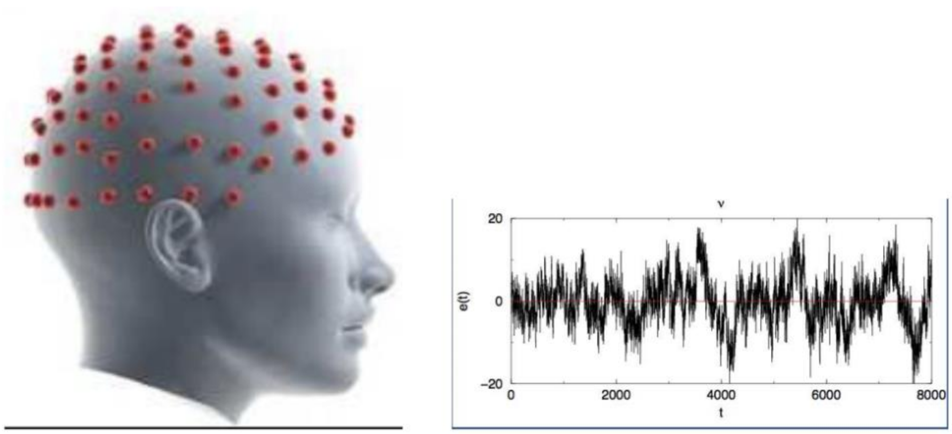

Figure 5. EEG components with time series from a single electrode.

To find patterns or algorithms that can help us make a distinction of categories in the mind, we use an encoding approach. We create a map from the stimuli to a neural response by cataloging how neurons respond to a wide variety of stimuli, however this doesn't provide any information about underlying temporal dynamics or patterns (Dayan and Abbott 2005).

Systematic investigation of the mind and brain in early brain imaging studies concentrated on labeling the neural substrates of psychological processes. Theories at this time considered the brain as an information processing system where modular neural structures were recruited to perform unique cognitive tasks (Fodor 1983). Accordingly, the first neuroimaging studies set out to localize cognitive functions to specific brain regions by finding significant "blobs". These brain-mapping studies provided a foundational understanding of the mind's neural organization and paved the way for modern conceptualizations of the mind and brain. Novel research has focused on understanding patterns of activation and connectivity over a region rather than mean activation in that region localized into blobs. Today, we know that the same neural structures are often involved in a diverse array of cognitive processes (Yarkoni et al. 2011). In addition, multiple structures are recruited to perform different computational tasks due to dynamic interactions between various brain structures. These structures work together or one exerts an influence over another in order to perform computations that are greater than the sum of the individual subcomponents (Bassett and Gazzaniga 2011), allowing complex higher-order cognitive functions to emerge from dynamic interactions between distributed networks of neural modules (Strogatz 2003). Evidence is ever increasing that supports this view that different cognitive processes result from unique interactions between neural networks (Bullmore and Sporns 2012; Davison et al. 2015; Gu et al. 2015; Hermundstad et al. 2013, 2014; Petersen and Sporns 2015). When a region of the brain is involved in various cognitive processes in conjunction with other areas, mapping of the one region's average neural activity provides limited information about the neural organization of the mind. Investigating how connected neural structures interact with one another in order to enable various cognitive processes provides clearer associations from neuroimaging data to cognitive theories and computational models, which offers a sounder foundation for linking brain activity and mental function (Poldrack 2012). Therefore, if we wish to understand how social experiences modulate brain states, then we must explore how functional connectivity differs over social interactions while acknowledging that structural connectivity permits functional relationships (Huskey 2016). While the field of neuroscience has seen an increasing focus on connectivity and pattern analysis to assess brain–network connectivity, there remains a residual influence of blobology and a resistance to networks.

Shared Cognitive States

Currently, two main theories for how we infer other people's mental states and understand the actions and emotions of others exist. Simulation Theory, in which simulation and empathy reside in the premotor cortex and insula and Theory-Theory which demands more deliberate theory-of-mind abilities involving the medial prefrontal cortex and the temporoparietal junction (Adolphs 2009).

The mirror neuron system (MNS) provides evidence for a simulationist account of social cognition and is believed to give us a "first-person grasp" of the motor goals and intentions of other individuals by internally replicating or simulating observed events through a mirror mechanism. While humans can reason about other's actions and emotions at a conceptual, declarative level, this hypothesis argues that explicit reflection is not required to directly grasp of the mind of others. As we witness others carrying out similar actions and experience similar emotions to ourselves, such as disgust or pain, our brain has the capacity to associate these third-person experiences as first-person experiences (Gallese, Keysers, & Rizzolatti, 2004). This mechanism was first supported by evidence in macaque monkeys, showing a system of F5 neurons in the motor cortex that responded both when monkeys observe or performed a goal-directed behavior (Gallese et al. 1996) and reproduced to a lesser extent in humans (Rizzolatti et al. 1996). Consequently, when conducting analysis with this approach, researchers aims to identify the computations used by shared circuits. Beyond individual neuron spiking activity, analysis of mu rhythm using EEG confirms a decrease in amplitude during both action execution and action observation, indicating cortical activity for both events (Fox et al. 2016).

While the Mirror Neuron Theory presents an attractive model of empathy, some researchers have noted the flaws in both original macaque experiments (Hickok 2009), as well as those involving human participants (Debes 2010) including misapplication of results about goal understanding to action understanding. An alternative to the MNS is the ToM or mentalizing network (MENT). This theory relates closely to our understanding of imitate facial gestures, as neither requires a simulation to construct a third-party's representation. Just as we can tease apart our own proprioceptive experience from the visual perception of another person's face using innate body schema to integrate sensory and motor systems (Gallagher and Meltzoff, 1996), conceivably there could be some cognitive schema that enables the interactive experience (Gallagher 2001). This instead provides evidence for a "Theory-Theory" account of social cognition, believed to give us an inferential, reflective, and third-person grasp of others' mental states due to an intentional schema or inductively formed representations of an intentional relation (Goldman 1989, Barresi and Moore 1996).

Temporal Dynamics

Thoughtful experimental design, recording techniques, and accurate detection and correction of artifacts, are important precursors to accurate results and their interpretation. When analyzing the neural data, it is common to explore the time-domain which shows how a signal changes over time. Converting the results of the time-domain using a Fourier Transform produces the frequency-domain doamain (Fig. 6) which allows us to see how each frequency band contributes to the signal by assuming that a waveform is comprised of oscillations. The frequency domain often an often offers a more intuitive understanding of the qualitative behavior of the system. Critiques have been made about not making use of frequency domain or in turn abusing it (Luck 2005). For example, contrary to some applications, power in a given frequency band is not evidence of an oscillation in that band, in fact false oscillations can be distinguished in this domain as the produce a broad band of power. These domains can be Fourier transformed again to result in a joint time–frequency domain (Fig. 7).

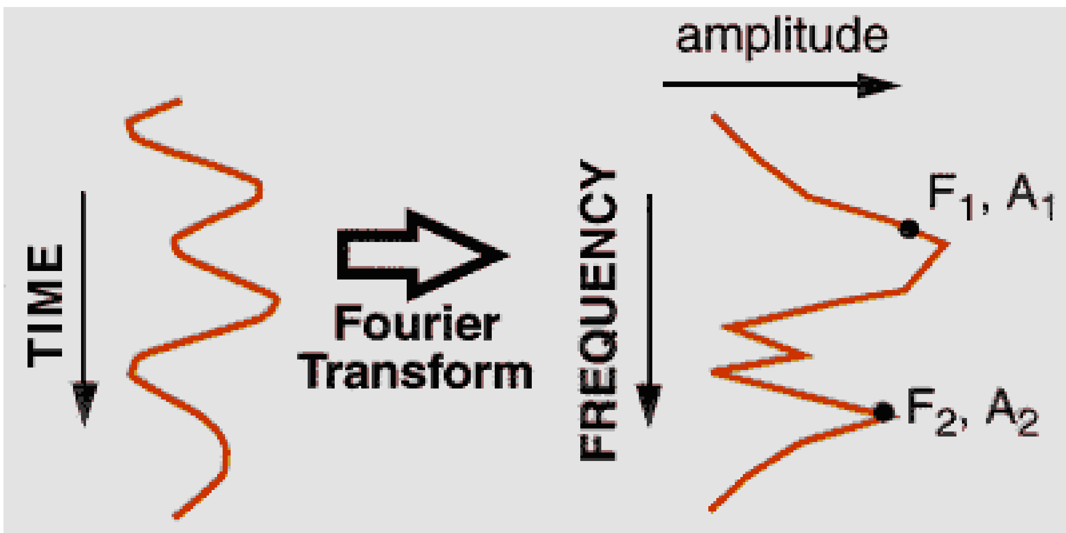

Figure 6. Conversion of time domain to frequency domain using the Fourier Transform.

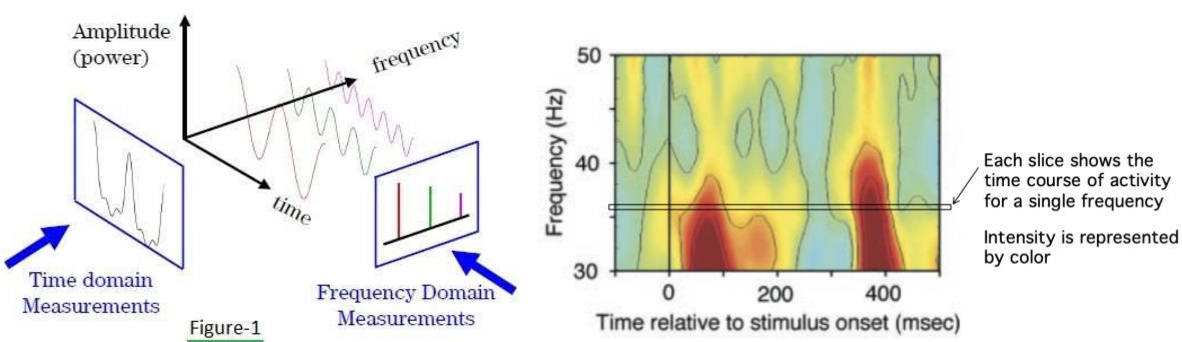

Figure 7. The time-frequency domain.

Proposed Solutions

Interactive Second Person Neuroscience

Detached observers of interaction, as seen in implementations of the Sally-Anne task, do not capture the complex processes involved in social cognition. To better understand the mechanism involved in mentally separating the Self and the Other, each with intentionality, joint attention, and goal-directed action (Moore and Paulus 2013) our approach recognizes the relationship and recurring transition between first and second-person perspectives of the state of the world (Longo and Tsakiris 2013). Interactions produce shared intentions and motivations with each joint action having a consequence. This project aims to move toward measuring actively engaged participation to generate more ecologically valid data (Schilbach et al. 2013), specifically in interpersonal understanding and emotional engagement in valuationally consequential decision making.

Producing controlled experimental conditions in a direct and unrestricted social environment generates a multitude of confounding factors. Our solution allows for a nondeceptive interaction with a partner that they not only meet, but are seated next to for the duration of the task.

Modeling Cognition

Mathematical psychologists rely on behavioral data to evaluate formal models of cognition, whereas cognitive neuroscientists rely on statistical models to understand patterns of neural activity. In this study, we utilized a combined brain and behavior approach to provide novel insights into the mechanisms underlying human cognition through modeling. Model-based cognitive neuroscience decomposes complex behavior into representations and processes and these latent model states can be used to explain the modulation of brain states under different experimental conditions both informing cognitive modeling and help to interpret neural measures. It attempts to go beyond prior approaches that consider, but not connect Marr's classic three levels of analysis, computational, algorithmic, implementation (Palmeri, Love, and Turner 2017). There are many possible paths for integrating theoretical modeling and cognitive neuroscience (Turner et al. 2017) as individually tailored, creative solutions are often needed to establish compelling links between multi-parameter models and complex sets of neural data that bypass the traditional barriers between models and the brain (Pratte and Tong 2017). Computational modeling examines latent neurocognitive processes and interactions, with their parameters revealing psychologically meaningful individual differences, making accurate parameter estimates critical.

When reasoning about value-based decision-making, we can measure behavior under varying reinforcement contingencies, where there is an adjustment of choice behavior in response to abrupt reversals in the prevailing reinforcer conditions. Individuals have been found to reach an average steady-state of responding after obtaining relatively few reinforcers under new conditions(Lau and Glimcher 2005). Reinforcement learning provides a computational approach to learning from interaction with our uncertain environment and the people within it because it incorporates a participant's memory of both past reinforcers and choice. In order to better understand how we weigh our different options and alter decisions response-by-response, we need to first be familiar with the amount of influence each past outcome has on our next decision. The weight given to decision outcomes should reflect their salience in predicting future outcomes. This salience or representation of reward expectation and probability in the brain can be seen behaviorally via the learning rate, a fundamental feature of behavior that determines how agents should make the decisions in the face of changing circumstances (Behrens et al. 2007). Exercising this connection produces a wealth of information about cause and effect, about the consequences of actions, and about what to do in order to achieve goals. Neither exploration of the environment nor exploitation of current knowledge can be pursued exclusively without failing at the task. Without direct instructions as to what actions to take or where the consequences, or complete models of the environment, uncertainty is prevalent. However by introducing the reward signals participants can create a policy for action, a reward signal, a value function, and, optionally, a model of the environment (Sutton and Barto 1998). Reinforcement learning provides both qualitative and quantitative frameworks for understanding and modeling adaptive decision-making in the face of rewards and punishments, using both model-based and model-free strategies.

Underlying Neural Networks

Real-world data is frequently incomplete, inconsistent, and is likely to be noisy or include errors. Large arrays of data often entail relatively small amounts of useful information, therefore we must extract information from the dataset to reveal the driving forces which underlie a set of observed phenomena. The phrase "garbage in, garbage out", attributed to Stephen ‘Wilf’ Hey, refers to the concept that irrelevant and redundant information produces nonsense output or "garbage", making detection of valuable patterns challenging. Preprocessing (EEGLAB Wiki - SCNN, Brunner, Delorme, and Makeig 2013; Delorme et al. 2011; Delorme and Makeig 2004) is a worthwhile data mining technique to enhance the representation and quality of the continuous EEG data. Cleaning, normalizing, and transforming raw data into an understandable format by filling in missing values, smoothing noisy data, identifying or removing outliers, and resolving inconsistencies can allow for superior interpretations of the data later on.

Currently, it is believed that independent component analysis (ICA) is an effective method for removing artifacts and separating sources of the brain signals from EEG recordings recorded across the scalp and postsynaptic signals from the brain. ICA attempts to exploit independence to perform blind source separation mixtures of signals from multiple brain generators to find their origins, exact dynamics, and relationship to brain function (Jung et al. 2001). Signals from the brain are composed of a linear combo of sinusoidal waves from various independent sources similar to auditory signals described in the cocktail party effect (Cherry 1953; Cherry and Taylor 1954). Teasing apart signals to identify discrete pieces of information can be accomplished through blind source separation. In ICA, the mixture of underlying factors or source signals is made statistically independent from each other, where one signal provides no information regarding the other signals, assuming that different physical processes yield independent outputs (Stone 2004). A commonly used alternative is principal component analysis, which find a set of signals which are uncorrelated with each other, capturing the most variance from the data in component one and so on (Joliffe and Morgan 1992; Jolliffe and Cadima 2016). However, this method assumes orthogonality between components and that the data comes from a highly-dimensional Gaussian distribution. It is also limited by EEG technology, such that the number of components cannot exceed the number of electrodes used.

Useful Approaches

One method of statistical inference, estimating parameters of a statistical model given observations, is Bayesian modeling. Bayes' theorem updates the probability of a hypothesis as access to more evidence or information increases. Bayesian inference develops the posterior probability from a prior probability and a "likelihood function" derived from a statistical model for the observed data. A posterior predictive distribution is used to do predictive inference, or predict the distribution of a new, unobserved data point. In frequentist inference, conclusions about data are drawn from a set of repetitions of such experience, each producing statistically independent results. For a frequentist, a probability function would be a simple distribution function, but Bayesian inference computes the posterior probability according to Bayes' theorem. In turn, a frequentist approach obtains the parameter estimates, maximum likelihood estimates (MLE), by finding the parameter values that maximize the likelihood function for each individual separately. These estimates are then used to specify the distribution of a data point. However, individual MLE estimates are often noisy especially when there is an inadequate quantity of data. A group-level analysis, which estimates a single set of parameters for the whole group of individuals, may generate more reliable estimates, but ignores individual differences.

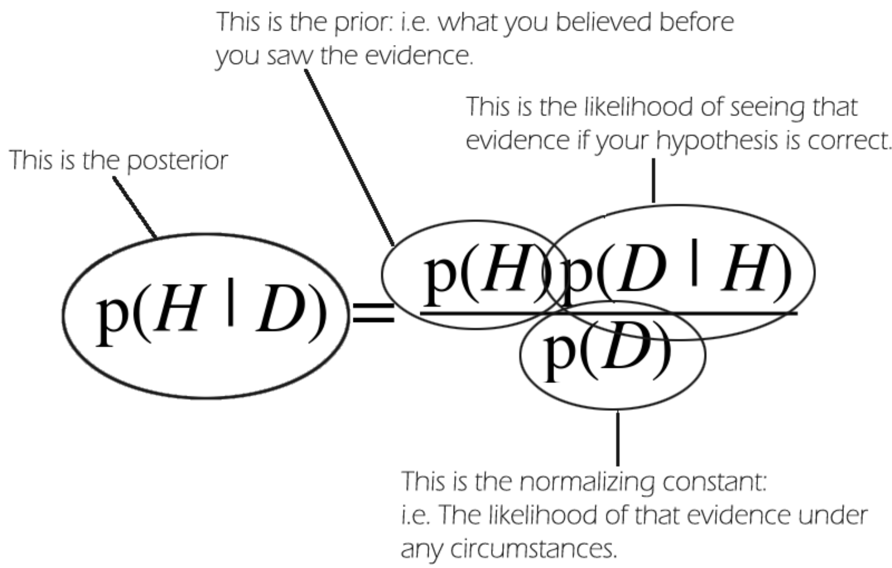

Figure 8. Bayes Theorem.

Bayesian updating is particularly important in the dynamic analysis of ordered data, which advantageous in most behavioral studies due to their design. Bayesian decision theory is a statistical system that tries to measure the probabilities and costs of decisions (Berger 1993; Bernardo and Smith 1994). An agent, or estimator, operating under such a decision theory uses the concepts of Bayesian statistics to estimate the expected value of its actions, and update its expectations based on new information. Due to recent advancements, the power of Bayesian analysis in neuroscientific research has grown significantly in the past few years. Bayesian inference is merely the reallocation of credibility across a space of candidate possibilities, at the same time providing complete information about the credible parameter values to assess our uncertainty. When contrasted the Bayesian approach with the t-test from Null Hypothesis Significance Testing, the latter falls short. Bayesian estimation for two groups provides complete distributions of credible values for the effect size, group means and their difference, standard deviations and their difference, and the normality of the data and it is also more intuitive than traditional methods of NHST. Due to Bayesian power analysis and the fact that Bayesian methods can accept the null value, there is a relative poverty of information provided by the NHST. It has foundational logical problems, such that the traditional t test assumes that the data in each group come from a normal distribution, which is not appropriate when the data contain outliers, and that NHST is based on sampling intentions and different intentions change the interpretation of the data. Bayesian methods instead yields an explicit posterior distribution over parameters unaffected by sampling intentions and can accommodate realistic data situations: non-normal data distributions, censored data, unequal variances, unbalanced sample sizes, nonlinear models, and multiple layers of hierarchical structure in a straightforward unified framework, providing the richest, most informative, and meaningful results for any set of data (Kruschke 2013).

The two main theories proposed for explaining the process of ToM development: Theory-Theory and simulation theory, are well studied behaviorally and neurally. Unfortunately there are few proposed computational models of ToM understanding. Support for mixed reasoning strategies has led to the search for a computational model that integrates theory-based and simulation-based strategies for false belief reasoning. A causal Bayesian framework rationally integrates both theories for false belief reasoning representing the internal models and their interactions, predicting that children's abilities to understand diverse beliefs and knowledge access multiplicatively contribute to their ability to understand false beliefs (Asakura and Inui 2016).

Brms is a package to fit models using a Bayesian approach in a multilevel context where prior distributions that actually reflect their belief can be applied. A wide range of distributions and link functions are supported and model fit can easily be assessed and compared with the Watanabe-Akaike information criterion and leave-one-out cross-validation. This allows for the modeling of data measured on different levels at the same time, thus taking complex dependency structures into account with linear predictors. The posterior distributions of more complex models including MLMs could previously not be found analytically, however recent advancements, along with the program's capabilities and speed, have allowed researchers to investigate population level, group level, and family specific parameter estimation using full Bayesian inference (Bürkner 2017).

hBayesDM is a user-friendly R package that allows researchers to perform hierarchical Bayesian analysis (HBA) of various computational models with a single line of coding (Ahn, Haines, and Zhang, 2017). HBA is regarded as the gold standard for parameter estimation, especially when the amount of information from each participant is small. However, many researchers interested in HBA often find the approach too technical and challenging to be implemented. hBayesDM aims to allow researchers with minimal knowledge of programming to be able to take advantage of advanced computational modeling and HBA.

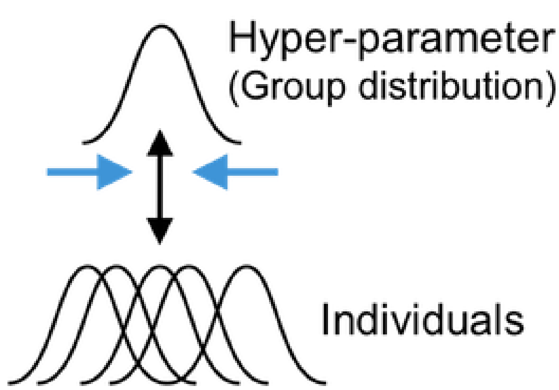

Figure 9. Hyperparameters of the sample used to estimate individual participant parameters.

HBA retains individual differences while using information across individuals to estimate both individual and group parameter estimates (Ahn et al., 2011), delivering full posterior distributions for each participant.

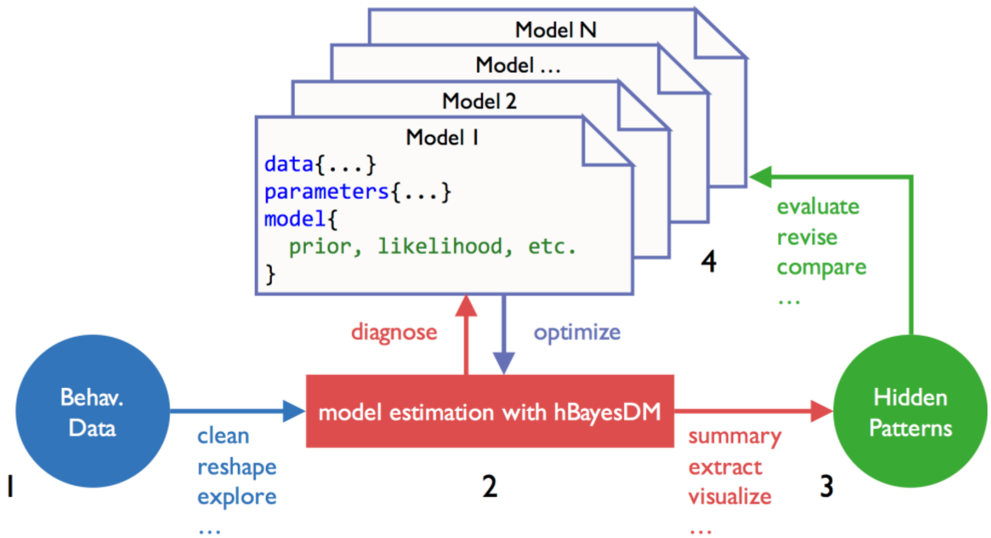

Figure 10. hBayesDM Model Estimation Process.

As cognitive models become more complex, the demands on the behavioral data and empirical tests of predictions derived from quantitative cognitive models increase as well, constantly necessitating more precise tests of hypotheses about brain–behavior relationships. Recently, mathematical models of cognition in the field of decision making in cognitive neuroscience have gained great traction, joining the informational value of behavioral and neural data an breaking the traditional barriers between models and the brain (Palmeri et al. 2017). By allowing both kinds of data to inform model selection researchers can examine common denominators underlying both behavioral data and neural data and ultimately relate two streams of data with potentially vastly different scales. There has been a steady ‘‘tightening’’ of the link between neural and behavioral data (de Hollander, Forstmann, and

Brown 2016; van Ravenzwaaij, Provost, and Brown 2017). Researchers are striving for the tightest possible linking, as this provides the greatest opportunity to investigate the underlying linking assumptions and uncover underlying psychological processes that simultaneously explain behavioral and neural data. Some methods combine both data sets in a single model-based analysis to use neural measurements as a model input in order to predict both behavioral measurements and a second set of neural measurements. Others', like the 'joint' model, instead use individual neural and behavioral models which are then joined by allowing their parameters to covary (van Ravenzwaaij et al. 2017). Two theoretical frameworks underlie the these relationships: composition of discrete latent state either being a mixture of discrete latent states or on a continuous dimension. Neurally-informed cognitive models more reliably recover under a discrete state representation than a continuous dimension representation. Neural data aids the identification of latent states in cognitive models, but unsurprisingly different frameworks for quantitatively informing cognitive models with neural information have different model recovery efficiencies (Hawkins et al. 2017).

Methods

A participant taught the state of the world ("Teacher") to a partner ("Learner") through their choices while the asymmetric knowledge distribution reversed over time with only one partner's awareness during each shift. Each participant starts with no knowledge about which option, either left or right box (Fig. 12), depicted below and during original task instruction as jars (Fig. 11A), will provide them with optimal reward and therefore the reward contingencies are learned through trial and error. At the beginning of every trial both players makes a prediction about the action of the other player and then their own choice. Cooperation is encouraged via a 10-fold payoff structure (Fig. 11B). For example, if both participants choose the high reward jar (green) and receive the high reward, their payoff is no longer 10 points, but 100. This is also true when both choose the low, 5 point jar (red), each receiving 50 points.

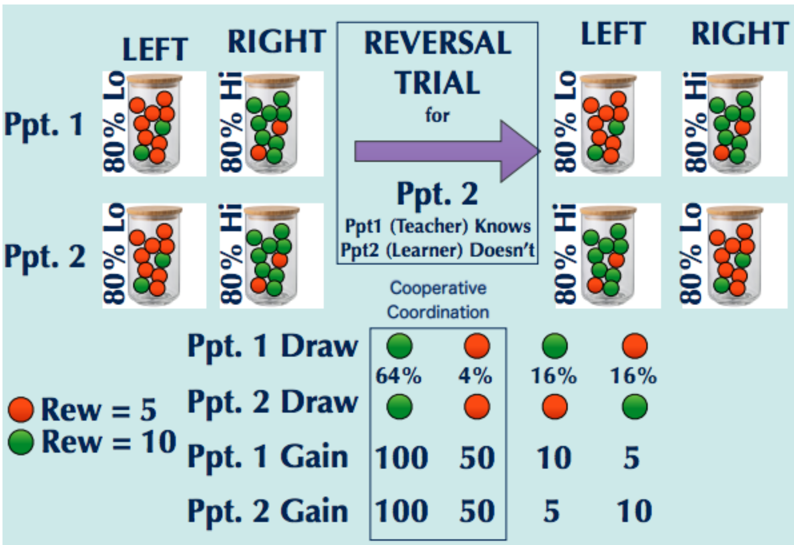

Figure 11. Reward structure of the task. Participants have two options or "jars" to choose from. Red balls are rewards with utilities of 5, and green balls are rewards with utilities of 10.

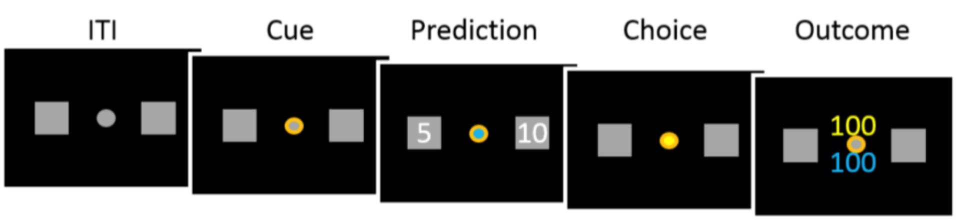

Figure 12. The standard trial sequence of the Interactive Teacher-Learner False Belief Task. The grey boxes represent two figurative jars containing different amounts of reward from which the participants can choose. Upon presence of the blue circle in the center of the screen, participants must predict the partner's action. The change to a yellow circle, indicates that the participant must make a choice for their own behavior. Choices are made through left and right mouse-click. On the outcome screen, numbers in yellow represent the participant's own reward and the blue numbers represent that of their partner.

Once we see that both players have been making decisions that indicate they understand the state of the world, we notify one player that the jars have been reversed for the other player, but that she unaware of this, this is called the reversal trial (Fig. 1). If the notified player, or "Teacher" is rational and implementing at least first level ToM, they will consider the knowledge and beliefs of the other player and will actually choose the suboptimal jar because they recognize that without the new knowledge, the other player is going to continue to choose the same jar as before. The Teacher recognizes that the learner is likely to choose the previously 10, now 5, valued jar and 1) would prefer to get 50 points rather than 10 and 2) needs to communicate that the jars have switched to get back to optimal 100-100 levels.

Participant predictions, choices, and current role were recorded. Presentation of the task (Fig. 12) used the Psychophysics Toolbox in MATLAB.

Self-Report Questionnaires.

Physical paper questionnaires written in German including the following measures were presented to the participants in random order. Open questions were also asked following completion of the task (e.g.: "How did you experience the task (impressions, thoughts feelings, etc.)" and "What did you think about your partner during the task?").

IQ.

The Leistungsprüfsystem serves as a short intelligence quotient test.

Emotional Intelligence.

The Trait Emotional Intelligence Questionnaire–Short Form (EMS) is a seven-point scale with 30 items, half being reverse-scored. It was developed to explore the construct of trait emotional intelligence.

Intolerance to Uncertainty.

The Intolerance to Uncertainty Scale (IUS) is a self-report, five-point scale with 27 items that all fall under the category of either Uncertainty Avoidance and Uncertainty Related Unfairness.

Personality.

The Temperament and Character Inventory Revised (TCI-R) is a psychometric questionnaire with 125 yes-or-no questions covering seven subscales: Harm-Avoidance, Novelty-Seeking, Reward-Dependence, Persistence, Self-Directedness, Cooperativeness, and Self-Transcendence.

Open-Ended Questions.

Open questions were also asked following completion of the task (e.g.: "How did you experience the task (impressions, thoughts feelings, etc.)" and "What did you think about your partner during the task?").

Participants.

All participants were screened using a test version of the task and were asked to answer a questionnaire on task comprehension. Exclusionary criteria included diagnosis of autism spectrum disorders and related neurodevelopmental disorders, psychiatric disorders, visual impairments, neurological disorders, and traumatic brain injury. We invited participants from the student populations throughout the greater Hamburg area. Pairs of participants were matched on gender and similar in age and education. For pilot data collection, participants came to the lab for two days of activity. During the first day self-report questionnaires were completed and the participants received training on the Interactive False Belief Teacher-Learner Task (IFBTFT). On the second day, the participants completed at least three sessions of the task. In the second study using EEG recordings, participants finished the test version of the task, the task itself, and the questionnaires in one day. All participants gave written consent and received debriefing after the experiment.

EEG Recording

Simultaneous, dyadic EEG recordings from both participants used a Biosemi, 128-electrode, active channel electroencephalographic system, sampled at 1024 Hz. Data from both participants were synchronized to one computer with their event codes from MATLAB. EEG data analysis used time-frequency decomposition in EEGLAB. Preprocessing consisted of separation of data into two participant-specific files, addition of electrode locations using a supplied electrode position file, rejection of electrode bridging (eBridge; Delorme et al. 2011), high pass (1 Hz) and low pass (120 Hz) filtering, removal of line noise using the CleanLine algorithm (NeuroImaging Tools & Resources Collaboratory), removal of artifactual data epochs and bad channels using Artifact Subspace Resampling, interpolation of removed electrodes, re-referencing to the average reference across all 128 electrodes, deletion of interpolated electrodes, adaptive mixture independent components analysis (AMICA; Palmer, J. A., Kreutz-Delgado, K. & Makeig, S., 2011), putative source localization of independent components using the MNI brain and equivalent source dipoles (DIPFIT), and selection of brain-related independent components.

Cognitive Modeling

Valuation Modeling

Two probabilistic reinforcement learning task model algorithms (discussed below) that allowed for multiple blocks in the data due to the noncontinuous reversal nature of the data were compared using four variations of outcome values (Table 1). Previous discrete choice modeling of this data has taken place (Lei, 2018), however models of reinforcement learning can provide a more principled approach to assessing behavior, due to freedom from external definitions that the subject is unaware of (learning criterion, point of reversal). After fitting both models on a trial-by-trial basis to each individual using hBayesDM, model comparison of the Leave-One-Out Cross-Validation, pareto-k values, chain conversion plots, distribution plots, and parameter rhat values took place. The best fitting model, Model 2, provided parameters for valuation. Correspondence between the Teacher's prediction and Learner's choice was evaluated.

Table 1. Four possible model outcome values (1 = reward, -1 = punishment)

Model	*5 points*	*10 points*	*50 points*	*100 points*
1	-1	-1	-1	1
2	-1	1	1	1
3	-1	-1	1	1
4	-1	1	-1	1

1) Fictitious Update Model (Gläscher, Hampton, and O'Doherty 2009)
Parameters: "eta" (learning rate), "alpha" (indecision point), "beta" (inverse temperature)

FICT is a incorporates knowledge participants have about the value of the chosen and as well as their belief, or fictitious prediction error, that the unchosen action will have the counterfactual outcome of the current trial (Hampton et al. 2007).

$$\mathbf{V}_{c,t+1} = \mathbf{V}_{c,t} + \boldsymbol{\eta} * (\mathbf{R}_t - \mathbf{V}_{c,t})$$
$$\mathbf{V}_{nc,t+1} = \mathbf{V}_{nc,t} + \boldsymbol{\eta} * (-\mathbf{R}_t - \mathbf{V}_{nc,t})$$

$\mathbf{V}_{c,t+1}$ = value at trial t + 1 for the currently chosen option
$\boldsymbol{\eta}$ = learning rate
$(\mathbf{R}_t - \mathbf{V}_{c,t})$ = prediction error
$\mathbf{V}_{nc,t+1}$ = value at trial t + 1 for the unchosen option
$(-\mathbf{R}_t - \mathbf{V}_{nc,t})$ = fictitious prediction error

2) Reward-Punishment Model (den Ouden et al. 2013)
Parameters: "(α^{pun})" (punishment learning rate), "(α^{rew})" (reward learning rate), "beta" (inverse temperature)

RP is a standard reinforcement learning model based on the classic Rescorla-Wagner model of conditioning that determines whether it is best to separate learning into reward and punishment parameters to capture behavioral strategies (Frank et al. 2007).

$$\mathbf{V}_{c,t} = \mathbf{V}_{c,t-1} + \boldsymbol{\alpha}^{\mathbf{pun}} * (\boldsymbol{\lambda}_{t-1} - \mathbf{V}_{c,t-1}) + \boldsymbol{\alpha}^{\mathbf{rew}} * (\boldsymbol{\lambda}_{t-1} - \mathbf{V}_{c,t-1})$$
$$\mathbf{V}_{nc,t} = \mathbf{V}_{nc,t}$$

$\mathbf{V}_{c,t}$ = value of currently chosen option at trial t
$\boldsymbol{\alpha}^{\mathbf{pun}}$ = punishment learning rate (0 on reward trials)
$\boldsymbol{\alpha}^{\mathbf{rew}}$ = reward learning rate (0 on punishment trials)
$\mathbf{V}_{nc,t}$ = value of unchosen option at trial t

Results

Teacher prediction and choice, as well as Learner prediction and choice were all best-fitted by a form of Model 2 where only outcomes of 5 were valued negatively. Using Leave-One-Out Cross-Validation, pareto-k values, consideration of the width and centrality of posterior parameter distributions, Teacher prediction and Learner choice were best modeled using the fictitious update model, whereas Teacher choice and Learner prediction were best characterized by the reward-punishment model.

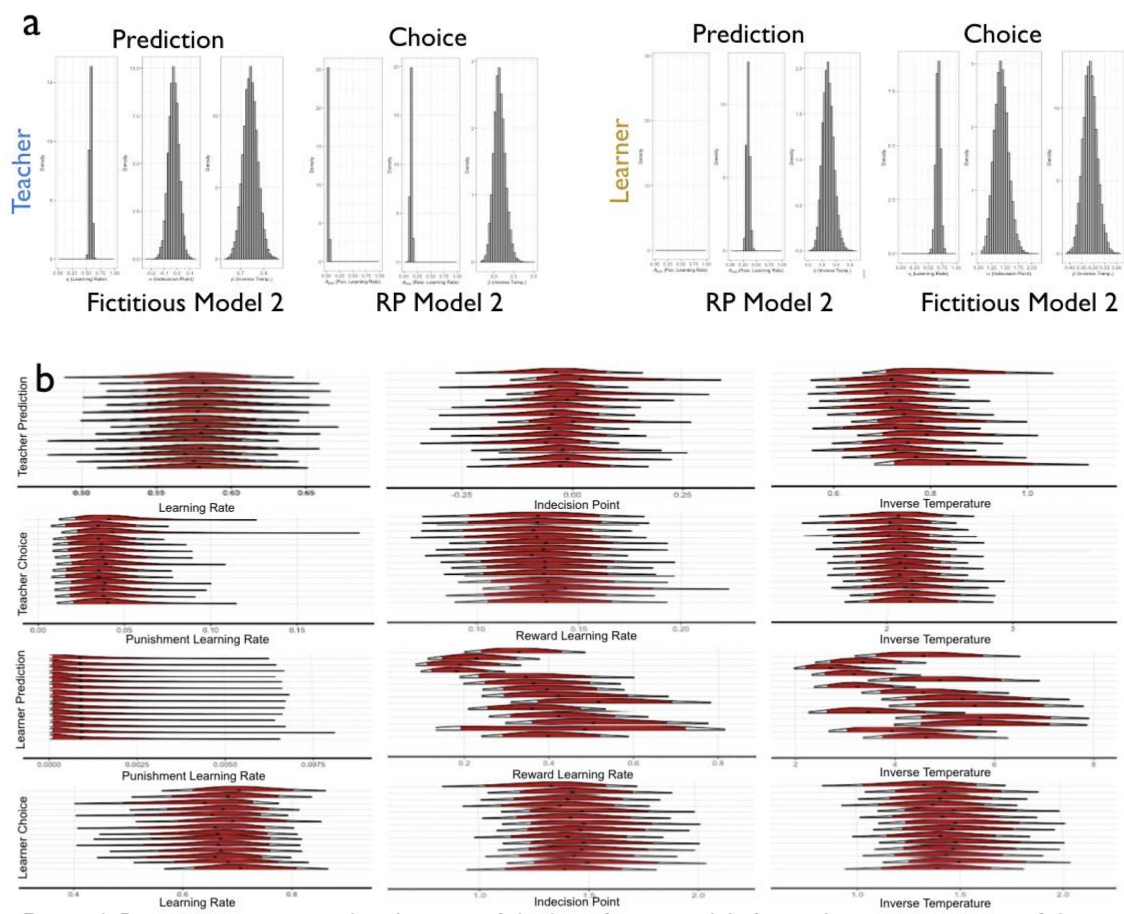

Figure 13. Posterior parameter distributions of the best fitting models for each participant state of the group (a) and 14 participants (7 pairs) who had corresponding independent components in their EEG (b).

However, posterior predictive checks of the models which simulate replicated data under the fitted model and then compare these to the observed data to look for systematic discrepancies between real and simulated data, showed that none of the Rescorla Wagner models provided a reliable representation of the patterns seen in the data.

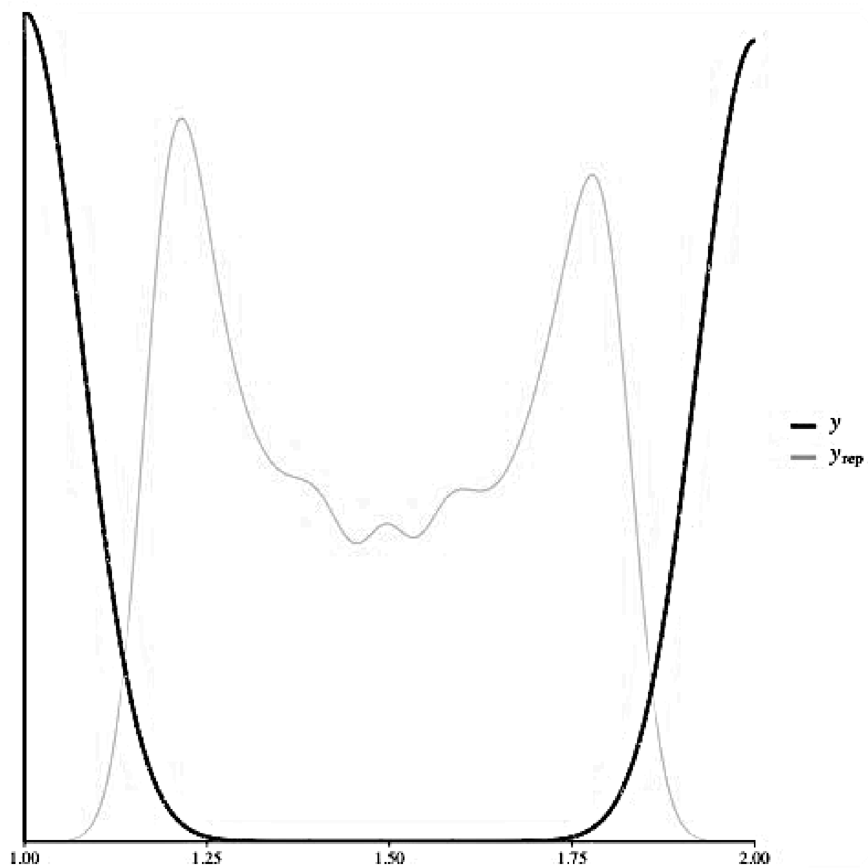

Figure 14. Posterior Predictive Checks of Best-Fitting Teacher Prediction Model; Fictitious Update Model 2.

Expected value and prediction error analyses estimated using the Rescorla-Wagner temporal difference learning model showed moderate similarity between Agent A's expected values during prediction about Agent B's choices and Agent B's expected values during her own choices. This provides some evidence for shared valuation of the fictitious prediction error for the chosen option and the fictitious indecision point (Fig. 15).

Figure 15. Shared Valuation Demonstrated by Time-Series of Individual Parameter Estimates Emerging from the Valuational Model. Histograms of 1000 randomized bootstraps with the actual r-value showing shared valuation between time series (red lines). Teacher (Ppt) Leading Learner (Lag = 1, Ppt + 1).

Discussion

Cognitive modeling supports the representation of feedback signals as a positive "stay" signal except for 5 points for all participants, independent of role of Teacher or Learner, such that both Teachers and Learners are operating under the same valuations. While Learners were expected to positively value outcomes of 10 points, which would signify that they now understand the arrangements of their probabilistic rewards and will be able to receive the optimal 100 points on the following trial, Teachers interestingly used this strategy as well. After reversal, Teachers are expected to aim to receive an optimal 50 points until the Learner signifies their understanding of the state of the world by choosing the "high" reward, resulting in the Teacher receiving 5 points. A positive value assigned to 10 points suggests that Teachers are potentially developing insights as to how many trials post-reversal it will take a Learner to understand the switch in reward contingencies. Teachers may become overhasty in an attempt to avoid receiving 5 points from when the Learner first understands the change in the reward contingencies. If Teachers are overhasty, they may have a mental model that the Learner receiving 5 points is an undeniable signal that the Learner will adjust their beliefs, just as with repeated encounters of 50 points.

However, none of the cognitive models do exceptionally well at defining narrow parameter distributions. Failure to yield well-defined parameter distributions to the reinforcement learning models implemented in this study indicates a need to consider other modeling options (ex. Drift-Diffusion Models, Working Memory Model, Interactive Partially Observable Markov Decision Processes Model (I-POMDP) (cite). In addition, the experience-weighted attraction model (EWA) and reward-punishment fictitious model could be modified to operate with data structures in multiple blocks. Further analyses will be needed to determine if greater similarity between the participants' expected values and prediction errors is associated with greater cooperative success and with self-reported characteristics such as cooperativity and perspective taking. A positive shift in indecision points for the Learner’s choice suggests that as Learners are less confident in their decision, they become more likely to choose the same choice as the previous trial.

This study was conducted with social interaction through a screen in which subjects did not actually take part in constant face-to-face social interactions with their partner. We believe that our lack of deception an close proximity of partners during the experiment controlled for this, but as there is greater activation in several social-cognitive, attention, and reward processing brain areas when participants interacted with a live experimenter, as compared to a video replays of social interaction (Redcay et al. 2010), we could see remaining artifacts in the neural data of ingenuine social interaction.

To understand what features of the model were producing the behavioral effects, we plan to examine how the best fitting cognitive model’s parameters corresponded with the EEG results, integrating the computational cognitive modeling and neural information.

Acknowledgements

I would like to express my appreciation for my advisor Prof. Michael Spezio for supporting my passion for neuroscience these past three years. Thank you to my collaborators at the Institute for Systems Neuroscience in Hamburg, specifically Dr. Jan Glascher and Tessa Rusch, who hosted me in their laboratory this past summer and to Yuqing Lei and Vanessa Hayes for their prior work in the data analysis process. I am grateful for funding support from the National Science Foundation's program in Collaborative Research in Computational Neuroscience (CRCNS; grant 1608278) and the Bundesministerium fuer Bildung und Forschung (grant 01GQ1603). I would like to thank my thesis reader, Prof. Lisette de Pillis, for her support. To my family, thank you for your encouragement and endless dedication to my education. Thank you also to my undergraduate collogues, Corinne Donnay, Naiti Bhatt, Colin Ferris, Sarah Horder, and Paige Brodrick for being full of advice, understanding, and patience. Finally, I would like to thank The Keck Science Department and Scripps College for the opportunity to complete a senior thesis, awarding me with the Claremont Colleges' Intercollegiate Neuroscience Summer Research Fellowship, and providing an institution in which I could learn from numerous transformative professors and students.

References

Adolphs, R. (2009). The Social Brain: Neural Basis of Social Knowledge. Annual Review of Psychology, 60(1), 693–716. https://doi.org/10.1146/annurev.psych.60.110707.163514

Ahn, W.-Y., Haines, N., & Zhang, L. (2017). Revealing Neurocomputational Mechanisms of Reinforcement Learning and Decision-Making With the hBayesDM Package. Computational Psychiatry, 1, 24–57. https://doi.org/10.1162/CPSY_a_00002

Anderson, J. R. (1996). The architecture of cognition. Mahwah, NJ: Lawrence Erlbaum Assos. Publ.

Asakura, N., & Inui, T. (2016). A Bayesian Framework for False Belief Reasoning in Children: A Rational Integration of Theory-Theory and Simulation Theory. Frontiers in Psychology, 7. https://doi.org/10.3389/fpsyg.2016.02019

Baron-Cohen, S. (2012). Evolution of a theory of mind? Retrieved December 4, 2018, from ResearchGate website: https://www.researchgate.net/publication/237401773_Evolution_of_a_theory_of_mind

Baron-Cohen, S., Leslie, A. M., & Frith, U. (1985). Does the autistic child have a “theory of mind” ? Cognition, 21(1), 37–46. https://doi.org/10.1016/0010-0277(85)90022-8

Barresi, J., & Moore, C. (1996). Intentional relations and social understanding. Behavioral and Brain Sciences, 19(01), 107. https://doi.org/10.1017/S0140525X00041790

Becchio, C., Manera, V., Sartori, L., Cavallo, A., & Castiello, U. (2012). Grasping intentions: from thought experiments to empirical evidence. Frontiers in Human Neuroscience, 6. https://doi.org/10.3389/fnhum.2012.00117

Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. Nature Neuroscience, 10(9), 1214–1221. https://doi.org/10.1038/nn1954

Berger, J. O. (1993). Statistical decision theory and Bayesian analysis (2nd ed). In Springer Series in Statistics (2nd ed). New York: Springer-Verlag.

Bernardo, J. M., & Smith, A. F. M. (1994). Bayesian theory. In Wiley Series in Probability and Mathematical Statistics. Chichester, Eng. ; New York: Wiley.

Bradford, E. E. F., Jentzsch, I., & Gomez, J.-C. (2015). From self to social cognition: Theory of Mind mechanisms and their relation to Executive Functioning. Cognition, 138, 21–34. https://doi.org/10.1016/j.cognition.2015.02.001

Brunner, C., Delorme, A., & Makeig, S. (2013). Eeglab – an Open Source Matlab Toolbox for Electrophysiological Research. Biomedical Engineering / Biomedizinische Technik. https://doi.org/10.1515/bmt-2013-4182

Bürkner, P.-C. (2017). brms : An R Package for Bayesian Multilevel Models Using Stan. Journal of Statistical Software, 80(1). https://doi.org/10.18637/jss.v080.i01

Carrington, S. J., & Bailey, A. J. (2009). Are there theory of mind regions in the brain? A review of the neuroimaging literature. Human Brain Mapping, 30(8), 2313–2335. https://doi.org/10.1002/hbm.20671

Cherry, E. (1953). Some experiments on the recognition of speech, with one and with two ears. Journal of the Acoustical Society of America, (25), 975–979.

Cherry, E., & Taylor, W. K. (1954). Some Further Experiments upon the Recognition of Speech , with One and with Two Ears. Journal of the Acoustical Society of America, (26), 554–559.

Dayan, P., & Abbott, L. F. (2005). Theoretical neuroscience: computational and mathematical modeling of neural systems (First paperback ed). In Computational Neuroscience (First paperback ed). Cambridge, Mass.: MIT Press.

De Dreu, C. K. W. (2013). Human Cooperation: Challenges for Science and Practice. Psychological Science in the Public Interest, 14(3), 117–118. https://doi.org/10.1177/1529100613496959

de Hollander, G., Forstmann, B. U., & Brown, S. D. (2016). Different Ways of Linking Behavioral and Neural Data via Computational Cognitive Models. Biological Psychiatry: Cognitive Neuroscience and Neuroimaging, 1(2), 101–109. https://doi.org/10.1016/j.bpsc.2015.11.004

Debes, R. (2010). Which empathy? Limitations in the mirrored "understanding" of emotion. Synthese, 175(2), 219–239. https://doi.org/10.1007/s11229-009-9499-7

Delorme, A., & Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. Journal of Neuroscience Methods, 134(1), 9–21. https://doi.org/10.1016/j.jneumeth.2003.10.009

Delorme, A., Mullen, T., Kothe, C., Akalin Acar, Z., Bigdely-Shamlo, N., Vankov, A., & Makeig, S. (2011). EEGLAB, SIFT, NFT, BCILAB, and ERICA: New Tools for Advanced EEG Processing. Computational Intelligence and Neuroscience, 2011, 1–12. https://doi.org/10.1155/2011/130714

den Ouden, H. E. M., Daw, N. D., Fernandez, G., Elshout, J. A., Rijpkema, M., Hoogman, M., … Cools, R. (2013). Dissociable Effects of Dopamine and Serotonin on Reversal Learning. Neuron, 80(4), 1090–1100. https://doi.org/10.1016/j.neuron.2013.08.030

Dodell-Feder, D., Koster-Hale, J., Bedny, M., & Saxe, R. (2011). fMRI item analysis in a theory of mind task. NeuroImage, 55(2), 705–712. https://doi.org/10.1016/j.neuroimage.2010.12.040

Dunbar, R. I. M. (1998). The social brain hypothesis. Evolutionary Anthropology: Issues, News, and Reviews, 6(5), 178–190. https://doi.org/10.1002/(SICI)1520-6505(1998)6:5<178::AID-EVAN5>3.0.CO;2-8

Ferree, T. ., Clay, M. ., & Tucker, D. . (2001). The spatial resolution of scalp EEG. Neurocomputing, 38–40, 1209–1216. https://doi.org/10.1016/S0925-2312(01)00568-9

Fox, N. A., Bakermans-Kranenburg, M. J., Yoo, K. H., Bowman, L. C., Cannon, E. N., Vanderwert, R. E., … van IJzendoorn, M. H. (2016). Assessing human mirror activity with EEG mu rhythm: A meta-analysis. Psychological Bulletin, 142(3), 291–313. https://doi.org/10.1037/bul0000031

Frank, M. J., Moustafa, A. A., Haughey, H. M., Curran, T., & Hutchison, K. E. (2007). Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning. Proceedings of the National Academy of Sciences, 104(41), 16311–16316. https://doi.org/10.1073/pnas.0706111104

Frith, C. D. (2007). The social brain? Philosophical Transactions of the Royal Society B: Biological Sciences, 362(1480), 671–678. https://doi.org/10.1098/rstb.2006.2003

Gallese, V., Fadiga, L., Fogassi, L., & Rizzolatti, G. (1996). Action recognition in the premotor cortex. Brain: A Journal of Neurology, 119 (Pt 2), 593–609.

Gläscher, J., Hampton, A. N., & O'Doherty, J. P. (2009). Determining a Role for Ventromedial Prefrontal Cortex in Encoding Action-Based Value Signals During Reward-Related Decision Making. Cerebral Cortex, 19(2), 483–495. https://doi.org/10.1093/cercor/bhn098

Gray, C. M., Maldonado, P. E., Wilson, M., & McNaughton, B. (1995). Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. Journal of Neuroscience Methods, 63(1–2), 43–54. https://doi.org/10.1016/0165-0270(95)00085-2

Hallez, H., Vanrumste, B., Grech, R., Muscat, J., De Clercq, W., Vergult, A., … Lemahieu, I. (2007). Review on solving the forward problem in EEG source analysis. Journal of NeuroEngineering and Rehabilitation, 4(1). https://doi.org/10.1186/1743-0003-4-46

Hampton, A. N., Adolphs, R., Tyszka, J. M., & O'Doherty, J. P. (2007). Contributions of the Amygdala to Reward Expectancy and Choice Signals in Human Prefrontal Cortex. Neuron, 55(4), 545–555. https://doi.org/10.1016/j.neuron.2007.07.022

Hawkins, G. E., Mittner, M., Forstmann, B. U., & Heathcote, A. (2017). On the efficiency of neurally-informed cognitive models to identify latent cognitive states. Journal of Mathematical Psychology, 76, 142–155. https://doi.org/10.1016/j.jmp.2016.06.007

Hickok, G. (2009). Eight Problems for the Mirror Neuron Theory of Action Understanding in Monkeys and Humans. Journal of Cognitive Neuroscience, 21(7), 1229–1243. https://doi.org/10.1162/jocn.2009.21189

Hillman, K. L., & Bilkey, D. K. (2012). Neural encoding of competitive effort in the anterior cingulate cortex. Nature Neuroscience, 15(9), 1290–1297. https://doi.org/10.1038/nn.3187

Holyoak, K. (1983). Toward a Unitary Theory of Mind: The Architecture of Cognition. Science, 222(4623), 499–500. https://doi.org/10.1126/science.222.4623.499

Huskey, R. (2016). Beyond Blobology: Using Psychophysiological Interaction Analyses to Investigate the Neural Basis of Human Communication Phenomena. In S. Kubitschko & A. Kaun (Eds.), Innovative Methods in Media and Communication Research (pp. 123–138). https://doi.org/10.1007/978-3-319-40700-5_7

Joliffe, I., & Morgan, B. (1992). Principal component analysis and exploratory factor analysis. Statistical Methods in Medical Research, 1(1), 69–95. https://doi.org/10.1177/096228029200100105

Jolliffe, I. T., & Cadima, J. (2016). Principal component analysis: a review and recent developments. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 374(2065), 20150202. https://doi.org/10.1098/rsta.2015.0202

Jung, T.-P., Makeig, S., McKeown, M. J., Bell, A. J., Lee, T.-W., & Sejnowski, T. J. (2001). Imaging Brain Dynamics Using Independent Component Analysis. Proceedings of the IEEE. Institute of Electrical and Electronics Engineers, 89(7), 1107–1122. https://doi.org/10.1109/5.939827

Koster-Hale, J., Richardson, H., Velez, N., Asaba, M., Young, L., & Saxe, R. (2017). Mentalizing regions represent distributed, continuous, and abstract dimensions of others' beliefs. NeuroImage, 161, 9–18. https://doi.org/10.1016/j.neuroimage.2017.08.026

Kruschke, J. K. (2013). Bayesian estimation supersedes the t test. Journal of Experimental Psychology: General, 142(2), 573–603. https://doi.org/10.1037/a0029146

Kruschke, J. K. (2015). Doing Bayesian data analysis: a tutorial with R, JAGS, and Stan (Edition 2). Boston: Academic Press.

Lau, B., & Glimcher, P. W. (2005). Dynamic Response-by-Response Models of Matching Behavior in Rhesus Monkeys. Journal of the Experimental Analysis of Behavior, 84(3), 555–579. https://doi.org/10.1901/jeab.2005.110-04

Legatt, A. D., Arezzo, J., & Vaughan, H. G. (1980). Averaged multiple unit activity as an estimate of phasic changes in local neuronal activity: effects of volume-conducted potentials. Journal of Neuroscience Methods, 2(2), 203–217. https://doi.org/10.1016/0165-0270(80)90061-8

Luck, S. J. (2005). An introduction to the event-related potential technique. In Cognitive Neuroscience. Cambridge, Mass: MIT Press.

Mithen, S. J. (1996). The prehistory of the mind: a search for the origins of art, religion, and science. London: Thames and Hudson.

Montague, P. R., King-Casas, B., & Cohen, J. D. (2006). IMAGING VALUATION MODELS IN HUMAN CHOICE. Annual Review of Neuroscience, 29(1), 417–448. https://doi.org/10.1146/annurev.neuro.29.051605.112903

Moore, C., & Paulus, M. (2013). A second-person approach cannot explain intentionality in social understanding. Behavioral and Brain Sciences, 36(04), 430–431. https://doi.org/10.1017/S0140525X12001987

Palmeri, T. J., Love, B. C., & Turner, B. M. (2017). Model-based cognitive neuroscience. Journal of Mathematical Psychology, 76, 59–64. https://doi.org/10.1016/j.jmp.2016.10.010

Poldrack, R. A. (2012). The future of fMRI in cognitive neuroscience. NeuroImage, 62(2), 1216–1220. https://doi.org/10.1016/j.neuroimage.2011.08.007

Pratte, M. S., & Tong, F. (2017). Integrating theoretical models with functional neuroimaging. Journal of Mathematical Psychology, 76, 80–93. https://doi.org/10.1016/j.jmp.2016.06.008

Rizzolatti, G., Fadiga, L., Matelli, M., Bettinardi, V., Paulesu, E., Perani, D., & Fazio, F. (1996). Localization of grasp representations in humans by PET: 1. Observation versus execution. Experimental Brain Research, 111(2), 246–252.

Rolls, E. T., McCabe, C., & Redoute, J. (2008). Expected Value, Reward Outcome, and Temporal Difference Error Representations in a Probabilistic Decision Task. Cerebral Cortex, 18(3), 652–663. https://doi.org/10.1093/cercor/bhm097

Ryynanen, O. R. M., Hyttinen, J. A. K., & Malmivuo, J. A. (2006). Effect of measurement noise and electrode density on the spatial resolution of cortical potential distribution with different resistivity values for the skull. IEEE Transactions on Biomedical Engineering, 53(9), 1851–1858. https://doi.org/10.1109/TBME.2006.873744

Schilbach, L., Timmermans, B., Reddy, V., Costall, A., Bente, G., Schlicht, T., & Vogeley, K. (2013). Toward a second-person neuroscience. Behavioral and Brain Sciences, 36(04), 393–414. https://doi.org/10.1017/S0140525X12000660

Sebanz, N., Bekkering, H., & Knoblich, G. (2006). Joint action: bodies and minds moving together. Trends in Cognitive Sciences, 10(2), 70–76. https://doi.org/10.1016/j.tics.2005.12.009

Stone, J. V. (2004). Independent component analysis: a tutorial introduction. Cambridge, Mass: MIT Press.

Sun, R. (2008). Introduction to Computational Cognitive Modeling. In R. Sun (Ed.), The Cambridge Handbook of Computational Psychology (pp. 3–20). https://doi.org/10.1017/CBO9780511816772.003

Sutton, R. S., & Barto, A. G. (1998). Reinforcement learning: an introduction. In Adaptive Computation and Machine Learning. Cambridge, Mass: MIT Press.

Thye, M. D., Ammons, C. J., Murdaugh, D. L., & Kana, R. K. (2018). Differential recruitment of theory of mind brain network across three tasks: An independent component analysis. Behavioural Brain Research, 347, 385–393. https://doi.org/10.1016/j.bbr.2018.03.041

Turner, B. M., Forstmann, B. U., Love, B. C., Palmeri, T. J., & Van Maanen, L. (2017). Approaches to analysis in model-based cognitive neuroscience. Journal of Mathematical Psychology, 76, 65–79. https://doi.org/10.1016/j.jmp.2016.01.001

van Dijk, E., Parks, C. D., & van Lange, P. A. M. (2013). Social dilemmas: The challenge of human cooperation. Organizational Behavior and Human Decision Processes, 120(2), 123–124. https://doi.org/10.1016/j.obhdp.2012.12.005

van Ravenzwaaij, D., Provost, A., & Brown, S. D. (2017). A confirmatory approach for integrating neural and behavioral data into a single model. Journal of Mathematical Psychology, 76, 131–141. https://doi.org/10.1016/j.jmp.2016.04.005

Wimmer, H. (1983). Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. Cognition, 13(1), 103–128. https://doi.org/10.1016/0010-0277(83)90004-5

Appendix A

Teacher Prediction

	Model	LOOIC
RP	1	7144.2
	2	8765.684
	3	8962.3
	4	6622.274
Fictitious	1	8597.9
	2	7679.7
	3	8896.3
	4	7112.2

Teacher Choice

	Model	LOOIC
RP	1	7928.0
	2	8355.8
	3	8528.3
	4	7692.1
Fictitious	1	7892.6
	2	7687.7
	3	7951.5
	4	7676.195

Learner Prediction

	Model	LOOIC
RP	1	4333.4
	2	3503.8
	3	3747.262
	4	3847.283
Fictitious	1	8522.001
	2	4534.744
	3	7072.413
	4	7251.7

Learner Choice

	Model	LOOIC
RP	1	7083.5
	2	7989.9
	3	8499.2
	4	7248.6
Fictitious	1	7579.5
	2	7191.0
	3	7725.7
	4	7320.2

Appendix B

Teacher Prediction

Reward-Punishment

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	0	0.0%	<NA>
(0.5, 0.7]	(ok)	6	12.0%	187
(0.7, 1]	(bad)	31	62.0%	29
(1, Inf)	(very bad)	13	26.0%	8

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	18	36.0%	1832
(0.5, 0.7]	(ok)	26	52.0%	401
(0.7, 1]	(bad)	6	12.0%	306
(1, Inf)	(very bad)	0	0.0%	<NA>

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	18	36.0%	1832
(0.5, 0.7]	(ok)	26	52.0%	401
(0.7, 1]	(bad)	6	12.0%	306
(1, Inf)	(very bad)	0	0.0%	<NA>

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	0	0.0%	<NA>
(0.5, 0.7]	(ok)	4	8.0%	279
(0.7, 1]	(bad)	42	84.0%	19
(1, Inf)	(very bad)	4	8.0%	12

Fictitious Update

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	3	6.0%	382
(0.5, 0.7]	(ok)	11	22.0%	142
(0.7, 1]	(bad)	29	58.0%	19
(1, Inf)	(very bad)	7	14.0%	14

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	28	56.0%	1813
(0.5, 0.7]	(ok)	22	44.0%	1347
(0.7, 1]	(bad)	0	0.0%	<NA>
(1, Inf)	(very bad)	0	0.0%	<NA>

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	12	24.0%	2003
(0.5, 0.7]	(ok)	26	52.0%	576
(0.7, 1]	(bad)	12	24.0%	146
(1, Inf)	(very bad)	0	0.0%	<NA>

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	0	0.0%	<NA>
(0.5, 0.7]	(ok)	14	28.0%	162
(0.7, 1]	(bad)	33	66.0%	17
(1, Inf)	(very bad)	3	6.0%	15

Teacher Choice

Reward-Punishment

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	1	2.0%	2308
(0.5, 0.7]	(ok)	19	38.0%	180
(0.7, 1]	(bad)	26	52.0%	43
(1, Inf)	(very bad)	4	8.0%	17

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	32	64.0%	4007
(0.5, 0.7]	(ok)	17	34.0%	952
(0.7, 1]	(bad)	1	2.0%	594
(1, Inf)	(very bad)	0	0.0%	<NA>

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	29	58.0%	3354
(0.5, 0.7]	(ok)	21	42.0%	361
(0.7, 1]	(bad)	0	0.0%	<NA>
(1, Inf)	(very bad)	0	0.0%	<NA>

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	1	2.0%	3818
(0.5, 0.7]	(ok)	20	40.0%	501
(0.7, 1]	(bad)	28	56.0%	48
(1, Inf)	(very bad)	1	2.0%	98

Fictitious Update

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	0	0.0%	<NA>
(0.5, 0.7]	(ok)	20	40.0%	187
(0.7, 1]	(bad)	27	54.0%	30
(1, Inf)	(very bad)	3	6.0%	9

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	5	10.0%	2312
(0.5, 0.7]	(ok)	32	64.0%	509
(0.7, 1]	(bad)	13	26.0%	50
(1, Inf)	(very bad)	0	0.0%	<NA>

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	8	16.0%	2
(0.5, 0.7]	(ok)	23	46.0%	0
(0.7, 1]	(bad)	18	36.0%	0
(1, Inf)	(very bad)	1	2.0%	0

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	0	0.0%	<NA>
(0.5, 0.7]	(ok)	21	42.0%	372
(0.7, 1]	(bad)	28	56.0%	43
(1, Inf)	(very bad)	1	2.0%	30

Learner Prediction

Reward-Punishment

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	0	0.0%	<NA>
(0.5, 0.7]	(ok)	10	20.0%	384
(0.7, 1]	(bad)	35	70.0%	31
(1, Inf)	(very bad)	5	10.0%	27

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	1	2.0%	2518
(0.5, 0.7]	(ok)	16	32.0%	432
(0.7, 1]	(bad)	31	62.0%	36
(1, Inf)	(very bad)	2	4.0%	19

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	0	0.0%	<NA>
(0.5, 0.7]	(ok)	8	16.0%	293
(0.7, 1]	(bad)	38	76.0%	40
(1, Inf)	(very bad)	4	8.0%	37

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	0	0.0%	<NA>
(0.5, 0.7]	(ok)	18	36.0%	210
(0.7, 1]	(bad)	30	60.0%	26
(1, Inf)	(very bad)	2	4.0%	19

Fictitious Update

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	0	0.0%	<NA>
(0.5, 0.7]	(ok)	0	0.0%	<NA>
(0.7, 1]	(bad)	45	90.0%	25
(1, Inf)	(very bad)	5	10.0%	6

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	0	0.0%	<NA>
(0.5, 0.7]	(ok)	2	4.0%	279
(0.7, 1]	(bad)	41	82.0%	23
(1, Inf)	(very bad)	7	14.0%	10

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	0	0.0%	<NA>
(0.5, 0.7]	(ok)	0	0.0%	<NA>
(0.7, 1]	(bad)	32	64.0%	13
(1, Inf)	(very bad)	18	36.0%	9

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	0	0.0%	<NA>
(0.5, 0.7]	(ok)	2	4.0%	248
(0.7, 1]	(bad)	35	70.0%	8
(1, Inf)	(very bad)	13	26.0%	4

Learner Choice

Reward-Punishment

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	3	6.0%	2685
(0.5, 0.7]	(ok)	32	64.0%	265
(0.7, 1]	(bad)	15	30.0%	41
(1, Inf)	(very bad)	0	0.0%	<NA>

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	22	44.0%	4255
(0.5, 0.7]	(ok)	25	50.0%	221
(0.7, 1]	(bad)	3	6.0%	1472
(1, Inf)	(very bad)	0	0.0%	<NA>

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	45	90.0%	5442
(0.5, 0.7]	(ok)	5	10.0%	2704
(0.7, 1]	(bad)	0	0.0%	<NA>
(1, Inf)	(very bad)	0	0.0%	<NA>

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	2	4.0%	3542
(0.5, 0.7]	(ok)	30	60.0%	613
(0.7, 1]	(bad)	18	36.0%	26
(1, Inf)	(very bad)	0	0.0%	<NA>

Fictitious Update

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	0	0.0%	<NA>
(0.5, 0.7]	(ok)	21	42.0%	254
(0.7, 1]	(bad)	27	54.0%	36
(1, Inf)	(very bad)	2	4.0%	39

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	35	70.0%	2494
(0.5, 0.7]	(ok)	15	30.0%	483
(0.7, 1]	(bad)	0	0.0%	<NA>
(1, Inf)	(very bad)	0	0.0%	<NA>

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	28	56.0%	3899
(0.5, 0.7]	(ok)	20	40.0%	2266
(0.7, 1]	(bad)	2	4.0%	420
(1, Inf)	(very bad)	0	0.0%	<NA>

		Count	Pct.	Min. n_eff
(-Inf, 0.5]	(good)	5	10.0%	1416
(0.5, 0.7]	(ok)	33	66.0%	216
(0.7, 1]	(bad)	11	22.0%	135
(1, Inf)	(very bad)	1	2.0%	19