1. Title:

Classification of Theory of Mind tasks and their computational models

2. Authors:

Tessa Rusch¹ & Jan Gläscher¹

¹ Institute of Systems Neuroscience, University Medical Center Hamburg-Eppendorf (UKE), Martinistr. 52, D-20251 Hamburg

3. Keywords

Theory of Mind, computational modeling, social decision making

4. Abstract

The ability to form a Theory of Mind (ToM) that is to theorize about others mental states and explain and predict behavior in relation to attributed intentional states, constitutes a hallmark of human cognition. These abilities are multi-faceted and include a variety of different cognitive sub-functions. Here, we focus on decision processes in social contexts and review a number of experimental and computational modelling approaches in this field. We classify experimental accounts and formal computational models with respect to two categories: ***interactivity*** and ***uncertainty***. Thereby, we aim at capturing the most relevant variations in ToM-related cognitive sub-functions and highlight the nuances of what broadly is referred to as social decisions processes. Finally, we discuss the most typical neuro-imaging findings in light of the categorization results.

5. Introduction

Humans are distinctly skilled at sophisticated social interactions. To successfully engage in social exchanges they rely on Theory of Mind (ToM): An individual's ability to attribute mental states such as beliefs, desires and intentions to others and predict and explain behavior in light of these mental states. The term *Theory of Mind* was first coined in the highly influential article "Does the Chimpanzee have a theory of mind?" by Premack and Woodruff (1978) and referred to "an individual imputing mental states to himself and others". Along with this definition a first task to probe this capacity in non-human primates was presented: To determine a chimpanzee's ToM abilities short videos of a human struggling with simple tasks were played to the animal. Subsequently, photographs of various items, including the one solving the actor's problem, were presented. The chimpanzee's ability to select the correct photograph was interpreted as it being capable of recognizing the problem, i.e. representing the state of affairs, as well as the actor's purpose in the scene, i.e. his intentions and goals. This highlighted that having a Theory of Mind requires a representation of the state of affairs *and* representing an individual's relationship to it, i.e. the individual's desires and beliefs in the situation (Wimmer & Perner, 1983).

Since this initial ToM task, countless experimental approaches probing and characterizing ToM capacities have been introduced by psychological and behavioral economics research (Houser & McCabe, 2014; Schurz, Radua, Aichhorn, Richlan, & Perner, 2014). Paired with neuroscientific methods neural networks relevant for ToM were successfully identified (Gallagher & Frith, 2003; Schurz et al., 2014; Siegal & Varley, 2002). Further, quantitative, testable descriptions of the covert cognitive processes constituting ToM were obtained by deploying computational models to explain behavior in these tasks and inform the analyses of neural signals (Charpentier & O'Doherty, 2018; Hampton, Bossaerts, & O'Doherty, 2008a; Hill et al., 2017; Xiao,

Geng, Riggins, Chen, & Redcay, 2019; W. Yoshida, Seymour, Friston, & Dolan, 2010). Despite the vast success of these approaches, a coherent picture of what Theory of Mind is, how humans (and other species) do it, and which neural mechanisms constitute it, is missing (Schaafsma, Pfaff, Spunt, & Adolphs, 2015). We argue, that in part this is due to the fact that different ToM tasks elicit largely different cognitive sub-functions.

One classic high level conceptual division of ToM is that between cognitive and affective processes (Baron-Cohen, 1988; Kalbe et al., 2010; Mitchell & Phillips, 2015). Cognitive ToM is primarily focused on perspective-taking and (strategic) reasoning about another person's beliefs and intention, thus allowing for predictions of the other's behavior. Affective ToM refers to cognitive processes regarding emotions of others such as empathy, emotion recognition and emotion simulation. Both cognitive and affective ToM processes have been investigated in great detail. Examining both lines of research at once would go beyond the scope of a single article. Therefore, in this current review, we exclusively focus on cognitive ToM processes. We examine how cognitive ToM tasks differ with respect to cognitive functions they elicit and classify experimental approaches and cognitive computational models along two overarching dimensions: ***interactivity*** and ***uncertainty***. We propose this categorization can provide a renewed view onto what constitutes ToM, how to interpret existing findings on the behavioral as well as the neural level, and what to consider when designing future studies.

First, we summarize a range of relevant ToM tasks from psychology, economics and decision neuroscience and categorize the different experimental approaches based on two features that determine the cognitive demands they place on participants as well as the type of reasoning they elicit: The first feature, ***interactivity*** denotes whether participants truly interact with or merely observe other individuals. Second,

tasks are classified with respect to the ***uncertainty*** they included. We discriminate between (1) ***environmental uncertainties***, i.e. is the social scene taking place in a stable and fully observable environment or are agents faced with uncertainty about the state of affairs, and (2) ***social uncertainties***, i.e. are agents unsure about others' internal variables as, e.g. subjective preferences, personality traits, etc. In the second step, we inspect different computational models for quantifying the cognitive processes individuals engage in when solving such tasks and categorize those according to the same basic dimensions as tasks. Lastly, we interpret neural findings with respect to these categories.

Figure 1| Observational ToM tasks comprising environmental uncertainty.

(A) In the false belief task a simple social scenario is presented to a participant. After observing a social scene that comprises a change in the physical environment that the observed agent is unaware of (inducing a false belief), participants have to predict that observed agent's behavior. To successfully do so, participants have to abstract from their own representation of the surrounding and adopt the observed agents (physically inaccurate) perspective. (B) The trajectory of an agent starting from the bottom right corner of a simple maze-like environment is present to participants. The observed agent's perceptual abilities are limited by occluding walls preventing him from overseeing the entire scene. Based on the path that the agent takes (here indicated by arrows) participants are asked to indicate the agent's subjective preferences over available goal states (here: purple, green and orange). (C) In a group decision game, participants need to learn the value of two dynamically changing probabilistically rewarded choice options. After making their own choice, co-players' selections who learn about the same choice options are revealed and participants are allowed to adjust their choice. Finally, feedback about the reward associated with the chosen option is presented. (D) Participants observe four different co-players with varying expertise in probabilistic value learning task. First, they choose between betting for or against these agents' success. Second they see the observed agent's choice (his prediction about whether the presented asset would increase or decrease in value). Last, they receive feedback about whether their bet was correct or not by either winning or losing money, allowing inference about the others' expertise and the value of assets. (E) Participants observe the actions of different agents whose preferences they learned in a pre-test training period. The outcome, i.e. the reward, associated with those actions is not revealed to the observer. To infer the underlying reward distribution, participants need to represent the observed agents learning processes and interpret their actions light of their preferences.

6. Tasks

Observation under environmental uncertainty

The probably most prominent task in ToM research is the so-called *false belief task* (Figure 1A). First formulated by Wimmer and Perner (1983) it consists of a short story

including the character Maxi who puts a chocolate bar on a shelf and leaves the scene. In Maxi's absence, his mother changes the state of affairs by moving the chocolate to a different location. Upon Maxi's return, the observing participant is asked where Maxi would search for his chocolate. The key feature in this task is the change in the state of affairs and Maxi's ignorance of that change, i.e. Maxi's false belief. To correctly predict Maxi's behavior observers need to abstract from their own (correct) knowledge about the situation, represent Maxi's (false) belief and respond based on their representation of Maxi's mental state. Countless variants of the false belief task have been deployed to assess the development of ToM abilities in children and impairments in individuals with autism spectrum disorder (ASD) or psychiatric populations (e.g. Baillargeon, Scott, & He, 2010; Baron-Cohen, Leslie, & Frith, 1985; Bora, Yucel, & Pantelis, 2009; Call & Tomasello, 1999; Dufour et al., 2013; Saxe & Kanwisher, 2003; Wimmer & Perner, 1983). Common to all variations is the use of short stories that require judgment about a false belief scenario. Yet, depending on task specifics, findings about when healthy children develop the ability to theorize about other minds differ: When explicitly asked children answer questions about an agent's false belief correctly from around four years on (Wimmer & Perner, 1983). However, already 13-month-old infants show correct anticipatory viewing behavior in such tasks (Surian, Caldi, & Sperber, 2007) pointing towards a much earlier onset of false belief understanding (Baillargeon et al., 2010). Children with ASD fail to show either, explicit false belief reasoning and anticipatory action (Baron-Cohen, 2000; Baron-Cohen et al., 1985; Senju et al., 2010). Further, they also perform poorly in so called level-2 visual perspective taking tasks (Hamilton, Brindley, & Frith, 2009). These consist of situations that require an understanding that an object simultaneously observed by two individuals may appear differently depending on the persons' perspectives. Similar to

false belief tasks, perspective taking requires putting oneself into some else's shoes and seeing the world from their eyes.

Following a similar approach as visual perspective taking tasks, Baker and colleagues (2017) used maze-like spatial layouts, an environment well suited for the application of formal decision models, to examine inferential processes about an observed individual's beliefs and preferences: An agent with unknown preferences is placed in an environment containing different choice options with varying subjective value to the agent (Figure 1B). At any given trial, only a subset of options is available in the environment. Additionally, occluding walls prevent the agent from overseeing the entire space, i.e. he has to move around to explore what options are currently available to choose the option that is most valuable to him. Participants observe the agent while taking a bird's eye perspective, i.e. as in the false belief task they are fully informed about the environmental properties, but the observed agent's preferences over choice options are unknown to them. In Figure 1B an exemplary situation with two out of three possible choice options (purple, green and orange) of varying subjective value are shown. From his initial position the observed agent can only see the orange option. He first moves around the occluding wall but then turns around and returns to the orange option. When asked to rate the agent's preferences based in this behavior, participants indicate that green is most valuable to the agent followed by orange and rate purple as least valuable. These judgments indicate that participants infer the agent's belief based on his perceptual experience, i.e. which choice options he can and cannot see, and attribute preferences by explaining behavior in light of these beliefs: "initially, he only sees orange but moves past it, hence orange is not the most valuable option. When finding out that the alternative is purple, he turns around. This means orange is more valuable than purple. Consequently, green is most valuable to him."

In both, the false belief and observational perspective taking task, fully informed participants watch agents acting in an uncertain environment. With respect to *interactivity* this means participants take a purely observational perspective. They are completely detached from the scenario and their judgements and predictions are entirely irrelevant to the characters in and the progress of the scene. Both tasks comprise *environmental uncertainty*, but only the observed agent is faced with uncertainty while participants take an all-knowing perspective. Thereby, a divergence between participants' own representation of the world and the observed agent's belief about the surrounding is induced. This mismatch is the key factor for triggering reasoning about another person's belief: In situations of identical own and other beliefs there is no reason to take the other's perspective as it provides not additional information. This means, to experimentally investigate the attribution of beliefs, a difference between own and other beliefs has to be created through environmental uncertainty and unequal distribution of information among individuals. If this prerequisite is not met, discrimination between own models of the world and models of other's mental states is impossible (or at least extremely difficult). In addition to *environmental uncertainty*, in the observational perspective taking task by Baker et al. (2017) participants also deal with *social uncertainty*: The observed agent's preferences are unknown, hence need to be inferred from overt behavior. Such a scenario requires two inference processes. For one, based on the information available to the observer the observed agent's beliefs needs to be tracked. Second, building on beliefs and observed behavior, underlying individual preferences need to be inferred.

Although not yet interactive, observational reinforcement learning paradigms require active engagement. Experimental setups examining observational learning deploy classic single-agent reinforcement learning problems such as probabilistic reversal learning or multi-armed bandit tasks: In these tasks, the goal is to choose

optimally among multiple competing choice options. However, participants receive only probabilistic information about choice-outcome associations, hence choosing optimally requires dynamic learning about the environment. During observational learning, instead of choosing actions and receiving outcomes themselves, the participants observe other individuals selecting between available options and obtain rewards (Hill, Boorman, & Fried, 2016; Selbing & Olsson, 2017). From observed choice-outcome associations, participants vicariously learn the underlying reward distributions. In such settings, the observed agent's choices determine the information the observer receives about the environment. Apart from that, the observed agent's actions and consequently his representation of the learning problem are irrelevant to the observer.

Extending vicarious learning into a more immersive setting Zhang & Gläscher (2019) examined the effect of observing multiple other learners while engaging in a probabilistic reversal learning task oneself: Participants choose between two choice alternatives, one associated with a high probability of winning the other with a high probability of losing. Reward contingencies switched are a variable number of trials. After making their own choice, the co-players' choices are revealed and participants get the opportunity to adapt their choice. Thereby, the task allows for learning from own and others' experience. (Figure 1C). Zhang & Gläscher found that the stronger the social information diverged from participants' own choice, i.e. the more co-players chose the opposite option to themselves, the more frequent participants switched their choice to go with the group.

In an engaging observational learning paradigm Boorman and colleagues (2013) rendered individual observed agents' learning success and consequently their mental states relevant to observing participants by introducing bets for or against different observed agent. Similar to a stock market scenario these agents had to predict different assets' changes in value. Participants had to bet on the correctness of four

individuals they observed while completing the task with varying success. Figure 1C shows an exemplary situation with four agents (green, purple, blue and orange) with different underlying fixed success rates (0.3, 0.4, 0.6 or 0.7) that are unknown to participants. In addition, participants themselves had to occasionally predict the assets' value development. To perform well in this task, participants needed to track the properties of a volatile environment (i.e. the value of assets) as well as the others' expertise in the task, i.e. the overall correctness of other's representation or belief about the current state of affairs and the correctness of their beliefs in light of one's own beliefs.

A second approach to motivate tracking another's learning process is hiding the observed agent's action outcomes from the observer while his preferences are known (Collette, Pauli, Bossaerts, & O'Doherty, 2017) (Figure 1D). In this case, instead of directly learning about reward structure from the action-outcome-association, values are inversely inferred from the actions and preferences of the observed agent.

Observational learning paradigms require participants' active engagement in a social scene but no direct interaction with other individuals yielding "partial" *interactivity*. Further, all approaches face participants with a high *environmental uncertainty*. In contrast to false belief tasks, to act successfully they themselves have to learn about the current state of affairs. Additionally, to varying degree tasks require tracing observed agents' beliefs about the surrounding, ranging from tracking the congruence between social information and own beliefs in the social influence and the inverse learning task (Collette et al., 2017; Zhang & Glascher, 2019) to tracking individual observed agents' representation of the environment (Boorman et al., 2013). None of the considered observational learning paradigms include social uncertainty as others' preferences are known or learned beforehand.

Interaction under full environmental certainty

When social scenarios expand from mere observation to direct interaction, interacting individuals' behavior and consequently their thought processes become interdependent. This means in situations in which an agent *A*'s actions are relevant to a second agent *B* while also *B*'s actions are relevant to *A* their reasoning processes may become recursive (in the sense of "A thinks that B thinks, that A thinks that B will do XYZ.") (Camerer, Ho, & Chong, 2004). To examine reasoning processes of this kind, behavioral game theory uses a range of simple yet very powerful interactive games. A game in this sense is a multi-agent decision situation where the actions of all participating agents affect each other, i.e. each individual's success (usually defined as maximizing the individual reward) depends on others' choices. Assuming complete and optimal rationality of all interacting agents (Gibbons, 1992), the field was initially concerned with computing optimal solutions to such games, so-called equilibrium states, from which to deviate would be detrimental for all agents. However, more recently, experimental economics and behavioral game theory focused on examining which factors affect human (often non-optimal) social decisions, driving optimal rationality into the background (Camerer, 2003).

Countless studies deploying a multitude of strategic games, examining the effects of instructions, framing, incentives, and many other highly relevant factors driving human interactive behavior exist (Houser & McCabe, 2014). Here, we are neither able to list all existing games nor summarize the effect of the different factors affecting strategic behavior listed above. We merely pick a sub-set of the well-known economic tasks and discuss them with respect to our two defining features: *interactivity* and *uncertainty*. Further, although classic economic research often tests strategic behavior in one-shot scenarios, we only consider recurring interactions as in everyday life humans tend to interact with the same individuals more than once. Repeated

Figure 2| Interactive tasks in stable and fully observable environments.

(A) The beauty contest game nicely illustrates the concept of recursive reasoning. The goal is to choose a number between 0 and 100 closest to $\frac{2}{3}$ of the average of number chosen by all other participants. Depending on the level of reasoning different values are chosen: level 1 ≈ 33 , level 2 ≈ 22 , and so forth.

(B) In the public goods game, a group of players is endowed with an initial amount. They can choose to secretly invest as much as wanted into a public fund which is then multiplied by a fixed factor and equal splits are returned to all players irrespective of their initial investment amount.

(C) (D) (E) Matrix Games, such as matching pennies, prisoner's dilemma and stag hunt are defined by a payout matrix that based on the two players' actions determines both players' rewards. Depending on the configuration of the matrix a competitive, cooperative or coordination scenario is created.

(F) The spatial stag hunt or grid game adds a spatial component to games defined by simple payout structures as the previous examples. To successfully coordinate in the spatial stag hunt, players need to take the path their co-player is taking through this environment into account, hence inference about the co-player's future actions is added to the decision process.

(G) In the trust game one player acts as the proposer, the second as the receiver. Actions are taken sequentially. The proposer can decide how much of an initial endowment to invest. The investment is multiplied by a known factor. The receiver can now decide how much of the multiplied investment to return to the proposer.

(H) The patent race game comprises two players, a rich and a poor. Both players can choose to bid as much of their capital for a price. The higher bid earns the price, both players' bid is lost.

(I) In the centipede game two players sequentially choose between "take" or "pass". When a player takes unequal rewards are distributed to both players, the taking player receives more. Importantly, with each move rewards increase but that player whose turn it is reward is always greater than the waiting player's.

interactions allow for sophisticated predictions about the other's behavior based on structures in their strategies or on built-up expectations about personalities, such as (mis-) trustworthiness.

The beauty contest game (Figure 2A) illustrates the concept of recursive reasoning in more detail: A group of individuals anonymously chooses between a number between 0 and 100. The goal in the task is to choose the number closest to $\frac{2}{3}$ of the average of the numbers chosen by all players. According to the cognitive hierarchy theory (Camerer, Ho, & Chong, 2003), people engage in reasoning processes of varying levels of sophistication to solve this problem: A very basic player, so called level 0, randomly chooses a number. A more sophisticated level 1 player assumes all others to act at level 0, resulting in an average of 50, and chooses 33 ($\frac{2}{3}$ of 50). At level 2 others are considered to be level 1 yielding an optimal response of 22 ($\frac{2}{3}$ of 33), and so on. Such recursive reasoning could in principle extend *ad infinitum* leading to the optimal equilibrium of 0. Nevertheless, On average, people select numbers between 25 and 40 suggesting a reasoning level of 1 or 2, but the variance in choices is large and groups of varying analytical training (e.g. economics PhDs compared to high school students) show highly different means (Camerer, 2003; Camerer, Ho, & Chong, 2015).

The public goods game (Figure 2B) deploys anonymous group decisions to study fairness and reciprocity in a situation where individual interests conflict with the group's welfare (Houser & McCabe, 2014; Khalvati et al., 2018): All members of an anonymous group are endowed with the same initial amount. They may choose how much of their individual money to secretly invest in a public fund. The money in this fund is then multiplied by a known factor and equal shares are distributed to all players irrespective of their endowment. Non-invested money is kept by the respective players. This means players can choose to cooperate by investing into the public fund but they can also "free-ride" by taking their share of the others' investment without investing themselves, maximizing only their individual reward. However, free-riding will

potentially raise mistrust in the group resulting in reduced group investments and eventually reduced reward for free-rides. Consequently, to maximize outcomes players need to consider their own actions' effect on the group's behavior. Overall, people invest about half of the initial endowment although the cooperation rate drops over repeated interactions (Ledyard, 1994).

Beauty contest and public good games immerses participants in an interactive scenario, where actions of all involved individuals directly affect each other. However, as the interaction takes place at the anonymous group level, recursive representations of other individuals and their reasoning processes is neither necessary nor possible. Instead, time-varying average group level variables, in the beauty contest the group's level of sophistication and in public goods games the rate of free-riding, need to be inferred. This requires observing and learning about these variables as interactions continue, i.e. resolving *social uncertainty*. Second, instead of explicit recursive reasoning about other individuals' cognitive processes the effect of own actions on the group behavior needs to be computed via recursive group reasoning. Neither game includes *environmental uncertainty* hence no learning about the state of affairs is required.

Dyadic games (i.e. those involving two players) allow for more direct interaction than anonymous group tasks. First, we consider strategic two-player games entirely defined by a single payoff matrix: Conditioned on both players' actions a payoff matrix alone determines individuals' rewards. Well known examples comprise "matching pennies" (Mookherjee & Sopher, 1994; Touhey, 1974) (as variant also known as "hide and seek", "inspection game", or "rock, paper, scissors"), "prisoner's dilemma" (Axelrod & Hamilton, 1981), and "stag hunt" (Rousseau & Cranson, 1984; Skyrms, 2001) (Figure 2C to E). Matching pennies is a zero-sum game fully defined by a competitive payoff structure, i.e. one player's win determines the opponent's loss and vice versa (Figure

2C). Stag hunt constitutes a cooperative game which requires coordination between partners to obtain the highest possible rewards: Jointly going for the large reward (i.e. both hunting the *stag* instead of going for smaller individual reward, the *hare*) yields the maximal payoff for both players. Thus, coordination is easily achieved (Figure 2E). The famous prisoner's dilemma lies somewhere in between competition and coordination. When actions are uncoordinated players' payoffs are asymmetric and the dilemma arises as one player's optimal outcome co-occurs with the second player's maximal loss (Figure 2D). Coordination averts one's big loss and leads to a symmetric but sub-optimal outcome for both players.

Extending the very basic decision processes in simple matrix games, spatial variants, also referred to as grid games, have been created by adding two-dimensional grid worlds to the basic payoff structure of fully cooperative and social dilemma type of games (Kleiman-Weiner, Ho, Austerweil, Littmann, & Tenenbaum, 2016; Shum, Kleiman-Weiner, Littman, & Tenenbaum, 2019; Yoshida, Dolan, & Friston, 2008). In these grid worlds, the underlying payoff structure of a game is contained, e.g. in a spatial stag hunt still jointly catching the stag yields the highest reward (Figure 2E) but it requires coordinated long-term action planning (Figure 2F). The rate of cooperation in grid games depends on the underlying payoff structure, i.e. high cooperation rate in pure coordination games and lower cooperation rate in dilemma type of scenarios (Kleiman-Weiner et al., 2016), and the partner's strategy, i.e. more cooperation with partners of higher sophistication level (Yoshida, Seymour, Friston, & Dolan, 2010; Yoshida et al., 2008).

The last group of interactive games considered here, including the "trust", "patent race" and "centipede" game, consists of simple bargaining environments (Sanfey, 2007) (Figure 2G to I). In the trust game (Berg, Dickhaut, & McCabe, 1995) one of two players, the proposer, is endowed with a certain amount of money (Figure

2G). The receiver's share is multiplied with a fixed factor. The receiver can now decide how much of the multiplied amount to return to the initial proposer. When the proposer invests a lot and the receiver returns a fair share of the multiplied investment both players mutually benefit. However, if the receiver does not return at least the invested amount, the proposer loses and likely starts to mistrust the co-player leading to lower investments and consequently lowery payoffs for both players. In a multi-round trust game players commonly follow a tit-for-tat strategy, i.e. they cooperate when the co-player cooperated in the previous round and likewise reciprocate non-cooperative behavior, but both, proposers' investments and receivers' returns slightly reduce over time, i.e. cooperation decreases (King-Casas et al., 2008, 2005). However, receivers show "coaxing" behavior when proposers investments become really low: They return a larger share of the multiplied investment to reassure the proposer of their trustworthiness (King-Casas et al., 2008). These findings indicate that behavior in the trust game relies on reasoning processes on how actions affect a co-player's impression of own trustworthiness and cooperativeness.

In the centipede game (Rosenthal, 1981) (Figure 2I) two players take turns at choosing between taking a pot of money or passing it on to the co-player. When passing on, and the co-player takes in the next round, the players outcomes are slightly lower than as when he would have kept it. However, after each round of passing by both players, both player's outcomes increase. As the trust task, the centipede game requires considerations on how own behavior will affect the co-player's future actions (e.g. rewarding/punishing of un-/fair or mis-/trusting behavior). Additionally, they might assume the other's reasoning about themselves and so on. For instance, a fully rational and optimal player solves the centipede game by *backward induction*: "in order to receive the maximum reward at the end the opponent has to see the advantage of passing, which implies that I have to pass before etc.". Such reasoning quickly reveals

that the optimal solution is to take the money in the first step thus insuring at least a minimal reward. However, humans often advance the game to a later stage (Hedden & Zhang, 2002; McKelvey & Palfrey, 1992) potentially due to their limited capacity for complete backward reasoning all the way through (Ho & Su, 2013) or because they recognize the mutual long term benefit of reciprocating altruistic behavior.

In the patent race game (Dasgupta & Stiglitz, 1980; Loury, 1979), two players, competing for a prize, receive initial endowments (Figure 2H). However, one player receives more making him “richer” than the opponent. Both players simultaneously bid for the prize. The one offering more wins the prize but loses his investment, the second player loses all, his investment and the prize. To maximize their returns, players need to invest as little possible but as much as necessary to overrule the opponent. Based on their opponent’s respective choice history participants can make predictions on their opponent’s next offering. Similar to the thought processes describe for the centipede and game, in the patent race game players might assume their partner is predicting themselves, triggering recursive reasoning processes.

During repeated dyadic tasks participants fully immerse in interaction with the other player. They can build sophisticated recursive models of their partners or opponents to predict future behavior from forgone actions. These include hidden personal variables such as their trustworthiness, fairness or selfishness, but also their sophistication, i.e. the depth, in recursive reasoning processes, which can be influenced and guided by own behavior. However, these tasks take place in very simple fully defined surroundings. There is no uncertainty with respect to others’ beliefs about the world. This means, game theoretic tasks create fully interactive situations that include a large variety of social uncertainty but no environmental uncertainty. Consequently, representing other’s beliefs about the state of the world is unnecessary but representations of the other’s beliefs about one’s own actions (i.e. a recursive

mental model of the other player) can be advantageous in predicting the other player's actions and adapt the own strategies to them.

Interaction under environmental uncertainty

Very few experimental approaches combine environmental uncertainty and interactivity in a single setting. One example is the advisor task (Behrens, Hunt, Woolrich, & Rushworth, 2008; Diaconescu et al., 2014) (Figure 3A): Two individuals, an advisor and a player, interact in a gambling task while the advisor is provided with more accurate information on outcome probabilities. He can send simple cues to the other player (the "advisee") on which action to take. After a correct choice the advisee moves forward on a progress bar and receives a small reward while moving backwards and losing a small amount following a wrong choice. Both participants have opposing goals: for each of them different regions in the player's progress bar are associated with additional bonuses of \$10 or \$20. These regions are fully disclosed to the advisor, but the advisee only knows in which region he himself receives bonuses, while being ignorant to the advisor's bonus distributions. This element of the task is reminiscent of the original false belief task, in which the testee has a more comprehensive knowledge

Figure 3| Interactive decision tasks under uncertainty

(A) In the advisor task one player takes to role of the advisor, the second of the player. The player has to make choices on a probabilistic lottery with high uncertainty. Correct choices move him forward on a progress bar, incorrect choices backwards. Additionally, when finishing in two defined regions additional bonuses of +10 or +20 are obtained. The advisor receives more accurate information about the outcome of the lottery and can choose to send advice to the player. However, the advisor's bonus regions are defined in competition to the player's, meaning that the advisor receives a bonus when the player finishes in a certain region that is different from the player's bonus region. This creates competing interests in the two players. **(B)** The tiger task simulates a simple game show scenario: Two players are faced with two doors. One hides a pot of gold (positive reward +10) the other a tiger (large negative punishment -100). They have to choose between three choice options: listen (providing probabilistic information about the tiger's location at a small cost of -1), open the left door, or open the right door. Depending on the reward configuration (cooperative -left- or competitive -right-) co-players have to coordinate or precede the opponent's action to maximize individual rewards.

about the world than the characters in the vignette. The competitive bonus in the advisor task incentivizes the advisor to provide false advice during certain phases of

the experiment, and veridical advices during others. The advisee knows about conflicting goals of the advisor and therefore one of his goals is to estimate the veracity of the advisor's signals. The advisor task includes a high level of environmental uncertainty, both players only receive probabilistic information about reward distributions, while at the same time being fully interactive. The majority of advisers show variable strategies over time: while trying to be helpful at the beginning they switch to misleading advice when reaching their individual goal. Advisees appear to detect this and follow advice in the beginning but less towards the end (Diaconescu et al., 2014).

A second powerful experimental setting combining interactivity and uncertainty is the multi-agent tiger task (Doshi & Gmytrasiewicz, 2004; Kumar, Rusch, Doshi, Spezio, & Gläscher, 2019). In a game show-like scenario two players have to learn which of two doors hides a pot of gold (reward of +10) and which hides a dangerous tiger (punishment -100). They can choose to open one of the doors or sample probabilistic information about the location of the tiger at a small cost (-1). If a door is opened the tiger is randomly reset and the game starts anew. A superimposed reward structure incentivizes players to cooperate or compete: In the cooperative setting, after each action players receive half of the partner's outcome while in the competitive scenario half of their partner's outcome is subtracted from their own outcome. This means, under competitive conditions, an opponent's loss results in an own win, under cooperative conditions the partner's losing results in an own loss. In the competitive setting players race to open the correct door before the opponent does but still need to sample enough information to avoid the tiger. The cooperative scenario incentivizes coordinated behavior over individual learning about the reward distribution. This results in less sampling of information in the competitive setting leading to riskier choices but also the potential to beat the opponent to the gold (Kumar et al., 2019).

The advisor and the multi-agent tiger game face participants with high environmental uncertainty. All participants receive only probabilistic information about state of the environment. This means they themselves form uncertain beliefs while they can also only estimate the other's belief about the state of affairs. In the advisor task, advisors' cues further add social uncertainty into the game. Advisors need to predict the effect that their cues exert on the advisee; the advisee on the other hand needs to infer the advisor's intent. In contradistinction, incentives are clearly set in the tiger game. However, individuals may differ with respect to risk. To successfully cooperate or compete players need to estimate and react to co-player's individual level of risk-seeking behavior. Due to full interactivity in both tasks thinking about the other's beliefs as well as their individual preferences or intentions potentially result in recursive reasoning processes.

7. Models

Figure 4| Localizing tasks and models with respect to our 2 dimensional classification space

We collapse physical and social uncertainty to visualize experimental tasks and intentionality and recursivity to localize computational models with respect to interactivity and uncertainty.

To specify the hidden cognitive processes underlying overt behavior cognitive neuroscience leverages quantitative computational models that are fit to the participants' behavioral responses. From these models latent variables capturing the significant computations are extracted and combined with neural data in model-based analyses approaches to examine how particular cognitive operations are carried out on the neural level (Gläscher & O'Doherty, 2010). We review a number of computational models used to elucidate social decision-making processes. As for experimental approaches, we classify models with respect to **interactivity** and **uncertainty**. **Interactivity** in the context of formal models will denote the extent to which models capture the other agent and will be sub-divided into two additional sub-categories: (1) **intentionality**, i.e. models that include intentional model of others in contrast to models that include the effects of others' actions as regularities in the environment, and (2) **recursivity**, i.e. can models capture process such as "Agent A thinks that agent B thinks that A thinks that XYZ is the case". The second category,

uncertainty, refers to models' capacities to capture the underlying but sometimes only partially accessible states of the environment.

Non-interactive models for decision making under uncertainty

The formal social decision frameworks included here are based on single-agent reinforcement learning (RL) models. RL problems are described by Markov Decision Process (MDPs). An MDP is defined by set of states $S = \{s^1, \dots, s^k, \dots, s^n\}$ defining the environment, a set of actions $A = \{a^1, \dots, a^k, \dots, a^n\}$ an agent can take, a reward function determining the reward based on states and actions $R(s_{t-1}, a_{t-1}, s_t)$, and a transition function $T = p(s_t | s_{t-1}, a_{t-1})$ determining the environmental dynamics, i.e. the transition probabilities between states given actions. Different decision strategies in the different states of the environment provides varying rewards to the agent. The goal in RL is to take those actions maximizing future rewards (Sutton & Barto, 2012). This can be achieved via choice-heuristics and learning the value of chosen (and sometimes unchosen) actions without explicitly representing the structure of the surrounding by mapping actions to rewards. Alternatively, agents might form a sophisticated representation of the environment, which means representing the transition probabilities between states, allowing for flexible goal directed decision making. The two decision types are widely referred to as *model-free* and *model-based* respectively and evidence for both decision systems has been found in humans and other organisms (Daw & Dayan, 2014; Daw, Niv, & Dayan, 2005; J. Gläscher, Daw, Dayan, & O'Doherty, 2010; Wan Lee, Shimojo, & O'Doherty, 2014).

Model-free learning and decision making can be captured by temporal difference (TD) learning (Sutton & Barto, 2012). In TD an agent learns solely by experience without representing the dynamics of the environment. At each time step t the value

$V(s_t)$ of taking a strategy in state $s_t \in S$ of the environment is updated based on the obtained reward:

$$V(s_t) = V(s_t) + \alpha[r_{t+1} + \gamma V(s_{t+1}) - V(s_t)]$$

with α being a learning rate that weights the influence of new observations and γ being the discount factor regulating the effect of future values. Although lacking an explicit representation of the surrounding, dynamic value updating allows adapting to an uncertain and changing environment, i.e. provides a basic representation of environmental uncertainty.

Basic model-free RL algorithms can capture vicarious learning mechanisms elicited in observational learning scenarios in which individuals do not act themselves, but learn by observing the actions as well as the outcome of others' actions (Burke, Tobler, Baddeley, & Schultz, 2010; M. R. Hill et al., 2016). More complex observational learning problems require extensions of classic single-agent RL models. To formalize indirect learning about the surrounding by observing others whose preferences are known to the observer but action-outcomes are hidden (Figure 2E) Collette and colleagues (2017) used an inverse RL framework. Inverse RL distinguishes from classic RL in that instead of trying to find an optimal strategy with respect to a given reward function, it aims at inferring the reward function that explains an agent's behavior best (Ng & Russell, 2000). Collette et al (2017) used inverse RL to capture how humans make sense of the world by observing other agents' behavior: Their computational model recovered the underlying reward distribution for which the observed agents' actions were optimal thereby capturing participants' learning about the environment by observing behavior.

A different adaptation of single-agent RL was presented by Boorman and colleagues (2013) to model the tracking of others' expertise (Figure 2D). They combined classic learning about the values of strategies with learning about the quality

of others' behavior (i.e. the correctness of others' actions). Observed agents' expertise, i.e. the probability of them making correct choices, was assessed in two sequential learning steps: First, the match between observed choice and own action-value estimate was assessed, second, the expertise estimate was updated based on the correctness of the observed action.

A third variation of classic RL was presented in a recent attempt to model social influence of group decisions on individual decision by Zhang & Gläscher (2019). In a two-phase group decision task participants could adapt their own choice after the decisions of the other group members had been revealed (Figure 1C). Zhang & Gläscher developed a computational model that combines learning from own experiences (via a classic RL approach) with learning from the other players by computing a value based on the recent reward history of the others. These two value signals were weighted into a single choice value determining the first decision. The switch or stay with the initial choice following the disclosure of group behavior was influenced by the difficulty of the first decision (approximated by the difference between the choice values at the first decision) and the coherence of the group decision.

All three variations of model free RL presented above can capture individuals' learning about an *uncertain* surrounding based on others' behavior in the world. In inverse RL this is achieved by learning "through the eyes" of an observed agent. Thereby, although inverse RL does not include an interactive component and consequently no recursivity, to some degree it can be considered to include an intentional model of the other's decision-making processes as the observed agent's learning is explicitly represented to infer the hidden aspects of the world. In contrast, neither the social influence model nor the expertise tracking model explicitly represent others learning, but track others' task-capacities irrespective of the decision process that lead up to their choices.

Model-based action planning relies on an explicit representation of the environmental dynamics captured by the transition function T . Representing these dynamics allows for flexible decision-making. However, planning an optimal path through a given surrounding requires a representation of the current state. This is complicated when agents cannot directly observe the current state but only receive incomplete information, e.g. when percepts are indistinct or information is ambiguous. Formally, decision making under state uncertainty is captured by Partially Observable Markov Decision Processes (POMDPs) (Kaelbling, Littman, & Cassandra, 1998). In a POMDP information are defined as a set of observations $\Omega = \{o^1, \dots, o^k, \dots, o^n\}$ an agent can make. Actions taken at a given time step a_t and the state s_{t+1} resulting from that action define observation probabilities $O(o_t|s_t, a_{t-1})$. To deal with state uncertainty an agent has to integrate observations to form a belief b_t about the possible states of the surrounding. Agents' beliefs are dynamically updated given their observations and prior beliefs with a Bayesian estimation function:

$$b_t(s_t) = \beta O(o_t|s_t, a_{t-1}) \sum_{s_{t-1}} T(s_t|s_{t-1}, a_{t-1}) b_{t-1}(s_{t-1}),$$

where $\beta = \frac{1}{\Pr(o_t|b_{t-1}, a_{t-1})} = \sum_{s_t} O(o_t|s_t, a_{t-1}) \sum_{s_{t-1}} T(s_t|s_{t-1}, a_{t-1}) b_{t-1}(s_{t-1})$ is a normalizing constant.

To capture core features of ToM, inference of beliefs and preferences from observed behavior, Baker et al. (2017) used an extended POMDP model ("Bayesian ToM model"). To compute a joint posterior probability representing an observer's beliefs about an observed agent's possible beliefs and desires they factorized the likelihood term of an overarching Bayesian learning rule into a model for the observed agent's beliefs and a model of the agent's planning process based on beliefs and desires. Thereby, their Bayesian ToM model formalized detached non-interactive representations of an observing individual's intentional representation of another's

beliefs about an uncertain surrounding based on the observed agent's imperfect perceptual capabilities and the agent's subjective preferences inferred from observed behavior.

Both, model-free and model-based single-agent RL frameworks capture environmental uncertainty. Additionally, using POMDPs and inverse RL the representation of other's intentional learning processes from a detached observational perspective can be formalized. Active engagement and interaction are not captured by any of these models.

Interactive decision models irrespective of environmental uncertainty

When multiple agents engage in interactions while their actions are relevant to one another, i.e. their individual rewards depend on the others' behavior, reasoning can become recursive in the form of "I think, that you think, that I think" and so on. Formally, this sort of thinking is captured by the k -level framework (Stahl, 1993) and cognitive hierarchy theory (Camerer et al., 2004). The depth of reasoning an agent engages in is referred to as its *level* with the parameter k determining the sophistication of depth of reasoning. The concept is well illustrated by the beauty contest game described before (Figure 2A). k -level frameworks are defined in a bottom-up fashion starting with a basic level 0 agent model. A level 0 agent does not reason about others. He might completely disregard any information about the other's behavior and treat it as environmental noise, or assume the co-player is acting according to a hidden distribution (Coricelli & Nagel, 2009; Devaine, Hollard, & Daunizeau, 2014b; Gmytrasiewicz & Doshi, 2005; Yoshida et al., 2008). From level 0 up, k -level models are built. Per definition a k -level agent represents the co-player at level $k-1$, i.e. a level 1 agent represents the other at level 0 while a level 2 agent assumes the co-player acts on level 1 attributing a level 0 model to himself, and so on. The specifics of level

0 models and consequently all higher-level definitions are determined by the decision problem and the underlying basic decision model.

The great advance of k-level frameworks is their ability to model agents' representations of others as intentional and reactive agents, i.e. as agents that act to achieve their goal within a social environment. Intuitively, this can be thought of as "model within a model". Many variants of k-level models exist capturing different social decision problems. Here, we provide an (incomplete) overview over these models and highlight their properties with respect to our two defining categories.

Although not considered a k-level model, *fictitious play* is a basic framework for interactive decision (Brown, 1951). In a fictitious play model, an agent observes the history of a co-player's actions and forms expectations about future actions based on the past choice frequency. In essence, the agent tracks the most frequently selected action in the past. With respect to these predictions the agent then chooses the own action maximizing his reward. The choice history is dynamical tracked via a simple updating rule essentially counting the co-player's frequency of taking any of the available actions. Fictitious play represents co-players' behavior via a very basic choice heuristic, i.e. represents others as sub-intentional level 0 agents. Moving up one step in the reasoning hierarchy, Hampton et al. (2008) introduced an *influence model*: Instead of directly predicting an opponent's behavior from his choice history, the model computes the influence of own behavior onto the opponent, assuming the opponent uses fictitious play. The agent then optimizes own choices with respect to this prediction.

Capturing a comparable level of sophistication as the influence model, a recent paper by Khalvati et al. (2018) utilized a sub-intentional social influence variant of the POMDP framework to model behavior in a public goods game (Figure 2B). In this version of the task, all other players in the group were displayed as identical avatars,

thus rendering the social interaction anonymous (i.e. the identity of the individual other players could not be inferred). As a consequence, a computational model for the task should not try to estimate the individual tendency to contribute to the public good; rather, only the mean contribution probability of the entire group determines the individual's actions. In their implementation of the POMDP, Khalvati and colleagues used a simple beta distribution that is updated on every trial to represent the belief of group contribution to the public good. Effectively, the model for each individual player represents the entire group as a level 0 player that chooses to contribute from an estimated belief distribution, but the other group players are not modelled themselves as intentional agents.

A full k-level framework for strategic decisions was introduced by Devaine and colleagues (2014a) to investigate peoples' different decision strategies when playing against a supposedly human opponent vs a random computer agent in a simple matching pennies task. Level 0 opponents' choice probabilities p_y are assumed to follow a time-varying hidden distribution x^t . Observing actions provides information about the mean μ_t^0 of the underlying distribution. Similar to a prediction error, μ_t^0 is updated at each time point with the difference in observed and expected action:

$$\mu_t^0 = \mu_{t-1}^0 + \sum_t (a_y^t - s(\mu_{t-1}^0))$$

with s denoting a sigmoid choice function. The opponent's decision probability $p_y^{t+1}(a)$ is then computed with a sigmoid decision function based on μ_t^0 and an unknown volatility σ^0 and choice temperature β^0 . Based on this action probability for his opponent the level 0 agent chooses his own action that maximizes his individual expected value. In a recursive fashion, predictions about the opponents' choices for a level 1 agent and higher are built up from this level 0 decision rule. The difficulty in making predictions is that both, the hidden variables governing action probabilities and

the opponent's sophistication level need to be inferred, resulting in a posterior distribution over the opponent's k-levels and variables.

Irrespective of the depth of reasoning hierarchy the models considered above are based on sub-intentional level 0 models, i.e. at the lowest level others are not represented as agents with desires, beliefs or intents. However, the k-level framework by Devaine et al. (2014a) allows for inferring hidden decision variables defining the opponent's behavior. Further, neither basic fictitious play, the influence model, nor the full k-level reasoning model include a representation of the surrounding. Therefore, they are only applicable to simple strategic games taking place in a stable environment such as matching pennies or the inspection game (Figure 2C) (Camerer et al., 2003; Devaine et al., 2014a; Hampton et al., 2008).

Yoshida and colleagues (2008) set out to capture strategic decision making in a complex social environment, the spatial stag-hunt game (Figure 2F). The spatial variation of the simple coordination game requires long term action planning with respect to the (fully observable) surrounding and the co-player's future actions. To integrate both agents acting in the environment they extend a basic MDP model by defining the state space as the product of both agent's admissible states. Consequently, also rewards are defined on the joint state space and the other's uncontrollable actions become a function the other's value potentially inducing cooperative coordinated or individualized behavior. In the next step, optimal strategies with different levels of recursion were computed in the extended multi-agent MDP framework. The resulting k-level MDP model can capture model-based goal directed agents recursively integrating other intentional agents' behavior into their decision process. However, the underlying MDP is fixed and environmental properties are not dynamically learned by agents. Hence, the k-level MDP represents no learning about the environment.

Interactive models for decision making under uncertainty

In this final model section, we introduce two computational frameworks that combine reinforcement learning with a model of other agents, allowing to formalize interactive social decision making under uncertainty. Although both models can be applied to larger multi-agent scenarios, for the sake of simplicity and comprehensibility we only describe a two-agent implementation. We take agent x 's perspective and refer to the second agent as agent y .

Experience weighted attraction (EWA) (Camerer & Ho, 1999) combines simple model free reinforcement learning and belief-based learning in a continuously weighted fashion: On the one hand, simple reinforcement learning estimates the values of own actions given the current rewards and update these value representations via an experienced-based reward prediction error without explicitly taking the actions of the other players into account. On the other hand, belief-based learning (essentially fictitious play) estimates the probability distributions by which the other agent chooses his actions, and then adapts the chosen action accordingly. These two individual forms of learning are implemented in EWA by specific parameter settings (Camerer & Ho, 1999). However, the power of EWA lies in the continuous weighting of these two forms of learning by a trade-off parameter δ (see below). Thereby, EWA combines a representation of the environment, i.e. learning the value of own actions given the current state of affairs, with a simple (i.e. sub-intentional) model of others' actions. The value of all available actions is updated according to a learning rule that combines several variables. The first variable $N(t)$ is equivalent to previous (pre-measurement) experience and updated according to the following:

$$N(t) = \rho N(t - 1) + 1.$$

The parameter ρ is a depreciation term that reflects how fast new reward associations can override prior experience. EWA updates the value of *all* available actions, but the

update distinguishes between the chosen ($a_x^k = a_x(t)$) and the non-chosen actions ($a_x^k \neq a_x(t)$), while holding the action of the other player ($a_y^j(t)$) constant.

$$V_x^k(t) = \begin{cases} \frac{\varphi N(t-1)V_x^k(t-1) + \delta \pi_x(a_x^k, a_y^j(t))}{N(t)}, & \text{if } a_x^k \neq a_x(t) \\ \frac{\varphi N(t-1)V_x^k(t-1) + \pi_x(a_x^k, a_y^j(t))}{N(t)}, & \text{if } a_x^k = a_x(t) \end{cases}$$

Thus, in case of the actual action ($a_x^k = a_x(t)$), the value of that action is update with the full joint reward $\pi(a_x(t), a_y(t))$ whereas the unchosen joint action $a_x^k \neq a_x(t)$ is updated with the delta-weighted joint reward $\delta * \pi(a_x(t), a_y(t))$. It is one of the key insights of EWA (detailed in Camerer & Ho, 1999) that this weighted updating of unchosen action is approximating belief-based learning. EWA has been successfully used in several behavioral economics experiments (see Camerer et al., 2004) and in a decision neuroscience context for modeling choice behavior in patent race games (Zhu, Jiang, Scabini, Knight, & Hsu, 2019; Zhu, Mathewson, & Hsu, 2011).

A combination of model-based reinforcement learning and intentional recursive representations of others' decision processes is given by Interactive Partially Observable Markov Decision Processes (I-POMDPs) (Gmytrasiewicz & Doshi, 2005). Recall that an (individual) POMDP agent x forms beliefs $b_x^t(s_t)$ about the physical states of the environment allowing him to plan actions in an uncertain, only partially accessible surrounding. I-POMDPs extent single-agent POMDPs to the interactive domain by expanding the physical state space S_{phys} by intentional models of another agent y (θ_y) yielding an interactive state space $IS_x = S_{phys} \times \theta_y$. Consequently, an I-POMDP agent's beliefs are no longer over S but IS : $b_x^t(is_t)$. The key component of the model θ_y agent x forms about the second agent y is that it includes y 's beliefs b_y^t . Hence, $b_x^t(is_t)$ is a probability distribution over the multidimensional space spanned by physical states and beliefs of y . Agent y 's beliefs are either over the physical states

space $b_y^t(s)$ (essentially equal to a single agent POMDP) or these beliefs themselves can be over an interactive state space including models of x ($IS_y = S_{phys} \times \theta_x$). The first case, i.e. agent y forming beliefs over physical states only, results in a level-1 model for agent x . He represents agent y as an intentional goal-directed agent that acts to maximize his reward in the world but does not represent or react to other agents in the surrounding. In the latter case, i.e. agent y forming beliefs about x 's beliefs, agent x represents y as reasoning about himself resulting in a level-2 model. Theoretically, this recursion could go on and beliefs could be nested infinitely yielding a higher and higher agent x models. However, as for simpler k-level frameworks, it is reasonable to assume bounded rationality. Dynamic belief updating in the I-POMDP framework follows the same basic Bayesian learning rule as the PODMP update, however it is extended to include an update of the other's belief (for details see Gmytrasiewicz & Doshi, 2005). This requires solving models of y in bottom up manner, essentially simulating y 's learning and decision process. Environmental uncertainty and uncertainty regarding the other's model of oneself hamper this process. The I-PODMP framework has been used to model behavior and reasoning processes in a multi-round trust task (Figure 2C) (Hula, Montague, & Dayan, 2015; Hula, Vilares, Lohrenz, Dayan, & Montague, 2018) and the centipede game (Figure 2I) (Doshi, Qu, Goodie, & Young, 2012)(Doshi, Qu, & Goodie, 2014). In its application to the trust game the I-PODMP allowed formalizing players' risk aversion and guilt as well the their recursive representations of their co-player's risk and guilt parameters which define the rate of cooperation, and breakdown and reestablishment of cooperation between players over the course of multiple interactions (Hula et al., 2018). Using I-PODMP models to capture decision processes in a centipede game showed that participants mostly adopt decision strategies of level-1 or level-2 and that the depth of recursive reasoning increases in more competitive scenarios (Doshi et al., 2012).

8. Neural responses

Imaging studies deploying false belief or perspective taking tasks most commonly associate reasoning about another person's mental states with activation in the temporal parietal junction (TPJ) (sometimes also posterior superior temporal sulcus, pSTS) and medial prefrontal cortex (mPFC) and precuneus (see Schurz et al., 2014 for a meta-analysis). On the other hand, (non-social) learning processes, as the computation of reward prediction errors and expected stimulus values, have been shown to activate the ventral striatum (see Garrison, Erdeniz, & Done, 2013 for a review) and ventromedial prefrontal cortex (vmPFC) (see Rangel & Clithero, 2014 for a review) respectively. Varying degrees of interaction between the between these two mechanisms as well as additional neural processes are found in the decision tasks we discussed with respect to interactivity and uncertainty.

In the observational learning task that required tracking the expertise of observed agents making choices in an uncertain surrounding (Figure 1D) Boorman et al. (2013) showed that expertise estimation was two-fold: first, when the observed agent's choice was revealed participants evaluated the quality of the agent's action with respect to their own expectation about the stimulus. Second, at feedback, they updated their estimate about the other based on the correctness of his action. Prediction errors about estimated expertise at choice, i.e. the mismatch between own expectation about stimulus values and the observed action, were found to covary with activity in TPJ, dmPFC and STS. Prediction errors at feedback correlated with activity in dlPFC and precuneus. Participants own reward expectation and reward prediction errors correlated with vmPFC and striatum.

Observational learning (Figure 1F) by Collette et al. (2017) focused on inverse RL, i.e. the inference of the reward distribution by observing actions of others. Model-based

fMRI results indicate that in this scenario the brain encodes the expected values of actions in the dmPFC in the observed agent's value space instead of the observing participant's own value space suggesting a simulation of the observed agent's perspective. TPJ activity reflected how much the observer updated his belief over the reward distribution. They argue that this indicates a more general function of TPJ in social situations.

Zhang & Gläscher (2019) found that learning in their social influence task (Figure 1C) combines learning from own experience with learning from others based on their outcome history. Own and observational value signals were found to be represented in vmPFC and ACC respectively closely corresponding to a previous report by Behrens and colleagues (2008), which the fidelity of social advice covaried with the same regions of the ACC, whereas they could also localize the own value to the vmPFC. Further, they found that the degree to which own decisions are modulated by the group behavior positively correlated with activity in TPJ.

Hampton et al. (2008) modeled participants' behavior in the inspection game, a competitive choice task similar to matching pennies, with the influence model. From the model they derive an expected value and a classic reward prediction error signal. Second, they calculated an "influence prediction error" - the difference between the expected degree of influence exerted on the opponent's action and the actual influence, once the outcome was revealed. They found evidence for two distinct updating mechanisms. First, RL updating was demonstrated by a correlation of expected value with activity in mPFC and by a correspondence between reward prediction errors and striatal activation. Second, value updating related to social influence was demonstrated by a correlation of pSTS activity with the model-derived influence prediction error. These findings were verified and reinforced in a study by Hill et al. (2017) who used the same design with repetitive TMS (transcranial magnetic

stimulation) stimulation over either TPJ or the vertex. They found that TPJ but not vertex stimulation disrupted the integration of own actions' effects and the opponent's decision strategy highlighting the importance of TPJ (pSTS) in the computations of strategic decisions.

Using the beauty contest game (Figure 2A) Coricelli & Nagel (2009) examined the neural signals related to higher level reasoning processes. The authors recorded fMRI while participant played the game either with other humans or a random computer. In most participants they found evidence for reasoning between level 1 and 3. On the neural level, they showed mPFC, posterior cingulate cortex, STS and TPJ activity to be greater in the human as compared to the computer condition. Furthermore, they find that mPFC, but none of the other regions covaries with the depth of participants' reasoning.

Zhu et al. (2011) computed standard RL prediction errors and belief prediction errors using an EWA model in a patent race game (Figure 2H). The belief prediction error codes participants' surprise about actions chosen by the opponent. Expectations about actions are computed based on the history of the opponent's actions. Reward prediction errors were found to be coded in ventral striatum while belief prediction errors were coded in ventral striatum as well as rostral ACC.

In a spatialized version of the stag hunt game (Figure 2F) Yoshida and colleagues (2010) examined the neural responses related to recursive action planning. They found that activity in mPFC coded participants' uncertainty about the co-players actions while the level of recursion they engaged in was represented in dlPFC, FEF and SPL.

In the advisor task (Figure 3A) Behrens et al. (2008) tested participants' continuous learning about a volatile environment and about the probability of an advisor with changing intentions providing them with correct advice. They used a Bayesian learning model to capture the changes in outcome probabilities and in the advisor's advice

fidelity. They found two adjacent but distinguishable regions in the ACC to track both types of volatility. Further, the computation of reward probability based on both types of learning signals was reflected in vmPFC activity weighted by the two parallel learning streams.

In summary, while there remain significant differences in brain regions that are recruited by these social decision-making tasks, a basic commonality seems to be that both TPJ and mPFC (and sometimes ACC) are often involved in representing information that is related to another player (esp. belief representation), while model-based often RL-related signals (like own expected values or reward prediction errors) that are derived from experiential learning or that related to the focal player herself often covary with brain activation in the (ventral) striatum and vmPFC. However, it cannot be ruled out that the heterogeneity of tasks and computational models plays a defining role in the divergence of brain activations across many studies. To resolve this interpretational ambiguity, task design for future studies is advised to include manipulations of environmental and social uncertainty and interactivity and choose a modeling framework capable of action planning under these circumstances.

9. Conclusion

When aiming at examining human ToM capacities several important aspects need to be considered. First, one should be aware that Theory of Mind is highly inclusive concept that implies various different cognitive sub-functions ranging from affective to more cognitive functions (Schaafsma et al., 2015). Further, as shown in the first section of this review, even within one “subsection” of these functions the cognitive processes differ strongly depending on the specifics of the social situation. We argue that the two dimensions of *uncertainty* and *interactivity* can provide an effective classification scheme for these types of social reasoning. When faced with environmental

uncertainty others' beliefs might differ from one's own creating the need to represent and track their beliefs in addition to representing the physical environment. A second source of uncertainty are individuals' hidden internal variables like subjective preferences or personality traits that can be inferred over time by observing behavior. The level of interactivity is highly important as it determines whether own actions have an effect on others, as well as the effect of others' behavior on oneself. This in turn determines whether recursive reasoning is needed. Tasks aimed at investigating ToM processes in their true richness should therefore include both dimensions. Further, for a complete and ecologically valid picture ToM tasks should include rich environments and face-to-face interactions to allow for combined examination of affective and cognitive ToM processes.

Computational cognitive models provide quantitative testable descriptions of the various hidden cognitive sub-functions constituting ToM. Applying models that capture the different aspects might allow us to disentangle the multitude sub-processes. Here, we put a strong focus on the “social reinforcement learning” and recapitulated how their abilities to capture strategic social interactions.

The need for true interactivity has also been highlighted by proponents of “second-person neuroscience”. Advocates of this few argue that cognition during interaction differs fundamentally from observational scenarios not only as recursive thinking is only elicited in the latter but also due to qualitative components like the *feeling of engagement with others* (Redcay & Schilbach, 2019; Schilbach et al., 2013). Further they emphasizes the importance of “multi-brain” or “hyperscanning” studies during which the neural activity of multiple interacting agents is recorded. Hyperscanning recordings have successfully been conducted using fMRI and EEG and revealed specific synchronizations between the neural activity of interacting partners during abstract communication and motion matching tasks (Dumas, Nadel, Soussignan, Martinerie, &

Garnero, 2010; Stolk, Verhagen, & Toni, 2015). In line with these views we argue for investigating ToM in rich interactive environments while simultaneously recording neural activity from all interacting individuals to be able to capture these multifaceted processes.

10. Acknowledgements

JG and TR were supported by the Collaborative Research Center TRR 169 “Crossmodal Learning” funded by the German Research Foundation (DFG) and the National Science Foundation of China (NSFC). JG was further funded by a Computational Neuroscience grant from the German Ministry of Education and Research (BMBF, 01GQ1603) and the United States National Science Foundation.

11. References

- Arslan, B., Taatgen, N. A., Verbrugge, R., & Stewart, T. C. (2017). Five-Year-Olds ' Systematic Errors in Second-Order False Belief Tasks Are Due to First-Order Theory of Mind Strategy Selection : A Computational Modeling Study, 8(February), 1–16. <https://doi.org/10.3389/fpsyg.2017.00275>
- Axelrod, R., & Hamilton, W. (1981). The evolution of cooperation. *Science*, 211(4489), 1390–1396. <https://doi.org/10.1126/science.7466396>
- Baillargeon, R., Scott, R. M., & He, Z. (2010). False-belief understanding in infants. *Trends in Cognitive Sciences*, 14(3), 110–118. <https://doi.org/10.1016/j.tics.2009.12.006>
- Baker, C. L., Jara-Ettinger, J., Saxe, R., & Tenenbaum, J. B. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, 1(4), 1–10. <https://doi.org/10.1038/s41562-017-0064>

- Baron-Cohen, S. (1988). Social and pragmatic deficits in autism: Cognitive or affective? *Journal of Autism and Developmental Disorders*, 18(3), 379–402.
<https://doi.org/10.1007/BF02212194>
- Baron-Cohen, S. (2000). Theory of mind and autism: A review. *International Review of Research in Mental Retardation*, 23, 169–184. [https://doi.org/10.1016/S0074-7750\(00\)80010-5](https://doi.org/10.1016/S0074-7750(00)80010-5)
- 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](https://doi.org/10.1016/0010-0277(85)90022-8)
- Behrens, T. E. J., Hunt, L. T., Woolrich, M. W., & Rushworth, M. F. S. (2008). Associative learning of social value. *Nature*, 456(7219), 245–249.
<https://doi.org/10.1038/nature07538>
- Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior*. <https://doi.org/10.1006/game.1995.1027>
- Boorman, E. D., O’Doherty, J. P., Adolphs, R., & Rangel, A. (2013). The Behavioral and Neural Mechanisms Underlying the Tracking of Expertise. *Neuron*, 80(6), 1558–1571. <https://doi.org/10.1016/j.neuron.2013.10.024>
- Bora, E., Yucel, M., & Pantelis, C. (2009). Theory of mind impairment in schizophrenia: Meta-analysis. *Schizophrenia Research*, 109(1–3), 1–9.
<https://doi.org/10.1016/j.schres.2008.12.020>
- Brown, G. (1951). Iterative solution of games by fictitious play. In *Activity analysis of production and allocation* (pp. 374–376). New York: Wiley.
- Burke, C. J., Tobler, P. N., Baddeley, M., & Schultz, W. (2010). Neural mechanisms of observational learning. *Proceedings of the National Academy of Sciences*, 107(32), 14431–14436. <https://doi.org/10.1073/pnas.1003111107>
- Call, J., & Tomasello, M. (1999). A non-verbal false belief task: the performance of

- chimpanzees and human children. *Child Development*, 70(2), 381–395.
- Camerer, C. F. (2003). Behavioural studies of strategic thinking in games. *Trends in Cognitive Sciences*, 7(5), 225–231. [https://doi.org/10.1016/S1364-6613\(03\)00094-9](https://doi.org/10.1016/S1364-6613(03)00094-9)
- Camerer, C. F., Ho, T.-H., & Chong, J.-K. (2004). A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics*, 119(3), 861–898. <https://doi.org/10.1162/0033553041502225>
- Camerer, C. F., Ho, T. H., & Chong, J. K. (2015). A psychological approach to strategic thinking in games. *Current Opinion in Behavioral Sciences*, 3(1), 157–162. <https://doi.org/10.1016/j.cobeha.2015.04.005>
- Camerer, C., Ho, T., & Chong, J.-K. (2003). A Cognitive Hierarchy Theory of One-Shot Games and Experimental Analysis. *Ssrn*, (December 2001). <https://doi.org/10.2139/ssrn.411061>
- Charpentier, C. J., & O'Doherty, J. P. (2018). The application of computational models to social neuroscience: promises and pitfalls. *Social Neuroscience*, 13(6), 637–647. <https://doi.org/10.1080/17470919.2018.1518834>
- Collette, S., Pauli, W. M., Bossaerts, P., & O'Doherty, J. (2017). Neural computations underlying inverse reinforcement learning in the human brain. *ELife*, 6, 1–20. <https://doi.org/10.7554/eLife.29718>
- Coricelli, G., & Nagel, R. (2009). Neural correlates of depth of strategic reasoning in medial prefrontal cortex. *Proceedings of the National Academy of Sciences*, 106(23), 9163–9168. <https://doi.org/10.1073/pnas.0807721106>
- Dasgupta, P., & Stiglitz, J. (1980). Uncertainty, Industrial Structure, and the Speed of R&D. *The Bell Journal of Economics*, 11(1), 1. <https://doi.org/10.2307/3003398>
- Daw, N. D., & Dayan, P. (2014). The algorithmic anatomy of model-based evaluation.

- Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1655), 20130478. <https://doi.org/10.1098/rstb.2013.0478>
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8(12), 1704–1711. <https://doi.org/10.1038/nn1560>
- Devaine, M., Hollard, G., & Daunizeau, J. (2014a). The Social Bayesian Brain: Does Mentalizing Make a Difference When We Learn? *PLoS Computational Biology*, 10(12). <https://doi.org/10.1371/journal.pcbi.1003992>
- Devaine, M., Hollard, G., & Daunizeau, J. (2014b). Theory of mind: Did evolution fool us? *PLoS ONE*, 9(2). <https://doi.org/10.1371/journal.pone.0087619>
- Diaconescu, A. O., Mathys, C., Weber, L. a E., Daunizeau, J., Kasper, L., Lomakina, E. I., ... Stephan, K. E. (2014). Inferring on the Intentions of Others by Hierarchical Bayesian Learning. *PLoS Computational Biology*, 10(9), e1003810. <https://doi.org/10.1371/journal.pcbi.1003810>
- Doshi, P., & Gmytrasiewicz, P. (2004). A Particle Filtering Algorithm for Interactive POMDPs. In *In Proceedings of Conference on Modeling Other Agents from Observation* (pp. 87–93). Retrieved from <http://u.cs.biu.ac.il/~galk/moo2004/proceedings/12.pdf>
- Doshi, P., Qu, X., & Goodie, A. (2014). Decision-Theoretic Planning in Multiagent Settings with Application to Behavioral Modeling. *Plan, Activity, and Intent Recognition: Theory and Practice*, 205–224. <https://doi.org/10.1016/B978-0-12-398532-3.00008-7>
- Doshi, P., Qu, X., Goodie, A. S., & Young, D. L. (2012). Modeling human recursive reasoning using empirically informed interactive partially observable markov decision processes. *IEEE Transactions on Systems, Man, and Cybernetics Part A: Systems and Humans*, 42(6), 1529–1542.

<https://doi.org/10.1109/TSMCA.2012.2199484>

Dufour, N., Redcay, E., Young, L., Mavros, P. L., Moran, J. M., Triantafyllou, C., ...

Saxe, R. (2013). Similar Brain Activation during False Belief Tasks in a Large Sample of Adults with and without Autism, 8(9).

<https://doi.org/10.1371/journal.pone.0075468>

Dumas, G., Nadel, J., Soussignan, R., Martinerie, J., & Garnero, L. (2010). Inter-Brain Synchronization during Social Interaction. *PLoS ONE*, 5(8).

<https://doi.org/10.1371/journal.pone.0012166>

Gallagher, H. L., & Frith, C. D. (2003). Functional imaging of 'theory of mind.' *Trends in Cognitive Sciences*, 7(2), 77–83. [https://doi.org/10.1016/S1364-6613\(02\)00025-6](https://doi.org/10.1016/S1364-6613(02)00025-6)

Garrison, J., Erdeniz, B., & Done, J. (2013). Prediction error in reinforcement learning: A meta-analysis of neuroimaging studies. *Neuroscience & Biobehavioral Reviews*, 37(7), 1297–1310.

<https://doi.org/10.1016/j.neubiorev.2013.03.023>

Gibbons, R. (1992). *Game Theory for Applied Economists*. Princeton University Press. <https://doi.org/10.2307/j.ctvcmxrzd>

Gläscher, J., Daw, N., Dayan, P., & O'Doherty, J. P. (2010). States versus Rewards: Dissociable Neural Prediction Error Signals Underlying Model-Based and Model-Free Reinforcement Learning. *Neuron*, 66(4), 585–595.

<https://doi.org/10.1016/j.neuron.2010.04.016>

Gläscher, J. P., & O'Doherty, J. P. (2010). Model-based approaches to neuroimaging: combining reinforcement learning theory with fMRI data. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(4), 501–510.

<https://doi.org/10.1002/wcs.57>

Gmytrasiewicz, P. J., & Doshi, P. (2005). A Framework for Sequential Planning in

- Multi-Agent Settings. *Journal of Artificial Intelligence Research*, 24, 49–79.
<https://doi.org/10.1613/jair.1579>
- Hamilton, A. F. d. C., Brindley, R., & Frith, U. (2009). Visual perspective taking impairment in children with autistic spectrum disorder. *Cognition*, 113(1), 37–44.
<https://doi.org/10.1016/j.cognition.2009.07.007>
- Hampton, A. N., Bossaerts, P., & O’Doherty, J. P. (2008a). Neural correlates of mentalizing-related computations during strategic interactions in humans. Supporting Information. *Proceedings of the National Academy of Sciences of the United States of America*, 105(18), 1–6.
<https://doi.org/10.1073/pnas.0711099105>
- Hampton, A. N., Bossaerts, P., & O’Doherty, J. P. (2008b). Neural correlates of mentalizing-related computations during strategic interactions in humans. Supporting Information. *Proceedings of the National Academy of Sciences of the United States of America*, 105(18), 1–6.
- Hedden, T., & Zhang, J. (2002). What do you think I think you think?: Strategic reasoning in matrix games. *Cognition*, 85(1), 1–36.
[https://doi.org/10.1016/S0010-0277\(02\)00054-9](https://doi.org/10.1016/S0010-0277(02)00054-9)
- Hill, C. A., Suzuki, S., Polania, R., Moisa, M., O’Doherty, J. P., & Ruff, C. C. (2017). A causal account of the brain network computations underlying strategic social behavior. *Nature Neuroscience*, 20(8), 1142–1149.
<https://doi.org/10.1038/nn.4602>
- Hill, M. R., Boorman, E. D., & Fried, I. (2016). Observational learning computations in neurons of the human anterior cingulate cortex. *Nature Communications*, 7(1), 12722. <https://doi.org/10.1038/ncomms12722>
- Ho, T.-H., & Su, X. (2013). A Dynamic Level- k Model in Sequential Games. *Management Science*, 59(2), 452–469. <https://doi.org/10.1287/mnsc.1120.1645>

- Houser, D., & McCabe, K. (2014). Experimental Economics and Experimental Game Theory. In *Neuroeconomics* (pp. 19–34). Elsevier. <https://doi.org/10.1016/B978-0-12-416008-8.00002-4>
- Hula, A., Montague, P. R., & Dayan, P. (2015). Monte Carlo Planning Method Estimates Planning Horizons during Interactive Social Exchange. *PLOS Computational Biology*, 11(6). <https://doi.org/10.1371/journal.pcbi.1004254>
- Hula, A., Vilares, I., Lohrenz, T., Dayan, P., & Montague, P. R. (2018). A model of risk and mental state shifts during social interaction. *PLoS Computational Biology*, 14(2), 1–20. <https://doi.org/10.1371/journal.pcbi.1005935>
- Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and Acting in Partially Observable Stochastic Domains. *Artificial Intelligence*, 101(1–2), 99–134. [https://doi.org/10.1016/S0004-3702\(98\)00023-X](https://doi.org/10.1016/S0004-3702(98)00023-X)
- Kalbe, E., Schlegel, M., Sack, A. T., Nowak, D. A., Dafotakis, M., Bangard, C., ... Kessler, J. (2010). Dissociating cognitive from affective theory of mind: A TMS study. *Cortex*, 46(6), 769–780. <https://doi.org/10.1016/J.CORTEX.2009.07.010>
- Khalvati, K., Park, S. A., Mirbagheri, S., Philippe, R., Sestito, M., Dreher, J.-C., & Rao, R. P. N. (2018). Bayesian Inference of Other Minds Explains Human Choices in Group Decision Making. *BioRxiv*, 419515. <https://doi.org/10.1101/419515>
- King-Casas, B., Sharp, C., Lomax-Bream, L., Lohrenz, T., Fonagy, P., & Read Montague, P. (2008). The rupture and repair of cooperation in borderline personality disorder. *Science*, 321(5890), 806–810. <https://doi.org/10.1126/science.1156902>
- King-Casas, B., Tomlin, D., Anen, C., Camerer, C. F., Quartz, S. R., & Montague, P. R. (2005). Getting to Know You: Reputation and Trust in a Two-Person Economic Exchange. *Science*, 308(5718), 78–83.

<https://doi.org/10.1126/science.1108062>

- Kleiman-Weiner, M., Ho, M., Austerweil, J., Littmann, M. L., & Tenenbaum, J. (2016). Coordinate to cooperate or compete: Abstract goals and joint intentions in social interaction. *Proceedings of the 38th Annual Conference of the Cognitive Science Society*, 1, 1679–1684.
- Kumar, S., Rusch, T., Doshi, P., Spezio, M., & Gläscher, J. (2019). Modeling cooperative and competitive decision-making in the Tiger Task. In *The Multi-disciplinary Conference on Reinforcement Learning and Decision Making*.
- Ledyard, J. O. (1994). Public Goods: A Survey of Experimental Research. In J. Kagel & A. Roth (Eds.), *The Handbook of Experimental Economics*. Princeton NJ: Princeton University Press.
- Loury, G. C. (1979). Market Structure and Innovation. *The Quarterly Journal of Economics*, 93(3), 395. <https://doi.org/10.2307/1883165>
- McKelvey, R. D., & Palfrey, T. R. (1992). An Experimental Study of the Centipede Game. *Econometrica*, 60(4), 803. <https://doi.org/10.2307/2951567>
- Mitchell, R. L. C., & Phillips, L. H. (2015). Neuropsychologia The overlapping relationship between emotion perception and theory of mind. *Neuropsychologia*, 70, 1–10. <https://doi.org/10.1016/j.neuropsychologia.2015.02.018>
- Mookherjee, D., & Sopher, B. (1994). Learning Behavior in an Experimental Matching Pennies Game. *Games and Economic Behavior*, 7(1), 62–91. <https://doi.org/10.1006/game.1994.1037>
- Ng, A. Y., & Russell, S. J. (2000). Algorithms for inverse reinforcement learning. *ICML*, 1, 2.
- Rangel, A., & Clithero, J. A. (2014). The Computation of Stimulus Values in Simple Choice. *Neuroeconomics*, 125–148. <https://doi.org/10.1016/B978-0-12-416008-8.00008-5>

- Redcay, E., & Schilbach, L. (2019). Using second-person neuroscience to elucidate the mechanisms of social interaction. *Nature Reviews Neuroscience*, 20(8), 495–505. <https://doi.org/10.1038/s41583-019-0179-4>
- Rosenthal, R. W. (1981). Games of perfect information, predatory pricing and the chain-store paradox. *Journal of Economic Theory*, 25(1), 92–100. [https://doi.org/10.1016/0022-0531\(81\)90018-1](https://doi.org/10.1016/0022-0531(81)90018-1)
- Rousseau, J. J., & Cranson, M. (1984). *A discourse on inequality* (6th ed.). Hamondsworth England: Penguin Books.
- Sanfey, A. G. (2007). Social decision-making: insights from game theory and neuroscience. *Science (New York, N.Y.)*, 318(5850), 598–602. <https://doi.org/10.1126/science.1142996>
- Saxe, R., & Kanwisher, N. (2003). People thinking about thinking people. The role of the temporo-parietal junction in “theory of mind.” *NeuroImage*, 19(4), 1835–1842. [https://doi.org/10.1016/S1053-8119\(03\)00230-1](https://doi.org/10.1016/S1053-8119(03)00230-1)
- Schaafsma, S. M., Pfaff, D. W., Spunt, R. P., & Adolphs, R. (2015). Deconstructing and reconstructing theory of mind. *Trends in Cognitive Sciences*, 19(2), 65–72. <https://doi.org/10.1016/j.tics.2014.11.007>
- 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(4), 393–414. <https://doi.org/10.1017/S0140525X12000660>
- Schurz, M., Radua, J., Aichhorn, M., Richlan, F., & Perner, J. (2014). Fractionating theory of mind: A meta-analysis of functional brain imaging studies. *Neuroscience & Biobehavioral Reviews*, 42, 9–34. <https://doi.org/10.1016/j.neubiorev.2014.01.009>
- Selbing, I., & Olsson, A. (2017). Beliefs about Others’ Abilities Alter Learning from Observation. *Scientific Reports*, 7(1), 16173. <https://doi.org/10.1038/s41598-017->

- Senju, A., Southgate, V., Miura, Y., Matsui, T., Hasegawa, T., Tojo, Y., ... Csibra, G. (2010). Absence of spontaneous action anticipation by false belief attribution in children with autism spectrum disorder. *Development and Psychopathology*, 22(2), 353–360. <https://doi.org/10.1017/S0954579410000106>
- Shum, M., Kleiman-Weiner, M., Littman, M. L., & Tenenbaum, J. B. (2019). Theory of Minds: Understanding Behavior in Groups Through Inverse Planning. Retrieved from <http://arxiv.org/abs/1901.06085>
- Siegal, M., & Varley, R. (2002). Neural systems involved in “theory of mind.” *Nature Reviews Neuroscience*, 3(6), 463–471. <https://doi.org/10.1038/nrn844>
- Skyrms, B. (2001). The Stag Hunt. *Proceedings and Addresses of the American Philosophical Association*, 75(2), 31. <https://doi.org/10.2307/3218711>
- Stahl, D. O. (1993). Evolution of Smartn Players. *Games and Economic Behavior*, 5(4), 604–617. <https://doi.org/10.1006/game.1993.1033>
- Stolk, A., Verhagen, L., & Toni, I. (2015). Conceptual Alignment: How Brains Achieve Mutual Understanding. *Trends in Cognitive Sciences*, xx, 1–12. <https://doi.org/10.1016/j.tics.2015.11.007>
- Surian, L., Caldi, S., & Sperber, D. (2007). Attribution of Beliefs by 13-Month-Old Infants. *Psychological Science*, 18(7), 580–586. <https://doi.org/10.1111/j.1467-9280.2007.01943.x>
- Sutton, R. S., & Barto, A. G. (2012). *Reinforcement Learning : An Introduction* (2nd ed.). Cambridge, MA, USA: MIT Press.
- Touhey, J. C. (1974). Decision Processes, Expectations, and Adoption of Strategies in Zero-Sum Games. *Human Relations*, 27(8), 813–824. <https://doi.org/10.1177/001872677402700807>
- Wan Lee, S., Shimojo, S., & O’Doherty, J. P. (2014). Neural Computations

- Underlying Arbitration between Model-Based and Model-free Learning. *Neuron*, 81(3), 687–699. <https://doi.org/10.1016/j.neuron.2013.11.028>
- Wimmer, H., & Perner, J. (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](https://doi.org/10.1016/0010-0277(83)90004-5)
- Xiao, Y., Geng, F., Riggins, T., Chen, G., & Redcay, E. (2019). Neural correlates of developing theory of mind competence in early childhood. *NeuroImage*, 184(June 2018), 707–716. <https://doi.org/10.1016/j.neuroimage.2018.09.079>
- Yoshida, W., Seymour, B., Friston, K. J., & Dolan, R. J. (2010). Neural Mechanisms of Belief Inference during Cooperative Games. *Journal of Neuroscience*, 30(32), 10744–10751. <https://doi.org/https://doi.org/10.1523/JNEUROSCI.5895-09.2010>
- Yoshida, Wako, Dolan, R. J., & Friston, K. J. (2008). Game Theory of Mind. *PLoS Computational Biology*, 4(12). <https://doi.org/10.1371/journal.pcbi.1000254>
- Zhang, L., & Glascher, J. P. (2019). A network supporting social influences in human decision-making. *BioRxiv*, (February), 551614. <https://doi.org/https://doi.org/10.1101/551614>
- Zhu, L., Jiang, Y., Scabini, D., Knight, R. T., & Hsu, M. (2019). Patients with basal ganglia damage show preserved learning in an economic game. *Nature Communications*, 10(1), 802. <https://doi.org/10.1038/s41467-019-08766-1>
- Zhu, L., Mathewson, K. E., & Hsu, M. (2011). Dissociable neural representations of reinforcement and belief prediction errors underlie strategic learning, (12), 1–6. <https://doi.org/10.1073/pnas.1116783109>