Book chapter, now published in *The Neural Bases of Mentalizing*

Full citation: Charpentier C.J., & O'Doherty J.P. (2021). Computational approaches to mentalizing during observational learning and strategic social interactions. In *The Neural Bases of Mentalizing* (Eds. Kevin Ochsner & Michael Gilead), 489-501. Springer: New York.
doi: 10.1007/978-3-030-51890-5_24

Computational approaches to mentalizing during observational learning and strategic social interactions

Caroline J. Charpentier^{1*}, John P. O'Doherty¹

¹*California Institute of Technology, Pasadena, CA, USA*

*Contact: ccharpen@caltech.edu

Abstract

In order to make decisions, we often seek and integrate information coming from other people, while at times also keeping track of the knowledge other people acquire from observing our own actions. In this chapter, we examine the computational mechanisms and the involvement of mentalizing when we learn from observing other people and when we engage in strategic social interactions in which two agents recursively represent each other's beliefs and intentions. We shed light on evidence that regions of the brain's mentalizing system play a key role in implementing these social learning computations, by representing (i) the mental states of other agents, (ii) how these mental states are dynamically updated over time and (iii) how other agents represent our own beliefs and intentions. We argue in favor of using a neuro-computational approach to study these processes, combining computational modelling to identify specific variables predicting behavior and model-based neuroimaging to understand how and where these variables are represented in the brain. We conclude by highlighting some open questions that remain to be addressed to provide a fully integrated account of mentalizing computations during social learning.

Keywords

mentalizing, social learning, decision-making, observation, strategic thinking, belief inference, computation, model-based neuroimaging

To navigate our social and connected world, it is key for individuals to be able to learn from other people. Whether it is learning a new skill by observing an expert perform it, learning to seek rewards and to avoid punishments, or making complex strategic decisions, learning from others is prevalent in our daily lives. In this chapter, we give an overview of the behavioral and neural computations at play when we attribute mental states to other agents in order to learn from them. We do so by focusing on two important social learning behaviors and describing the role of mentalizing computations in these processes: observational learning, which involves integrating information received from another agent into one's own beliefs, and strategic thinking, which involve recursive belief inference between agents in order to win a competition or reach a common goal.

1. The role of mentalizing during observational learning

The goal of observational learning is to learn which actions and decisions in the environment are good – i.e. likely to lead to positive outcomes – or bad – i.e. likely to lead to negative outcomes – by observing other people perform those actions. Any species endowed with the ability to engage in observational learning has an evolutionary advantage, as it allows individuals to learn about threatening outcomes without having to experience them directly. It can also allow individuals to learn without observing outcomes at all.

Recent research on observational learning in humans has shed light on three possible strategies that can be employed during observational learning (Charpentier & O'Doherty, 2018; Dunne & O'Doherty, 2013). The first strategy is vicarious reward learning, in which the observer learns by observing another agent making decisions and experiencing outcomes, rather than directly experiencing those outcomes themselves. Similar to experiential learning, the observer is able to form associations between actions and outcomes, but does so through observation. They can then use these learned associations in order to make decisions. A simple computation encoded in the brain to underlie vicarious reward learning is an observational reward prediction error (oRPE), calculated as the difference between the other agent's expected and actual outcome. These oRPEs have been found to be represented in several brain areas, namely vmPFC (Burke et al., 2010; Suzuki et al., 2012), dorsal striatum (Cooper et al., 2012) and ACC (M. R. Hill et al., 2016). A second observational learning strategy is action imitation, which allows people to learn simply from observing the actions performed by another agent and to repeat or copy the most frequently taken action. Computationally, action imitation learning can be explained in a reinforcement learning framework, with an action prediction error (APE) – the difference between the expected and actual actions of the other agent – reinforcing previously chosen actions positively and unchosen actions negatively. Observers are therefore more likely to perform actions that were also performed by the agent being observed. Neuroimaging results confirmed that APEs are tracked in the brain, specifically in the dmPFC, dlPFC and inferior parietal lobule (Burke et al., 2010; Suzuki et al., 2012). Finally, the third and more complex observational learning strategy falls under the term “emulation”. In emulation learning, observers learn by inferring the other agents' intentions, goals, beliefs and hidden mental states. The exact computational form of such an inference process is still being investigated, but recent literature suggests that it could take place as a Bayesian inference process, whereby prior beliefs about the other agent's goals are combined with the evidence received from observing the agent's decisions to produce an updated posterior of the inferred beliefs (Charpentier et al., 2020; Charpentier & O'Doherty, 2018; Collette et al., 2017; Devaine et al., 2014; Diaconescu et al., 2014; Farmer & Kashdan, 2012). Interestingly, brain areas that have been identified as playing a role in implementing these belief update computations overlap with the mentalizing network: TPJ, pSTS and dmPFC (Behrens et al., 2008; Boorman et al., 2013; Charpentier et al., 2020; Collette et al., 2017). For example, in Collette et al (2017), the inference model that best explained

participants' behavior was an inverse reinforcement-learning (RL) model, whereby instead of learning the value of an action from observing outcomes (classical RL), individuals infer the outcome distribution from observing another agent's actions (inverse RL). The dmPFC was found to contribute to this mechanism by representing the value of the predicted outcomes in agent-referential space (i.e. from the point of view of the agent, not the participant). In addition, the TPJ and pSTS were found to track a learning signal, specifically the entropy or "surprise" predicted by the inverse RL model when observing the other agent's chosen action. While more computationally expensive than vicarious reward learning and action imitation, emulation learning is very adaptive and flexible, can integrate over multiple social signals, and allows the observer to learn from an agent that had different preferences, goals, or even a competing agenda, which will be discussed in more detail in the second part of this chapter.

These strategies can also be used in combination: vicarious reward learning and action imitation (Burke et al., 2010; Suzuki et al., 2012), imitation and emulation (Charpentier et al., 2020), experiential and social learning (Pärnamets & Olsson, 2020; Zhang & Gläscher, 2020). Yet, an important outstanding question concerns how it is that people decide or arbitrate between strategies. For example, in a situation where outcomes cannot be directly observed, people can only rely on imitation or emulation in order to learn from observing another agent. The factors that influence the decision to rely on one strategy over the other remain to be elucidated. An interesting hypothesis is that of an arbitration mechanism that would be influenced by the relative estimated reliability of each strategy in the current environment (similar to Lee, Shimojo, & O'Doherty, 2014). In Charpentier et al. (2020), we find that people arbitrate between imitation and emulation learning solely based on the reliability of the emulation system, which depended on the uncertainty of the emulation prediction. This reliability signal was represented in the brain, specifically in the ventrolateral prefrontal cortex (vlPFC), but also in the TPJ and ACC. Many factors could play a role in pushing this arbitrator around, such as uncertainty, expertise, or trust in the other agent. For example, if inferring the other agent's intentions becomes more difficult because of increased uncertainty about the evidence provided by observing their behavior, emulation would become more computationally demanding, and it is likely that learning behavior would preferentially rely on imitation. Inversely, if the other agent is deemed incompetent or untrustworthy, simply imitating them may lead to a lot of mistakes and emulation may be favored.

Given this short overview of observational learning, it seems clear that mentalizing plays a role during this process, and does so mainly through emulation learning, which relies on inferring the mental state of another person, whether it is their goals, preferences, beliefs or intentions. In contrast, vicarious reward learning and action imitation function with simple associative computations, either between action and outcome (vicarious reward learning) or between actions performed by others and actions performed by the self (imitation), suggesting that mentalizing probably doesn't play a role in those strategies. Nonetheless, we note that there may be some degree of overlap in the brain regions involved in the different strategies, for example between imitation and emulation neural signals. Specifically, action prediction errors were found to be encoded in the dmPFC during imitation learning (Burke et al., 2010; Suzuki et al., 2012). The main hypothesis, however, remains that action imitation occurs through the representation of another person's actions in the mirror neuron system, which is active both when an individual performs an action and when they observe another person performing that same action and include regions of the premotor cortex and intraparietal sulcus (Catmur et al., 2009; Lametti & Watkins, 2016; Rizzolatti et al., 1996; Rizzolatti & Craighero, 2004). This is further supported by a meta-analysis of over 200 fMRI studies (Van Overwalle & Baetens, 2009) comparing the mirror and mentalizing systems and suggesting that the two systems appear to be complementary – rather than one system subserving the other – because they are rarely found to be active together.

The recent findings of Charpentier et al. (2020) also confirm this hypothesis, with distinct neural correlates of imitation and emulation update signals, mapping onto the mirror and mentalizing systems, respectively. One situation that may trigger a transition from the mirror to the mentalizing system is when people observing body motions in other people are deliberating about and inferring the goals of these behavioral executions. Additional evidence supports this functional distinction between the two systems during action understanding, with the mirror neuron system suggested to be involved in automatic action identification, perception and understanding of how actions are implemented and the mentalizing system supporting a more controlled representation of why actions are performed by others and understanding the underlying motives and goals (Spunt & Lieberman, 2012, 2013).

2. The role of mentalizing in strategic social interactions

Our ability to learn from observing others can also be applied to cases where there is a mutual and repeated interaction with one or multiple other agents. In everyday social interactions, we don't only rely on others in order to gather information about the world, but we also engage in strategic interactions in which there is an incentive to infer and exploit another person's knowledge or even an incentive to lie or deceive each other in order to maximize our own rewards and outcomes (D. Lee & Seo, 2016). These behaviors occur in many classic strategy games, such as poker and chess, but also in decisions to cooperate and decisions to engage in prosocial behavior.

An interesting framework to study strategic social interactions in the lab, and to model participants' behavior, is game theory of mind (Camerer, 2003; Yoshida et al., 2008). This framework combines predictions of game theory and optimal behavior together with the social component of repeated mutual interactions between agents. It provides a model of how behavior in such social interactions can be optimized through recursive belief inference, specifically making the assumption that "I represent your value function and thoughts, your representation of mine, your representation of my representation of yours, and so on ad infinitum" (Yoshida et al, 2008). In most strategic interactions, an agent's optimal behavior would be to infer their opponent's degree of sophistication – i.e. levels of recursive beliefs inference – and then play using one degree of sophistication higher than their opponent's. An example to illustrate the different degrees of sophistication is the beauty-contest game (Camerer et al., 2015; Coricelli & Nagel, 2009; Ho et al., 1998; Nagel, 1995). In the original version of the game (Keynes, 1936), competitors have to pick the 6 prettiest faces from 100 photographs. The winner is the competitor whose choice is the closest to the average preferences of all competitors. As Keynes pointed out at the time, "it is not a case of choosing those which are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree – to anticipating what average opinion expects the average opinion to be". In lab experiments, the game has been formalized as "p-Beauty Contests", whereby many players have to choose a number between 0 and 100, and the winner is the person who is closest to the mean of all numbers multiplied by p , with $0 < p < 1$. In the most common setup ($p=2/3$) a player with no degree of sophistication (level 0) will choose a number randomly, so 50 on average (the expectancy of the uniform distribution of possible answers). A level-1 player will think of other players as being level 0 and will choose $2/3$ of 50, so 33 on average. A more sophisticated level-2 player will model others as level-1 and choose $2/3$ of 33, so 22 on average. As people progress in their degree of sophistication, they will eventually reach the Nash equilibrium, which would be to choose 0. Studies have shown that most people in the normal population are level-1 or level-2 (Camerer et al., 2015; Ho et al., 1998; Nagel, 1995). Now the question is: does mentalizing play a role in determining the degree of sophistication of an individual who makes inferences about others? Preliminary evidence using the

beauty contest game suggested that it does, as the dmPFC, a region of the mentalizing network, was found to be more active in participants who engage in high relative to low level of inference (Coricelli & Nagel, 2009).

Additional evidence for the role of mentalizing in strategic social interaction has come from studies using a stag-hunt game. In this task, participants interact with another agent and either decide to hunt a rabbit for a small individual payoff or collaborate with the other agent to hunt a stag for a large payoff (Yoshida et al., 2008). Two types of computational models, both implementing recursive belief inference, were fit to the data. One type was a “fixed model”, in which the degree of sophistication was assumed to be fixed for each agent throughout the task, and the other type was a “theory of mind model”, in which the degree of sophistication is updated after each player’s move. The theory of mind model was found to fit participants’ data best, suggesting a role for mentalizing when players need to dynamically infer their opponent’s strategy and policy. In a subsequent fMRI study, the authors found evidence that signals related to the theory of mind model are represented in the brain (Yoshida et al., 2010). Specifically, the uncertainty associated with the inference about the other agent’s strategy was found in the dmPFC and variations in the other agent’s estimated degree of sophistication were associated with activation in the left dlPFC. It is worthwhile to note that this region is outside the classical mentalizing network, and may instead simply reflect the recruitment of executive processes (Chung et al., 2014) needed for the increasingly complex inference associated with higher sophistication. The social specificity of this process thus remains open to investigation.

In another set of studies, participants played a competitive game called the “inspector game” in pairs (Hampton et al., 2008; C. A. Hill et al., 2017). This game is a variant of “matching pennies”, in which one participant is assigned the role of the employer and the other participant the role of the employee (**Figure 1A**). The employer’s choices are to inspect or not inspect the employee; while the employee’s choices are to work or shirk. The incentives of each player are different, such that the employer has an incentive to not inspect if the employee works or to inspect if the employee shirks, while the employee prefers to shirk if not inspected or work if inspected. Therefore, in order to maximize their outcomes, each participant has to try and predict what the other participant will do next in order to choose the best action for themselves in consequence. To assess the role of mentalizing in such learning, the authors fitted 3 types of models to the data, assuming either (i) no mentalizing, (ii) an elementary form of mentalizing, or (iii) a more sophisticated form of mentalizing (Hampton et al., 2008). The model assuming no mentalizing was a simple reinforcement learning (RL) model, predicting that participants would choose the action that gave the most reward in the recent past. This is equivalent to a level-0 strategy and would be very easy for an opponent to exploit. The model assuming an elementary form of mentalizing employed a strategy called “fictitious play”. This strategy learns about the opponent’s past actions to predict the upcoming action, thus leading the participant to choose according to that predicted action. Using this strategy involves some representation of the other agent’s intentions, similar to a level-1 inference. Finally, the more sophisticated form of mentalizing was called the “influence” learning model, and is equivalent to a level-2 belief inference, in which the participant not only represents their opponent’s past actions like in the fictitious play model, but also tracks how their own actions influence the opponent’s next play (**Figure 1B**). This latter model was found to best explain participants’ behavior on the task, confirming a role for mentalizing and second-order representations of others’ mental states in this strategic social learning task. Regions of the brain’s mentalizing system were also found to track several signals related to the “influence” model. The expected reward associated with the action selected by the participants was encoded in the mPFC at the time of choice, and this signal was better explained by the influence model than by the simpler models, which make different predictions about the expected reward signal. At the time of outcome, activity in the pSTS bilaterally tracked the update in the opponent’s inferred

strategy as predicted by the influence model, and activity in the dmPFC was associated with the degree to which the influence model outperformed the fictitious play model (**Figure 1C**). These fMRI results are therefore consistent with an implication of the mentalizing system, since two key computations were represented in two regions typically involved in mentalizing.

Figure 1. A computational social neuroscience approach to study the role of mentalizing in strategic social interactions. A state-of-the-art combination of three methods can be utilized to demonstrate the role of mentalizing in a strategic social learning task. (**A**) In the competitive “inspector game”, two players make repeated decisions which have different payoffs depending on the choice of the other player. (**B**) Behaviorally, the role of mentalizing can be tested by comparing different computational models of behavior on the task, showing that the strategy requiring the highest degree of mentalizing outperforms strategies that use some or no mentalizing. (**C**) Neurally, the mentalizing network is recruited and tracks relevant computations predicted by the winning model. (**D**) Finally, manipulating activity in parts of the mentalizing network can show its causal involvement in strategic social interactions.

Adapted from Hampton et al. (2008) and C. A. Hill et al. (2017).

An interesting approach was used in a more recent study to determine whether regions of the mentalizing system are causally involved during strategic social reasoning. In this study, the same inspector game task described above (Hampton et al., 2008) was used in combination with theta-burst repetitive transcranial magnetic stimulation (rTMS) to disrupt neural excitability in the rTPJ and examine whether mentalizing processes are impaired as a result (C. A. Hill et al., 2017) (**Figure 1D**). Specifically, participants who received rTMS over the rTPJ, compared to a control group who received rTMS over the vertex, were less likely to switch actions and therefore became more predictable for the opponent to exploit. The ability to reason about the influence of the player’s own actions on the opponent’s response (second-order beliefs), was also found to be significantly reduced in the participants whose rTPJ activity was disrupted. The influence update signal in the rTPJ/pSTS was also

reduced by the stimulation, suggesting that disrupting neural excitability in this region impaired its ability to efficiently encode the necessary social learning signal. Interestingly, the authors also examined long-range effects of rTPJ stimulation by examining neural activity in the dmPFC and vmPFC, as well as functional connectivity between the rTPJ and these regions. Replicating previous findings from Hampton et al. (2008), individual differences in the dmPFC influence update signal were found to predict how likely participants were to rely on the influence over fictitious strategy. This relationship was not affected by rTPJ stimulation, suggesting that the representation of this 2nd order influence model in the dmPFC does not exclusively depend on inputs from the rTPJ. However, functional connectivity between the rTPJ and frontal regions was found to be affected by the stimulation in two ways: (i) reduced functional connectivity between rTPJ and a more dorsal region of the dmPFC at the time of feedback (relative to baseline) and (ii) reduced modulation of functional connectivity between rTPJ and vmPFC by the influence update signal. In summary, this study provided crucial evidence to support a causal role for rTPJ in both the behavioral and neural computations associated with mentalizing; in other words, demonstrating that the mentalizing system is necessary for people to be able to learn how their own actions influence their opponent's future behavior.

Mentalizing regions were also found to play a role in a slightly different type of strategic social interaction – advice giving (Hertz et al., 2017). In a task aimed at eliciting this phenomenon, the participant plays the role of one of two advisers who give advice to a client and have to compete for social influence in order for the client to choose them over the other adviser. Activity in the rTPJ was found to represent whether the participant was chosen by the client or not, which, according to the model, played a role in subsequent strategic influence over the client. Activity in the mPFC encoded relative merit, or advice accuracy relative to the other adviser. In a multi-round economic exchange game (Xiang et al., 2012), one player is an investor deciding which fraction of a \$20 endowment to share with a trustee, the fraction is tripled, and the trustee decides which fraction of that triple amount to repay to the investor. Computational modelling of behavior allowed classifying each investor, who also underwent fMRI, into level-0 (about 50% of investors), level-1 (about 25%) or level-2 (about 25%) players. Different patterns of neural activations were found in the 3 groups; specifically, the rTPJ was found to track 1st order interpersonal prediction errors (when repayments were revealed) more strongly in level-2 compared to level-0 players.

While all the work describe above has focused on human mentalizing and social learning, there is also limited – but nonetheless interesting – evidence that non-human primates can engage in complex strategic social interactions. Two recent studies show that monkeys (i) can predict their opponent's actions and counter a possible exploitation by the opponent (Seo et al., 2014), and (ii) can recursively infer another agent's intentions to decide whether to cooperate or not (Ong et al., 2021).

In the first study (Seo et al., 2014), the authors recorded from monkeys' dmPFC neurons while the animals performed a biased matching pennies game against a computer opponent. In the game both players choose between two targets. If they choose the same target, the monkey wins a point; if they choose different targets the animal either loses a point (risky option) or gets nothing (safe option). The computer opponent's behavior was such that if the monkey chose the risky or safe option more frequently than predicted by the optimal strategy, this behavior was exploited by the computer. Therefore, the monkey has an incentive to not be too predictable. This is exactly what the behavioral model showed. Contrary to the predictions of simple reinforcement learning, the animals' actions did not only depend on their previous outcomes, but also on their previous actions, suggesting that they learn to change their action patterns in order to not be exploited by the computer. In addition, when the computer's actions were predictable, the monkeys were able to exploit them to maximize their payoffs. Neurons in the dmPFC were found to represent the integration of both past outcomes and

past choices, as predicted by higher-order inference about the opponent. Stay versus switch choices were decoded from dmPFC activity, such that the difference in decoding accuracy for switch versus stay choices was predictive of the extent to which the monkey's switch choices deviated from the simple reinforcement learning algorithm. In other words, the more decisions to switch were consistent with strategic thinking, the more dmPFC neurons' activity could decode those decisions.

In another study, Ong and colleagues sought to provide evidence for a TPJ homolog regions in non-human primates and to test its role in strategic interactions (Ong et al., 2021). To do so, they recorded from middle STS (mSTS) neurons while monkeys played a version of the "chicken" game. The game is somewhat similar to the stag-hunt game described above (Yoshida et al., 2008, 2010). In this game, two monkeys are facing each other and moving a joystick to either go straight or yield to the side. If they both go straight they will "crash" into each other and receive no reward. If they both yield, they will get a medium cooperation reward. If one monkey yields and the other goes straight, the monkey who yields gets a small reward and the monkey who goes straight gets a large reward. This task allowed testing whether the monkeys would rather coordinate in order to obtain a cooperative reward, or rather compete to pursue an individual reward at the expense of their opponent. Interestingly the payoffs varied across trials such that a mixed strategy switching between cooperating and competing was optimal. Behavioral results showed that monkeys largely avoided going straight and crashing, suggesting they relied on the other player's behavior to also guide their choice. Specifically, computational models of behavior were tested with different degrees of sophistication. The best-fitting model was found to be the one with the most sophistication, including both a representation of the other monkey's maximum payoffs and learning about the other monkey's strategy via a strategy prediction error (SPE), suggesting an engagement of mentalizing function. Interestingly, mSTS neurons were found to selectively respond to reward obtained cooperatively, but not to rewards obtained selfishly. Some neurons in both mSTS and ACC were also found to encode the opponent's strategy, as predicted by the model. Overall, this very promising line of work suggests that non-human primates also engage in some form of mentalizing, which relies on similar brain networks as humans, to learn from another agent in the context of strategic interactions.

Those two studies provide evidence for a role of mentalizing in social interaction in macaque monkeys. In another study (Devaine et al., 2017), the authors were able to compare mentalizing abilities from seven non-human primate species – specifically lemurs, macaques, mangabeys, orangutans, gorillas and chimpanzees – to test whether mentalizing abilities and degree of sophistication are better explained by social network complexity (as indexed by group size) or by cognitive capacity (as indexed by brain volume). All animals from the seven species (39 in total) played simple dyadic games against artificial players with different degrees of sophistication. Using computational models of behavior, mentalizing abilities on these games were found to be more strongly associated with brain volume rather than social network complexity, suggesting that mentalizing abilities seem to be limited by neurobiological factors and overall cognitive capacity. In addition, comparing the animals' performance with human players, the authors also conclude that great apes' mentalizing abilities still fall short of that of humans.

3. Conclusions and future directions

In this chapter, we have explored evidence suggesting that the brain's mentalizing network – dmPFC, pSTS and TPJ – is involved not only in learning from another agent by inferring their intentions, goals, and beliefs, but can also perform complex computations of mental state inference during strategic social interactions. We highlight that computational models of belief inference, model-based neural activations in the mentalizing network and causal manipulation of these neural computations (**Figure**

1B-D) constitute three valuable methods for examining the role of mentalizing in observational learning and strategic social thinking, especially when used in combination. In this neuro-computational approach, specific mathematical variables predicting behavior are extracted from a computational model and can be directly regressed against brain activity, thus refining our understanding of how exactly a particular process is implemented in the brain. This approach, across the many studies described in this chapter, has provided us with a novel perspective about how different areas of the mentalizing network represent specific mentalizing computations.

Overall, more studies are needed to provide a more integrated account of the computational mechanisms associated with observational and strategic social learning, both at the behavioral and neural level. We would like to highlight some open questions that have yet to be addressed:

- Is mentalizing required for social learning? Recent evidence suggests that by disrupting activity in the brain's mentalizing network (C. A. Hill et al., 2017), as noted, or by studying a clinical population with disrupted mentalizing ability (Rosenthal et al., 2019), we can show that mentalizing is necessary for some particular social learning processes. However, this evidence is still extremely limited and preliminary, thus more studies are needed to generalize these findings to a range of observational learning and strategic social interaction tasks.
- How specific are the computations associated with a particular social learning strategy? In many studies using computational modelling of behavior, one 'winning' model, and the computations predicted by this model, are selected because of their greater explanatory power. However, the specificity of these computations is rarely tested and it is possible that several models would result in the same behavioral and neural computations, thus questioning the specificity of the particular 'winning' model.
- Is mentalizing involved in the arbitration between two social learning strategies? When decisions in a social learning task are found to be a combination of two strategies, it is unclear whether mentalizing abilities, and the mentalizing network, play a role in arbitrating between these strategies. For example, is mentalizing required to decide between relying on imitation versus emulation when learning from another agent? Or to decide between level-1 and level-2 reasoning during strategic interactions?
- How do the different brain regions involved in these processes functionally interact? Very few studies to date have examined functional connectivity between regions of the mentalizing network in the context of social learning computations. Preliminary evidence (C. A. Hill et al., 2017; Zhang & Gläscher, 2020) suggests that connectivity between the TPJ and the prefrontal cortex would be a good candidate to investigate further.
- Does mentalizing play a different role when the goal of social learning is to obtain rewards versus avoid threats? Most studies covered in this chapter examine observational or strategic learning tasks in which the participant's goal is usually to maximize some positive outcomes (e.g. monetary rewards). However, investigations of social learning to avoid threat are much less common (for an example, see Pärnamets et al., 2020), and it is unknown whether and how behavioral and neural computations would differ between positive and negative contexts.

REFERENCES

- Behrens, T. E. J., Hunt, L. T., Woolrich, M. W., & Rushworth, M. F. S. (2008). Associative learning of social value. *Nature*, 456(7219), 24524–24529. <https://doi.org/10.1038/nature07538>
- 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>
- 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>
- Camerer, C. F. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.
- 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>
- Catmur, C., Walsh, V., & Heyes, C. (2009). Associative sequence learning: The role of experience in the development of imitation and the mirror system. *Philosophical Transactions of the Royal Society: Biological Sciences*, 364(1528), 2369–2380. <https://doi.org/10.1098/rstb.2009.0048>
- Charpentier, C. J., Iigaya, K., & O'Doherty, J. P. (2020). A Neuro-computational Account of Arbitration between Choice Imitation and Goal Emulation during Human Observational Learning. *Neuron*, 106(4), 687–699.e7. <https://doi.org/10.1016/j.neuron.2020.02.028>
- 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>
- Chung, H. J., Weyandt, L. L., & Swentosky, A. (2014). The Physiology of Executive Functioning. In S. Goldstein & J. Naglieri (Eds.), *Handbook of executive functioning* (pp. 13–27). Springer.
- Collette, S., Pauli, W. M., Bossaerts, P., & O'Doherty, J. P. (2017). Neural computations underlying inverse reinforcement learning in the human brain. *eLife*, 6, e29718. <https://doi.org/10.7554/eLife.29718>
- Cooper, J. C., Dunne, S., Furey, T., & O'Doherty, J. P. (2012). Human dorsal striatum encodes prediction errors during observational learning of instrumental actions. *Journal of Cognitive Neuroscience*, 24(1), 106–118. https://doi.org/10.1162/jocn_a_00114
- 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>
- Devaine, M., Hollard, G., & Daunizeau, J. (2014). The social Bayesian brain: Does mentalizing make a difference when we learn? *PLoS Computational Biology*, 10(12), e1003992. <https://doi.org/10.1371/journal.pcbi.1003992>
- Devaine, M., San-Galli, A., Trapanese, C., Bardino, G., Hano, C., Saint Jalme, M., Bouret, S., Masi, S., & Daunizeau, J. (2017). Reading wild minds: A computational assay of Theory of Mind sophistication across seven primate species. *PLoS Computational Biology*, 13(11), e1005833. <https://doi.org/10.1371/journal.pcbi.1005833>
- Diaconescu, A. O., Mathys, C., Weber, L. A. E., Daunizeau, J., Kasper, L., Lomakina, E. I., Fehr, E., &

- 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>
- Dunne, S., & O'Doherty, J. P. (2013). Insights from the application of computational neuroimaging to social neuroscience. *Current Opinion in Neurobiology*, 23(3), 387–392. <https://doi.org/10.1016/j.conb.2013.02.007>
- Farmer, A. S., & Kashdan, T. B. (2012). Social anxiety and emotion regulation in daily life: spillover effects on positive and negative social events. *Cogn Behav Ther*, 41(2), 152–162. <https://doi.org/10.1080/16506073.2012.666561>
- Hampton, A. N., Bossaerts, P., & O'Doherty, J. P. (2008). Neural correlates of mentalizing-related computations during strategic interactions in humans. *Proceedings of the National Academy of Sciences*, 105(18), 6741–6746. <https://doi.org/10.1073/pnas.0711099105>
- Hertz, U., Palminteri, S., Brunetti, S., Olesen, C., Frith, C. D., & Bahrami, B. (2017). Neural computations underpinning the strategic management of influence in advice giving. *Nature Communications*, 8(1), 1–12. <https://doi.org/10.1038/s41467-017-02314-5>
- 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, 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, 12722. <https://doi.org/10.1038/ncomms12722>
- Ho, T. H., Cambrer, C., & Weigelt, K. (1998). Iterated Dominance and Iterated Best Response in Experimental “p-Beauty Contests.” *American Economic Review*, 88(4), 947–969.
- Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money*. Harcourt Brace and Co.
- Lametti, D. R., & Watkins, K. E. (2016). Cognitive neuroscience: The neural basis of motor learning by observing. *Current Biology*, 26(7), R288–R290. <https://doi.org/10.1016/j.cub.2016.02.045>
- Lee, D., & Seo, H. (2016). Neural basis of strategic decision making. *Trends in Neurosciences*, 39(1), 40–48. <https://doi.org/10.1016/j.tins.2015.11.002>
- Lee, S. W., 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>
- Nagel, R. (1995). Unraveling in Guessing Games: An Experimental Study. *The American Economic Review*, 85(5), 1313–1326. <https://doi.org/http://www.aeaweb.org/aer/>
- Ong, W. S., Madlon-Kay, S., & Platt, M. L. (2021). Neuronal correlates of strategic cooperation in monkeys. *Nature Neuroscience*, 24(1), 116–128. <https://doi.org/10.1038/s41593-020-00746-9>
- Pärnamets, P., Espinosa, L., & Olsson, A. (2020). Physiological Synchrony Predicts Observational Threat Learning in Humans. *Proceedings of the Royal Society B*, 287(20192779). <https://doi.org/10.1098/rspb.2019.2779>
- Pärnamets, P., & Olsson, A. (2020). Integration of social cues and individual experiences during instrumental avoidance learning. *PLoS Computational Biology*, 16(9), 1–24. <https://doi.org/10.1371/JOURNAL.PCBI.1008163>
- Rizzolatti, G., & Craighero, L. (2004). The mirror-neuron system. *Annual Review of Neuroscience*,

- 27(1), 169–192. <https://doi.org/10.1146/annurev.neuro.27.070203.144230>
- Rizzolatti, G., Fadiga, L., Gallese, V., & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3(2), 131–141. [https://doi.org/10.1016/0926-6410\(95\)00038-0](https://doi.org/10.1016/0926-6410(95)00038-0)
- Rosenthal, I. A., Hutcherson, C. A., Adolphs, R., & Stanley, D. A. (2019). Deconstructing Theory-of-Mind Impairment in High-Functioning Adults with Autism. *Current Biology*, 29(3), 513–519. <https://doi.org/10.1016/j.cub.2018.12.039>
- Seo, H., Cai, X., Donahue, C. H., & Lee, D. (2014). Neural correlates of strategic reasoning during competitive games. *Science*, 346(6207), 340–343. <https://doi.org/10.1126/science.1256254>
- Spunt, R. P., & Lieberman, M. D. (2012). Dissociating modality-specific and supramodal neural systems for action understanding. *Journal of Neuroscience*, 32(10), 3575–3583. <https://doi.org/10.1523/JNEUROSCI.5715-11.2012>
- Spunt, R. P., & Lieberman, M. D. (2013). The Busy Social Brain: Evidence for Automaticity and Control in the Neural Systems Supporting Social Cognition and Action Understanding. *Psychological Science*, 24(1), 80–86. <https://doi.org/10.1177/0956797612450884>
- Suzuki, S., Harasawa, N., Ueno, K., Gardner, J. L., Ichinohe, N., Haruno, M., Cheng, K., & Nakahara, H. (2012). Learning to simulate others' decisions. *Neuron*, 74(6), 1125–1137. <https://doi.org/10.1016/j.neuron.2012.04.030>
- Van Overwalle, F., & Baetens, K. (2009). Understanding others' actions and goals by mirror and mentalizing systems: A meta-analysis. *NeuroImage*, 48(3), 564–584. <https://doi.org/10.1016/j.neuroimage.2009.06.009>
- Xiang, T., Ray, D., Lohrenz, T., Dayan, P., & Montague, P. R. (2012). Computational phenotyping of two-person interactions reveals differential neural response to depth-of-thought. *PLoS Computational Biology*, 8(12), e1002841. <https://doi.org/10.1371/journal.pcbi.1002841>
- Yoshida, W., Dolan, R. J., & Friston, K. J. (2008). Game theory of mind. *PLoS Computational Biology*, 4(12), e1000254. <https://doi.org/10.1371/journal.pcbi.1000254>
- Yoshida, W., Seymour, B., Friston, K. J., & Dolan, R. J. (2010). Neural mechanisms of belief inference during cooperative games. *The Journal of Neuroscience*, 30(32), 10744–10751. <https://doi.org/10.1523/JNEUROSCI.5895-09.2010>
- Zhang, L., & Gläscher, J. (2020). A brain network supporting social influences in human decision-making. *Science Advances*, 6(34), 1–20. <https://doi.org/10.1126/sciadv.abb4159>