

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Help me help you: A computational model for goal inference and action planning

Permalink

<https://escholarship.org/uc/item/2rk2b4sm>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors

Kumar, Abhilasha

Steyvers, Mark

Publication Date

2023

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Help me help you: A computational model for goal inference and action planning

Abhilasha A. Kumar
Bowdoin College
a.kumar@bowdoin.edu

Mark Steyvers
University of California, Irvine
mark.steyvers@uci.edu

Abstract

Helping is an inherently cooperative behavior, but the cognitive mechanisms underlying this behavior remain relatively underexplored. In this paper, we introduce a novel gamified paradigm for understanding a variety of cognitive behaviors associated with helping. Principals are assigned secret goals in a block-based grid (e.g., move all blue blocks to room C), and helpers can either pass their turn or make a move that could help the principal. We show that principals make useful and *pragmatic* first moves and helpers accumulate evidence over time before initiating a helpful move. We also introduce a preliminary set of computational models based on recursive pragmatic inference and utility maximization that attempt to account for these behavioral findings.

Keywords: helping; assistance; pragmatic inference; Bayesian reasoning

Introduction

Imagine that you are preparing for a long-distance move. As you are packing, a helper that you hired to assist with the move arrives. But there’s a catch: the helper does not speak your language and can only observe and learn from your actions. How do you communicate your goals to the helper, in a way that *helps them* assist you effectively? Now consider the perspective of the helper. Armed with the knowledge that some things need to be moved but no clear goal, they heavily rely on your actions to understand your intentions. In this game of charades, should the helper focus on your actions, or also why you may have chosen one move over another?

The ability to collaborate and help each other is considered central to the success of our species (Tomasello et al., 2005a,b). Helping is an interesting behavior from a cognitive perspective, because it not only requires understanding the perspective of another person’s beliefs and goals, but also coordinating actions that serve their goals. Moreover, from the perspective of the person being helped, prioritizing actions that effectively communicate the goal early on may be critical, and in turn require understanding the perspective of the helper. This ability to infer what others are thinking is broadly referred to as theory of mind (ToM; Premack & Woodruff, 1978). ToM has been suggested to underlie social interaction and cooperation, and is assumed to develop early on among human infants (Warneken & Tomasello, 2006, 2007). While there is a rich literature on the development of ToM and altruistic helping, relatively little is known about ToM and its connection to helpful behaviors among human adults. Furthermore, the ability to infer or derive conclusions about an

agent’s goals or beliefs that go beyond the available evidence, i.e., pragmatic inference, appears to be a critical component of ToM (Bergen et al., 2016; Goodman & Stuhlmüller, 2013), but the extent to which pragmatic inference contributes to helpful behaviors is a relatively unexplored question. Finally, although language is ubiquitous in several contexts, learning from observation or “non-verbal task learning” (NTL; Barbu et al., 2010; Bentivegna et al., 2004; Kunda, 2019) is a critical component of social learning, and inferring task goals via multiple modalities (visual, linguistic, etc.) is an active area of research (Hinrichs & Forbus, 2014; Kirk et al., 2016). Therefore, specifying the link between pragmatic inference and helping behaviors in non-verbal contexts will not only advance research on social cognition, but may also inform the development of artificial assistive agents (Bobu et al., 2020; Puig et al., 2020)

Consider the block-based environment shown in Figure 1. The space is divided into three “rooms” (A, B, and C), each of which are further subdivided into two sub-rooms (e.g., A1, A2). There are blocks of different colors in each room. Imagine that one agent (the “principal”) has been assigned a particular goal (e.g., move all *blue* blocks to room B1) and can move one block on each turn to achieve the goal. The second agent (the “helper”) has no knowledge of this goal, but can choose to help the principal on each turn by moving a block, or passing their turn. What kind of strategy should the principal or helper employ to ensure that the goal is achieved in a collaborative manner? For example, as shown by the *aqua* and *purple* arrows, should the principal move the blue block in A1 to C1, or the green block in B2 that is covering a blue block to B1? From a purely goal-based perspective, both moves bring the principal one step closer to achieving the goal. However, one of these moves (*aqua*) is more likely to clearly communicate the goal to the helper than the other, and is therefore a more *pragmatic* move. From the helper’s perspective, if they see the *aqua* move, is that sufficient evidence for them to infer the principal’s goal, or do they need more information before they can confidently assist the principal by making the *orange* move?

In this paper, we investigated whether inverse planning and pragmatic inference contribute to helpful behaviors in humans, by proposing and evaluating a model of Bayesian social reasoning within a utility-based framework (based on prior work by Baker et al., 2009; Ho et al., 2021; Goodman &




Figure 1: A block configuration in the current study. Rooms A, B, and C are divided into subrooms (A1, A2, etc.). The principal is assigned a secret goal (e.g., move all blue blocks to room C) and can move one block per turn. The helper can pass their turn or move a block, which could be helpful to the principal. The aqua and purple arrows show potential principal moves, and the orange arrow shows a move considered by the helper after the aqua move has been made.

Frank, 2016). Our approach is related to previous approaches discussed above, but asks a novel question: do humans employ pragmatic inference when engaging in helpful behaviors? We situate this question in the context of a collaborative two-player assistance game shown in Figure 1, where principals are assigned secret goals on a two-dimensional grid of colored blocks. On a given turn, principals move blocks to achieve their goal, and helpers can either pass their turn or assist the principal in achieving the goal by making a move. Grid-based stimulus paradigms such as ours have been extensively used in the planning and pragmatic inference literature (e.g., Ho et al., 2021; Krych-Appelbaum et al., 2007; McCarthy et al., 2021). Such paradigms allow for fine-grained experimental control over the scope of actions and hypotheses and are an ideal framework for eliciting a broad array of behaviors which lend themselves well to computational modeling.

In the following sections, we describe our methodology and discuss the behavioral findings from a pre-registered experiment (<https://osf.io/q2p6b>) conducted online, where principals and helpers played several rounds of assistance games with different goals. We then describe a novel computational modeling framework that seeks to explain the behavioral patterns. In doing so, we motivate novel intuitions for how humans seek and provide help in collaborative contexts.

Methods

Participants

Two hundred participants were recruited in dyads from Amazon Mechanical Turk (with “Masters” level qualification) as well as the psychology subject pool at XYZ. Participants from MTurk received \$6 for their participation, and students were either compensated in course credit or \$5 gift cards. 89 dyads completed the game, and based on exclusion criteria (incom-

plete games, too many moves, etc.), we excluded 9 additional dyads, leading to a final sample of 80 dyads ($N = 160$), with 10 games per dyad, i.e., 800 total games.

Design and procedure

The game was programmed in nodeGame (Baliotti, 2017) and played online. All games were based on the block configuration shown in Figure 1, with three rooms (A, B, and C), each further sub-divided into two smaller rooms. Each configuration consisted of ten colored blocks of three types (red, blue, and green). Roles (principal and helper) were randomly assigned at the start of the game and remained constant for the duration of the game. Players were encouraged to be cooperative (i.e., split the work as much as possible). Each dyad played ten rounds with different goals, where goals could be of five different types: moving a block from one room to another (move), covering all blocks of a certain color with another block (cover), uncovering all blocks of a certain color (uncover), clearing a particular room by moving all blocks inside a room out of it (clear), and filling up a particular room by covering all possible white spaces in the room (fill). Goals were pseudo-randomly generated for each dyad, such that each dyad received two goals of each type (move, cover, uncover, fill, clear) presented in a randomized order. The goals were randomly chosen from a “goalspace” of 48 total goals, and a goal was resampled from the same subtype if it required less than 3 moves to complete based on the randomly generated starting configuration. The starting configuration of the grid was randomized across participant dyads.

Behavioral results

To analyze the moves made by the principal and helper, we first computed the minimum number of moves required to achieve different goals across the games. For example, in Figure 1, moving all blue blocks to C requires a minimum of 8 moves, which includes moving 5 unobstructed blue blocks,

moving the 1 green block that is covering the blue block in B2, and then moving the 2 now-unobstructed blue blocks. Next, we computed a *utility* measure for each move based on whether it advanced the goal, and classified moves into three categories: useful, inconsequential, and harmful. A move was considered *useful* when it decreased the minimum number of moves required to achieve the goal, *inconsequential* when the move had no immediate impact on the minimum number of moves ¹, and *harmful* when the move increased the minimum number of moves required to achieve the goal. We then examined the types of moves made by the principal and helper over the course of the game.

Principals make *useful* moves

Figure 2 displays the types of moves made by principals across all games. As shown, principals made significantly more useful moves compared to harmful or inconsequential moves, which was confirmed by a significant effect of move type, $F(2,130) = 162.85, p < .001$.

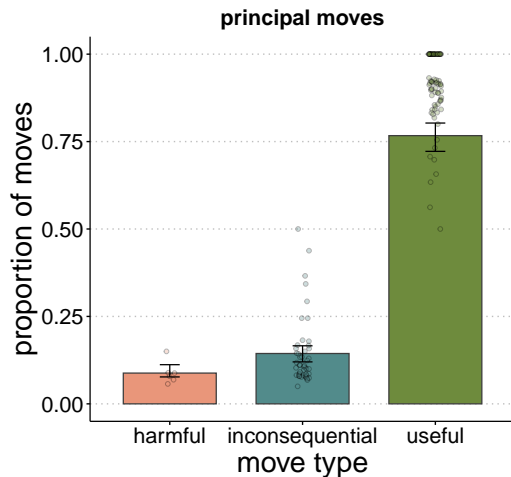


Figure 2: Types of moves made by the principal. Error bars denote 95% bootstrapped confidence intervals.

Principals make *pragmatic* first moves

Next, we examined whether the moves made by the principal were *pragmatic* by any means. To understand what counts as a pragmatic move, consider two moves shown via the curved arrows in Figure 1: moving a blue block from A1 to C1 (aqua arrow) and moving the green block in B2 to uncover the blue block underneath (purple arrow). From a utility perspective, both these moves are “useful”, as they advance the goal of moving all blue blocks to room C. However, the first move intuitively feels like a better move than the second. To formalize this intuition, we evaluated whether a given move eliminated other possible goals from the hypothesis space considered by a potential helper. Specifically, when the utility of two

¹Note that even though inconsequential moves do not advance or harm the minimum moves in that instance of the game, they do present an opportunity cost, i.e., more moves will have to be made in total to achieve the goal




Figure 3: Mean rank of useful first moves made by the principal; lower ranks indicate a move that eliminates more goals, i.e., a *pragmatic* move. Error bars denote 95% bootstrapped confidence intervals.

moves is identical and positive, a pragmatic move is one that minimizes the total number of possible goals from a given goalspace. For example, in Figure 1, moving the blue block from A1 to C1 (aqua move) is consistent with the following six goals: (1) moving blue blocks to C, (2) clearing A1, (3) filling C1, (4) covering all red blocks, (5) clearing A, and (6) move blue blocks to C1. On the other hand, moving the green block above the blue block in B2 to B1 (purple move) is consistent with the following 12 goals: (1) moving blue blocks to A1, (2) moving green blocks to B1, (3) moving blue blocks to C2, (4) moving blue blocks to A, (5) moving blue blocks to A2, (6) moving blue blocks to C1, (7) moving blue blocks to C, (8) uncovering all blue blocks, (9) clearing B2, (10) moving blue blocks to B1, (11) covering all red blocks, and (12) filling B1. Therefore, even though both moves are “useful”, moving the blue block from A1 to C1 is a more *pragmatic* move, compared to the moving the green block from B2 to B1, because it eliminates more goals.

To understand whether principals were making pragmatic moves in the assistance game, we ranked each move based on whether it was a useful move and how many *other* goals it served. Then, we computed the rank of the actual empirical move made by the principal, and compared this rank to the mean rank of 1000 randomly sampled useful moves (with replacement). In this way, we were able to evaluate whether the useful moves made by the principal were more pragmatic than what would be expected by chance. Figure 3 displays the overall patterns. As shown, the mean rank of useful first moves was significantly lower than what would be expected by chance. This effect was confirmed by a significant effect of rank (empirical vs. random), $\chi^2(1, N = 1634) = 9.98, p = .002$. Therefore, when principals made useful first moves, these moves were likely to be pragmatic moves that eliminated more goals and assisted the helper.




Figure 4: Types of moves made by the helper. Error bars denote 95% bootstrapped confidence intervals.

Helpers tend to pass their turn a quarter of the time

We also examined the different types of moves made by the helper during the course of the game. As a reminder, after the first principal move, the helper had the option to either pass their turn or help the principal by moving a block. Similar to our analysis of principal moves, we classified each action taken by the helper into whether it was useful, harmful, or inconsequential. In addition, although a “pass” move is technically an inconsequential move because it does not change the number of moves required to complete the goal, we considered pass moves as a special case in these analyses. Figure 4 displays the overall pattern. Helpers generally made more useful moves, compared to other moves, $F(3,258) = 85.71$, $p < .001$. However, helpers also chose to pass nearly a quarter of the time.

Helpers accumulate evidence for goals over time

We next looked at whether this behavior of passing their turn varied over the course of the game. As shown in Figure 5, over 40% of the *first* moves made by helpers were pass moves. The frequency of passing decreased after the first move, and helpers began to make useful moves over time. This suggests that helpers preferred to gather more evidence about the underlying goal in the initial stages of the game, before they decided to help the principal.

Summary of behavioral results

Overall, our analysis of the gameplay showed that principals and helpers successfully cooperated to achieve goals in the games. We found that principals were able to understand the goal and made pragmatic first moves that best communicated the secret goal to the helper. We also found that helpers were hesitant to actively provide help during the initial stages of the game, and instead chose to pass their first turn. This behavior




Figure 5: Proportion of helper moves as a function of move order. Error bars denote 95% bootstrapped confidence intervals.

likely reflects uncertainty about the goal space and/or the secret goal. In the following section, we introduce and evaluate a series of candidate models that seek to explain the behavior of the principal and helper in the game, and therefore provide a cognitively motivated computational perspective on helping behaviors.

Computational modeling

Our computational framework formalizes helpful actions as a combination of Bayesian pragmatic inference and utility maximization. We first describe how the agents and actions within the assistance game map are encoded within a Bayesian framework. Next, we describe a recursive reasoning process that underlies the actions of agents in this environment, although we focus on relatively shallow levels of recursive reasoning here. Finally, we compare a series of computational models in the extent to which they account for behavior in the current experiment, and specifically focus on *first* moves made by the agents.

Overall framework

We assume that each assistance game begins with a configuration c and a bounded goal space G consisting of 48 possible goals. The agents, the principal and helper, take turns to achieve a secret goal, $g \in G$, which is known only to the principal. For any given configuration, M is the set of all possible moves. We further assume that the helper and principal recursively reason about the other agent’s intentions and make inferences based on their actions.

We define $\text{MINIMUMMOVES}(g, c)$ as the minimum number of moves required to achieve a goal g for a configuration c , i.e., the sum of total number of moves required to move any unobstructed blocks that would achieve the goal and the moves required to uncover the obstructed blocks and move them to the desired location. We also define the function $s(m, c)$, which returns a new configuration after a move m has

been made on configuration c . Finally, we define $u(m|g, c)$, which represents the utility for move m given goal g and configuration c by computing the difference between the minimum moves before and after the move via

$$u(m|g, c) = \text{MINIMUMMOVES}(g, c) - \text{MINIMUMMOVES}(g, s(m, c)) \quad (1)$$

As described in the behavioral analyses, the utility of a move is 1 (useful), 0 (inconsequential), or -1 (harmful) for a given goal and configuration.

Baseline principal model Within this framework, one simple model of the principal is to maximize the utility of a given move under the specific goal g and configuration c . However, we assume that there is some decision noise in the process of selecting moves to allow for the possibility of moves that are inconsequential or (occasionally) even harmful. We model the probability of move selection by a softmax,

$$p(m_i|g, c) = \left(e^{\frac{U(m_i|g, c)}{\tau_d}} \right) / \sum_j \left(e^{\frac{U(m_j|g, c)}{\tau_d}} \right) \quad (2)$$

Baseline helper model We assume a helper that has two tasks: inferring a goal distribution, and then evaluating how different moves are consistent with the most likely goals. Specifically, after receiving a specific move $m_i \in M$ from the principal, the helper first computes the probability of each goal $g_k \in G$ via

$$p(g_k|m_i) \propto p(g_k) \prod_{m \in H} P(m_i|g_k) \quad (3)$$

where H denotes the set of moves so far. Next, the helper evaluates the expected utility of a move m_i based on the likelihood of different goals via:

$$E(m_i) = \sum_{g_k \in G} p(g_k|m_i) u(m_i|g_k) \quad (4)$$

and then computes the probability of making each move via:

$$p(m_i) = \left(e^{\frac{E(m_i)}{\tau_h}} \right) / \sum_j \left(e^{\frac{E(m_j)}{\tau_h}} \right) \quad (5)$$

where τ_h denotes the temperature parameter for the helper, which estimates the noise in selecting the appropriate move. In this model, we consider passing as an inconsequential move (i.e., $u = 0$), such that passing is possible but not “useful” in advancing the goal.

Careful helper model While the model above is a reasonable approximation of the cognitive processes that a helper may employ, it is possible that uncertainty may also arise in the decision to make a move versus pass a turn. Therefore, we evaluate a second model that assigns a probabilistic decision rule to making a move. Specifically, after inferring a goal distribution via Equation 3, the helper assesses the difference

between the two most likely goals, and assigns some probability to passing based on the difference in the goals scaled with a parameter α , as follows:

$$p(\text{pass}) \propto \ln \frac{1 - \alpha(g_{\max} - g_{\text{next}})}{\alpha(g_{\max} - g_{\text{next}})} \quad (6)$$

Thus, the higher the difference between the most and next most likely goal, the lower the probability of passing, which in turn influences making a move, and probabilities obtained from Equation 5 are renormalized accordingly.

Pragmatic principal model We model the pragmatic principal by assuming that they make inferences about how their actions may in turn inform the inferences made by the helper about the underlying goal. Therefore, the pragmatic principal attempts to choose actions that would increase the likelihood of inferring the *true* goal, g , via:

$$p(m_i|g) = \left(e^{\frac{p(g_k=g|m_i)}{\tau_p}} \right) / \sum_j \left(e^{\frac{p(g_k=g|m_j)}{\tau_p}} \right) \quad (7)$$

where τ_p denotes the temperature parameter for the pragmatic principal.

Qualitative model performance

First, we qualitatively examined the probabilities of different moves from Figure 1 based on the baseline and pragmatic principal models². As shown in Table 1, while the baseline principal considers each useful move (aqua and purple) as equal, the pragmatic principal assigns a higher probability to the pragmatic move (aqua), consistent with the behavior of the principal in the game.

Table 1: Move preference in principal and helper models

principal models		
	aqua move	purple move
baseline	0.023	0.023
pragmatic	0.007	0.005
helper models		
	pass move	orange move
baseline	0.001	0.045
careful	0.976	0.001

We also examined how the two helper models evaluated different moves after the first principal move. For example, after the principal makes the aqua move in Figure 1, the helper might consider passing their turn or moving a block. We evaluated the probability of passing versus making one such move, the orange move, i.e., moving the blue block in A2 to C1, in the two helper models³. As shown in Table 1, while the baseline helper prioritizes making the orange move, the careful helper prioritizes passing, even though the orange

² τ_d , τ_p , and τ_h were set to .2 for this demonstration

³ τ_d and τ_h were set to .2, and α was set to 1 for this demonstration




Figure 6: Simulated proportion of principal moves. Error bars denote 95% bootstrapped confidence intervals.

move is useful. This is consistent with the actual behavior in the game, where helpers chose to pass their initial turns, even when they could have made useful moves.

Quantitative model performance

Finally, we obtained best-fitting parameter values at the participant level for the first moves made by principals and helpers, and then simulated first moves made by agents based on these parameters. As shown in Figure 6, simulations from both principal models mirrored the behavioral pattern in Figure 2. Additionally, we found that the pragmatic principal model assigned lower ranks to empirical moves and preferred them over other moves, compared to the baseline model, confirmed by a significant effect of model on ranks, $t(1419.33) = -18.08$, $p < .001$. Similarly, as shown in Figure 7, while the baseline helper prioritized useful first moves over all other moves, the careful helper prioritized pass moves over other first moves, therefore better mirroring the behavioral pattern in Figure 5. Additionally, the careful helper also assigned lower ranks to pass moves compared to the baseline helper, confirmed by a significant interaction between model and move type, $\chi^2(3, N = 1600) = 300.53$, $p < .001$.

Discussion

In this paper, we introduced a novel experimental paradigm to study helping behaviors and evaluated a computational modeling framework that conceptualized helping and being helped as an inverse reasoning problem. We now discuss the key insights from the behavioral patterns and computational modeling of these assistance games.

Overall, both agents (principals and helpers) were able to successfully cooperate and make useful moves that advanced the goal for a given configuration. Importantly, we found that the first moves chosen by principals were not only useful, but also *pragmatic*. Specifically, when faced with a choice between several useful moves, all of which would advance the secret goal, principals chose moves that eliminated more

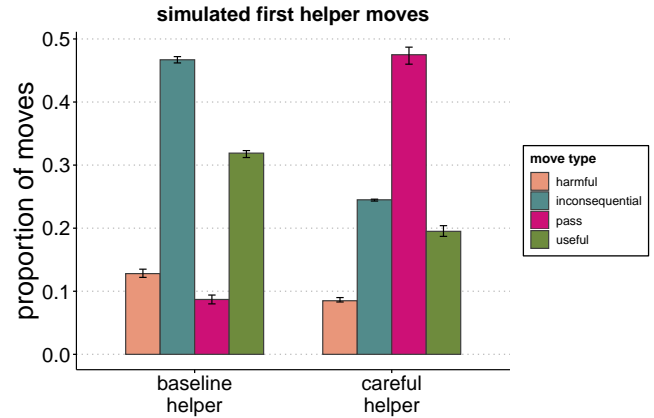


Figure 7: Simulated proportion of helper moves. Error bars denote 95% bootstrapped confidence intervals.

goals and were therefore more likely to nudge the helper in the right direction. These behavioral patterns were also supported by a computational model that framed this problem of being helped as an inverse reasoning problem, where (pragmatic) principals reasoned about which move would have the highest likelihood of communicating the secret goal to the helper. This behavior is interesting because the principals could have chosen randomly among the useful moves (which would still achieve the goal), and suggests that principals were pragmatically reasoning during the game. These results are consistent with prior work on pedagogical and communicative demonstrations (Ho et al., 2021; Shafto et al., 2014), but extend it to contexts where agents are not explicitly asked to “show”, “demonstrate”, or “teach”. This tendency to engage in demonstrative behaviors without explicit instructions may be related to evolutionary benefits of engaging in cooperative behaviors (Bshary & Bergmüller, 2008).

On the other hand, despite *pragmatic* first moves made by principals, helpers preferred to wait and evaluate the evidence before choosing to make moves. This suggests that helpers were possibly weighing the likelihood of different goals in their decision to pass or act, and chose to act only when they were confident of their inference. It is also possible that helpers were unclear about the hypothesis/goal space, given the ill-defined problem space, and we hope to investigate this issue in the future, in addition to exploring the entire gameplay. Overall, these patterns are consistent with prior work on information sampling, where agents tend to gather evidence before committing to a plan (Ma et al., 2021), and may also be related to the helper’s trust in the knowledge and helpfulness of the principal (Eaves Jr & Shafto, 2012). In future work, we hope to more carefully evaluate our model predictions for the entire duration of the game as well as extend this basic experimental and computational framework to address broader questions about the cognitive underpinnings of helpful behaviors.

References

- Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, *113*(3), 329–349.
- Baliotti, S. (2017). nodegame: Real-time, synchronous, online experiments in the browser. *Behavior Research Methods*, *49*(5), 1696–1715.
- Barbu, A., Narayanaswamy, S., & Siskind, J. M. (2010). Learning physically-instantiated game play through visual observation. *2010 IEEE International Conference on Robotics and Automation*, 1879–1886.
- Bentivegna, D. C., Atkeson, C. G., & Cheng, G. (2004). Learning tasks from observation and practice. *Robotics and Autonomous Systems*, *47*(2-3), 163–169.
- Bergen, L., Levy, R., & Goodman, N. (2016). Pragmatic reasoning through semantic inference. *Semantics and Pragmatics*, *9*.
- Bobu, A., Bajcsy, A., Fisac, J. F., Deglurkar, S., & Dragan, A. D. (2020). Quantifying hypothesis space misspecification in learning from human–robot demonstrations and physical corrections. *IEEE Transactions on Robotics*, *36*(3), 835–854. doi: 10.1109/TRO.2020.2971415
- Bshary, R., & Bergmüller, R. (2008). Distinguishing four fundamental approaches to the evolution of helping. *Journal of evolutionary biology*, *21*(2), 405–420.
- Eaves Jr, B. S., & Shafto, P. (2012). Unifying pedagogical reasoning and epistemic trust. *Advances in child development and behavior*, *43*, 295–319.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Science*, *20*(11), 818–829.
- Goodman, N. D., & Stuhlmüller, A. (2013). Knowledge and implicature: Modeling language understanding as social cognition. *Topics in cognitive science*, *5*(1), 173–184.
- Hinrichs, T. R., & Forbus, K. D. (2014). X goes first: Teaching simple games through multimodal interaction. *Advances in Cognitive Systems*, *3*, 31–46.
- Ho, M. K., Cushman, F., Littman, M. L., & Austerweil, J. L. (2021). Communication in action: Planning and interpreting communicative demonstrations. *Journal of Experimental Psychology: General*.
- Kirk, J., Mininger, A., & Laird, J. (2016). Learning task goals interactively with visual demonstrations. *Biologically Inspired Cognitive Architectures*, *18*, 1–8.
- Krych-Appelbaum, M., Law, J. B., Jones, D., Barnacz, A., Johnson, A., & Keenan, J. P. (2007). “i think i know what you mean”: The role of theory of mind in collaborative communication. *Interaction Studies*, *8*(2), 267–280.
- Kunda, M. (2019). Nonverbal task learning. *Proceedings of the Seventh Annual Conference on Advances in Cognitive Systems*.
- Ma, I., Ma, W. J., & Gureckis, T. M. (2021). Information sampling for contingency planning. In *Proceedings of the 43rd annual conference of the cognitive science society*.
- McCarthy, W. P., Hawkins, R. D., Wang, H., Holdaway, C., & Fan, J. E. (2021). Learning to communicate about shared procedural abstractions. *arXiv preprint arXiv:2107.00077*.
- Premack, D., & Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, *1*(4), 515–526. doi: 10.1017/S0140525X00076512
- Puig, X., Shu, T., Li, S., Wang, Z., Liao, Y.-H., Tenenbaum, J. B., ... Torralba, A. (2020). *Watch-and-help: A challenge for social perception and human-ai collaboration*. arXiv. Retrieved from <https://arxiv.org/abs/2010.09890> doi: 10.48550/ARXIV.2010.09890
- Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive psychology*, *71*, 55–89.
- Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005a). In search of the uniquely human. *Behavioral and Brain Sciences*, *28*(5), 721–735. doi: 10.1017/S0140525X05540123
- Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005b). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, *28*(5), 675–691. doi: 10.1017/S0140525X05000129
- Warneken, F., & Tomasello, M. (2006). Altruistic helping in human infants and young chimpanzees. *Science*, *311*(5765), 1301–1303. doi: 10.1126/science.1121448
- Warneken, F., & Tomasello, M. (2007). Helping and cooperation at 14 months of age. *Infancy*, *11*(3), 271–294. doi: <https://doi.org/10.1111/j.1532-7078.2007.tb00227.x>