

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Guiding Inference: Signaling intentions using efficient action

Permalink

<https://escholarship.org/uc/item/05z6b1wh>

Journal

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

Authors

Royka, Amanda L

Török, Georgina

Jara-Ettinger, Julian

Publication Date

2023

Peer reviewed

Guiding Inference: Signaling intentions using efficient action

Amanda Royka (amanda.royka@yale.edu)^{*1}, Georgina Török (torok@mpib-berlin.mpg.de)^{*2}, Julian Jara-Ettinger (julian.jara-ettinger@yale.edu)¹

¹Department of Psychology, Yale University, ²Max Planck Institute for Human Development; * Equal contribution

Abstract

People have a remarkable capacity to infer others' goals and intentions based on how they behave. Yet, humans are also motivated to ensure that others can infer their mental states easily and accurately. Past work has shown that people achieve this by introducing inefficiencies to their behavior, which reveal its underlying goal (e.g., exaggerating one's movements so as to make their purpose obvious). We hypothesized that inefficiency is not a constitutive feature of signaling, and that people will often signal their goals and intentions solely through efficient action. We test this idea in a signal-design experiment where participants need to reach an instrumental goal while also making that goal as inferable as possible. In line with our hypothesis, people shape their behavior to increase inferability without jeopardizing efficiency (Experiment 1). Using a computational model, we show that these efficient signals are well-designed to guide observers' inferences about the relevant instrumental goal. Moreover, observers' intuitions about which paths were produced to signal correlate with the proportion of times that the paths were generated in the signaling condition of our first experiment (Experiment 2). Our results show that humans not only exploit opportunities to reveal their goals without deviating from efficient action, but that these signals allow observers to understand the instrumental and signaling goals underpinning the movement.

Keywords: goal inference; signaling; theory of mind

Introduction

Much of our social cognition is built on a basic capacity to make sense of other people's behavior. Is she going to get lunch from the Turkish or Indian food truck? Is he waving at me or is he swatting a fly? Is my teammate going to pass the ball left, or right? This ability to interpret behavior is supported by our *Theory of Mind*—our basic capacity to understand other people's actions in terms of unobservable mental states such as beliefs, desires, goals, and intentions (Gopnik & Meltzoff, 1994). Consistent with this, researchers have found that this capacity emerges early in infancy (Gergely & Csibra, 2003; Liu et al., 2017) and supports a range of complex human behaviors including language use (Goodman & Frank, 2016; Jara-Ettinger & Rubio-Fernandez, 2021), pedagogy (Gweon, 2021), and moral reasoning (Young et al., 2007).

While the ability to infer mental states is critical to social cognition, this picture neglects a major feature of our mental life. Humans are not just "mindreaders," making inferences about some disinterested third party: We also shape our behavior to make our own mental states better understood. For example, when completing cooperative tasks, people exaggerate their behavior to make their goal more obvious to

observers (e.g., when reaching towards one of two targets, moving in a higher or wider path than you otherwise would if completing the task alone; McEllin et al., 2018; Pezzulo & Dindo, 2011; Sacheli et al., 2013; Vesper & Richardson, 2014).

Evidence of people shaping their behavior to make themselves understood has also been documented in the domain of gesture. For gestures to be successful, observers must recognize their communicative goal (otherwise, the gesture will not convey its intended message). Recent accounts have hypothesized that the recognition of gesture is underpinned by an expectation that communicative actions will be inefficient such that they reveal that the actions are not directed towards non-communicative goals in the environment. Indeed, people infer that a movement is more likely to be communicative when they detect inefficiencies shaped to reveal the absence of a world-directed goal (Royka et al., 2022). Conversely, when creating novel symbolic gestures, people spontaneously generate signals that reveal the lack of a non-communicative goal through the introduction of inefficiencies in the movement (Royka et al., 2021; Scott-Phillips et al., 2009). Finally, inefficiencies are also a signature of how people reveal and recognize intentions to communicate, particularly when the movement risks looking like it has a world-directed goal (e.g., exaggerating how to remove a pen lid; Brand et al., 2002; Trujillo et al., 2018), and a feature of how we help reveal our own knowledge (Ho et al., 2016).

Taken together, this work shows the important role that inefficiency plays in creating actions that make our mental states more easily understood. This and related work has therefore led to the idea that inefficiency is central to signaling behavior (Dockendorff et al., 2019; Pezzulo et al., 2013, 2019; Royka et al., 2022). While inefficiency is an important and flexible signaling strategy, this work leaves open the question of why inefficiency is so prevalent in signaling behavior. One possibility is that inefficiency is a common, but context-sensitive strategy for making our behavior more legible for observers because it eases their inferential burden. Alternatively, inefficiency could be an intrinsic property strictly necessary for signaling behavior independent of context (e.g., perhaps also serving the role of alerting observers that the movement contains relevant information for them; Wilson & Sperber, 2004).

Importantly, however, people can signal their goals through movement that does not deviate from efficient action (Ho et al., 2016). Imagine, for example, that you are working at an office and can either walk to the copy room through a

hallway that also leads to the breakroom or through a hallway that only leads to the copy room. If you see your boss, and want her to know that you are heading to the copy room, then you should choose to use the hallway that only leads there, rather than using the other hallway which might lead your boss to infer you are headed to the breakroom. Even though you are not overtly indicating that your intention is to head to the copy room rather than the breakroom, your boss would be able to infer that you intend to do so simply by observing your actions. This suggests that we may be able to convey aspects of our mental states without resorting to incorporating inefficiencies into our actions.

Here, we propose that the structure of signaling reflects a motivation to make our intentions understood, but that inefficiency will emerge only when efficient action is insufficient to reveal our mental states. Therefore, people should opt for legible efficient action when this is possible and observers should still recognize the intention to signal even when viewing efficient actions.

Experiment 1

In this experiment, we test whether people will choose to signal through efficient action and whether those signals are shaped to guide observer inferences. To examine this question, we presented participants with a grid world in which they had to move to one of several possible goals while also making their target goal obvious to an observer. Critically, participants had the chance to convey their goal while simultaneously using an efficient trajectory.

This set-up is a conceptual replication of prior work (Ho et al., 2016), in which people used efficient trajectories to show observers where in a grid world a reward was located (Experiment 1). Here, we specifically examine whether any inefficient signals emerge. Additionally, we include a more diverse range of grid environments to provide further evidence that people flexibly use efficient action to signal their instrumental goal.

Methods

Participants 80 participants from the US (as indicated by their IP addresses) were recruited through the Prolific research platform. All participants passed a three-question quiz about the directions before completing the task.

Stimuli. The stimuli consisted of simple 7x15 grid worlds, containing fruits and obstacles (Figure 1). Each grid contained three fruit icons, one of which was designated as the target fruit for the round. In addition to the fruit, the grid world included obstacles (represented as rocks).

Two types of grids were created for the experiment: *easy grids* and *difficult grids*. For the easy grids, the participant could disambiguate which of the three fruits was their target fruit in a single move. For example, the target fruit could be 5 squares right and 3 squares down from the agent's starting position, but another fruit could also be four squares down from the agent's starting position. In this case, choosing to move down first would be consistent with efficiently moving

towards two possible fruits, while choosing to move right is only consistent with efficiently moving towards the actual target fruit. Therefore, an individual considering their observer's likely inferences while completing the task should move right first. For the difficult grids, all possible combinations of two efficient moves away from the starting position were confounded with two possible goals. Therefore, participants could only disambiguate what their target fruit was across multiple movements. We also created mirrored versions of each grid in which the same set-up was reflected across the x- and y-axis (see <https://osf.io/kqrwdw> for full stimuli set).

Procedure The study was implemented using the p5.js library (<https://p5js.org>) for JavaScript. Each round, participants had to move an agent around the grid world, using arrows that they clicked on with their mouse; agents could only move in the four cardinal directions. The agent's starting point for each round was pre-determined such that the target fruit could never be reached by moving in a single direction.

Participants were told that they would be playing a game in which they had to navigate between obstacles to collect pieces of fruit. Each round, they were assigned a target fruit (which was the same across participants for a given round) and told that the round would end once they reached their target fruit. The other two fruits were irrelevant to the explicit rules of the game and nothing happened if participants moved the agent over them.

Participants in the *signaling condition* were then told that their trajectories would later be shown to other participants. Those other participants would have to guess what the target fruit was as quickly as they could. Participants in the signaling condition then saw an additional screen clarifying that these observers would be able to see the full grid, but would not know what the target fruit was and that the observers would see their trajectory one move at a time (to clarify that the observers would not have information about how fast the agent moved between grid squares). Participants in the *alone condition* were told that their goal was simply to get to the target fruit in all rounds and there was no mention of observers.

After completing eight rounds of the game, participants in the signaling condition were asked whether or not they tried to make it easier for future participants to guess their target fruit (the question indicated that their response would not impact their compensation). Finally, all participants were also asked whether they found any part of the task confusing and to explain their strategy.

Results and Discussion

Did participants signal through efficient action? In this task, participants in the signaling condition had to complete a goal (moving to their target fruit), while also making that goal obvious to future observers. There are many possible ways in which participants could have signaled what their goal was. However, by design, this task enabled participants to choose

to efficiently move towards the goal in a manner that also quickly disambiguated what their goal was.

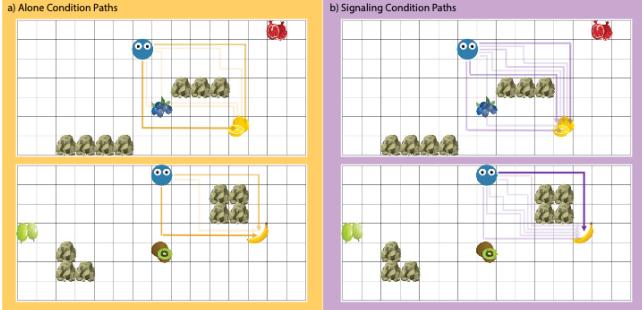


Figure 1: Examples of game grids and participant-generated paths in the (a) alone condition in which participants were told simply to move to the target fruit and the (b) signaling condition in which participants were told to move to the target fruit in a way that would make it easy for an observer to guess their goal. Grids in the top row are difficult grids and grids in the bottom row are easy grids.

If participants in the signaling condition chose to signal through efficient action, then they should have moved to their target fruit in a similar number of moves relative to participants in the alone condition. For each participant-generated trajectory, we subtracted the smallest number of moves required to reach the target goal from the actual number of moves made by the participant to create an inefficiency score. An inefficiency score of 0 indicates that for that round, the participant moved efficiently to the target fruit. We then analyzed whether participants in either condition significantly deviated from efficient action using a linear mixed effects model predicting inefficiency score based on the condition (alone vs. signaling) and grid type (easy vs. difficult) with random intercepts for participants (the maximal random effects structure that allowed the model to converge; Barr et al., 2013). In line with our proposal that people will signal through efficient action, there was no significant difference in the number of moves taken to reach the target goal across conditions ($\beta_{\text{Condition}}=.025, p=.832$), grid types ($\beta_{\text{Grid}}=.063, p=.504$) or their interaction ($\beta_{\text{Condition:Grid}}=-2.776e^{-17}, p=1$), suggesting that participants in both the signaling and alone conditions were moving efficiently towards the target fruit.

While this analysis indicates that participants in the signaling condition were moving efficiently, they may have disregarded the directions and not signaled due to inattention, confusion, or indifference. However, in our post-test survey, all participants in the signaling condition indicated that they were in fact choosing their movements to clarify their goal for the observer, providing initial self-report evidence that people were indeed attempting to signal through efficient action.

Most participants in the signaling condition chose to convey their goal while also navigating along the shortest route possible, and only a small minority deviated from efficient action (5.62% of paths). Interestingly, even among

the minority inefficient signalers, the inefficiencies appear to originate from mistakes such as accidentally moving too far in one direction, and having to reverse direction, rather than disambiguating inefficiencies similar to those observed in previous studies (Ho et al., 2016; McEllin et al., 2018; Royka et al., 2021; Sacheli et al., 2013; Vesper & Richardson, 2014). These initial findings suggest that people will readily opt to signal through efficient action even when they are not prompted to prioritize efficiency.

Nonetheless, these initial results do not reveal whether people actually behaved differently when trying to signal their goal to an observer compared to when they were just navigating to the goal. It is possible that participants favored the same trajectories across the two conditions since eventually all paths move close enough to the target and far enough from distractors that the goal becomes evident. However, if people are trying to reveal their goal to an observer, they should use paths that quickly disambiguate their goal for observers.

To examine the strategies that people employed while signaling through efficiency, we took a computational approach. By modeling how observer inferences change after each successive move in a trajectory, we can evaluate whether signalers' trajectories more quickly lead observers to infer the correct goal relative to trajectories produced by non-signalers in the alone condition.

Computational Framework

To quantify how participants' actions revealed their goals, we implemented a computational model of how observers make goal inferences based on actions in space. To do this, we used a computational framework known as *inverse planning*, which has been shown to capture human goal inferences with quantitative accuracy in simple grid worlds like the ones we used (Baker et al., 2009, 2017). At a high level, inverse planning performs goal inference by building a hypothesis space of possible goals (in this case, each fruit), and assumes that agents move rationally and efficiently towards their goals (an expectation that structures action understanding from early in infancy; Gergely & Csibra, 2003; Liu et al., 2017; Scott & Baillargeon, 2013).

Modeling efficient behavior in space. Formally, we model agents acting in our environment as a probabilistic Markov Decision Process (MDP), a common framework for capturing how agents move in space (see Jara-Ettinger et al., 2020 for details on probabilistic MDPs). The agent could sequentially take one of four actions $A = \{\text{move up, move down, move left, move right}\}$. We assumed that taking actions always incurs a negative cost of 1, and that one of the fruits has a numerical reward of 10 (set to be high enough to outweigh the cost of navigation for obtaining it). This representation enabled us to create a probability distribution over possible action plans (known as *policies*), which would lead the agent to maximize the reward function, making it a useful representation of how we expect agents to act.

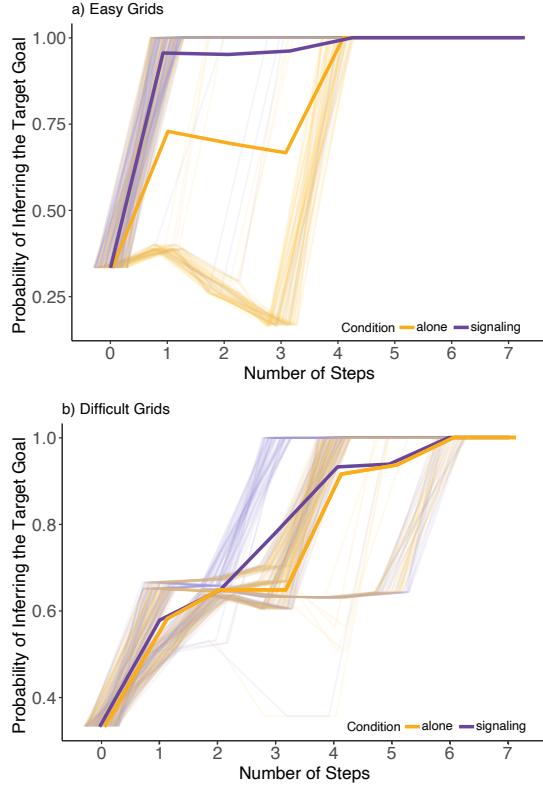


Figure 2: Lines indicate how the probability that an observer would infer the target goal (y-axis) changed with each additional step participants made (x-axis). Lines that are faded represent individual trajectories made by participants; darker lines represent the averages across participants in the alone condition (shown in gold) and signaling condition (shown in purple). For clarity, we present the results from the (a) easy grids and (b) difficult grids separately.

Goal inference. The MDP described above enables us to calculate the probability of an agent’s action, given a reward function. To model how different actions reveal the agent’s goal, we modeled an observer that can see the agent’s actions, but does not know the reward function that the agent is acting under. This observer assumes that the agent is acting under one of three possible reward functions, where each reward function has a constant negative cost of 1 for moving in space, and a numerical reward of 10 for one of the fruits. For simplicity, we refer to each of the reward functions as a goal, as it specifies an agent moving towards each of the three fruits, and each of the three goals has the same prior probability of being the target. With this model, we can measure how well a path reveals the target goal, by calculating the observer’s posterior belief that the agent is pursuing that goal, given by

$$p(G_T|a) \propto p(a|G_T)p(G_T)$$

where G_T is the target goal, and a is the observed sequence of actions (normalized by considering the likelihood of the same actions under the pursuit of the other two goals).

Results and Discussion

Using this model, we calculated the observer’s posterior belief that each participant trajectory was pursuing the target fruit (hereafter referred to as the target goal) for each individual step in the trajectory. Because we are interested in whether the efficient paths produced in the signaling condition were indeed revealing the underlying goal better than those in the alone condition, here we focused on analyzing only the efficient paths (94.38% of paths; $n=604$ of 640 total paths). If participants in the signaling condition are shaping their efficient actions to guide observers’ inferences, then our observer model should generate higher posterior beliefs that the agent is pursuing the target goal after fewer steps, relative to trajectories from the alone condition. Figure 1 depicts participant-generated trajectories across conditions and Figure 2 shows the model’s probability of attributing the target goal as a function of the observed action. As Figure 2 shows, participant actions revealed the target goal faster in the signal condition than in the alone condition, but this effect was more pronounced in the easy grids.

To analyze whether there was a significant difference in how effectively participants revealed their goal, we used a linear mixed effects model predicting the posterior belief that the agent is pursuing the target goal at each successive step made by the participant based on the condition (alone vs. signaling), the number of steps taken so far, and their interaction with random intercepts for individual trajectory (the maximal random effects structure that allowed the model to converge; Barr et al., 2013). Consistent with the pattern shown in Figure 2, the probability of inferring the target goal significantly increased with additional steps ($\beta_{Step}=.061$, $p < .001$), as participants are getting closer to their goal over time and further from the distractor fruits. Critically, the probability of inferring the correct goal was significantly higher for paths in the signaling condition ($\beta_{Condition}=.160$, $p < .001$). Additionally, the number of steps caused a greater increase in the probability that an observer would infer the correct goal in the alone condition relative to the signaling condition ($\beta_{Step:Condition} = -.022$, $p < .001$), which indicates that the difference in the inferability of goals between the two conditions decreases with additional steps. This provides evidence that participants in the signaling condition revealed their target more quickly than participants in the alone condition, even though both were moving efficiently.

Taken together, these results show that participants in the signaling condition were creating different trajectories relative to participants in the alone condition, conceptually replicating prior work (Ho et al., 2016). Importantly, the signaling trajectories more quickly revealed the target goal, suggesting that signalers accounted for how their actions would affect observers’ goal inferences.

Experiment 2

The signaling paths generated in Experiment 1 allowed our observer model to quickly infer their instrumental goal. However, these trajectories were not just shaped by an instrumental goal, they were also shaped by the intention to

signal. Do people have intuitions about when others are signaling through efficient action?

Intuitively, it seems as though we can sometimes tell when someone is trying to make their behavior understood even when they are moving efficiently towards a goal. Imagine that you are the boss of a company and you just admonished an employee for slacking off. If you then see the employee choose to walk to the copy room along a hallway that exclusively leads to the copy room rather than using a hallway that also leads to the breakroom, you might infer that the employee is trying to make it clear that they are going to the copy room.

However, because this kind of signal would be perfectly confounded with the actions that someone might take to simply accomplish a goal without intending to signal, inferring signaling intent for efficient action is not trivial. As such, people may be reluctant to attribute signaling intentions to efficient action, and instead may expect signals to have some inefficiency (Royka et al., 2022; Trujillo et al., 2018).

Here, we examine whether people are willing to attribute signaling intent to efficient action and whether their intuitions track how often those actions were actually used to signal.

Methods

Participants 41 participants from the US (as indicated by their IP addresses) were recruited through the Prolific research platform. All participants passed a three-question quiz about the directions before completing the task.

Stimuli. Participants saw videos of 24 paths generated by participants from Experiment 1. We chose three paths per grid that varied based on how frequently they were produced in the alone and signaling conditions (see <https://osf.io/kqrwd> for videos and detailed path choosing procedure).

Procedure Participants were told that they would watch videos of different people playing a game. In the game, players had to collect a certain fruit on the grid. Participants learned that there were two versions of the game: an *alone* version in which the players were just moving to collect their target fruit and a *signaling* version in which the player was trying to make it as obvious as possible to an observer which fruit they were going to collect.

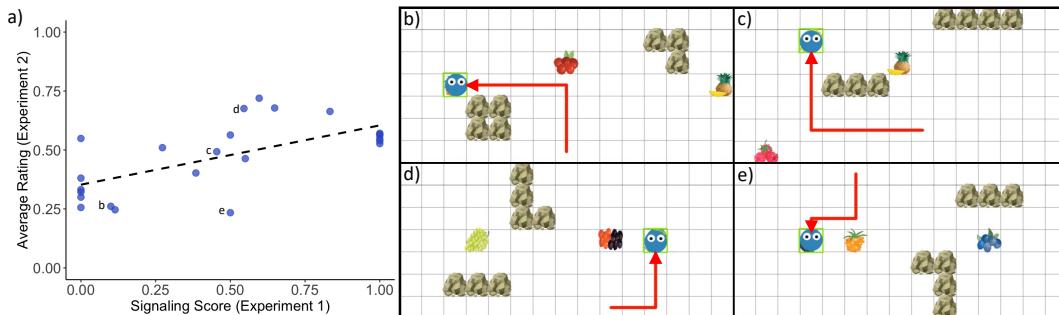


Figure 3: (a) Average judgements from Experiment 2 (which version of the game they thought that player was playing on a scale of 0, *definitely the alone version*, to 1, *definitely the signaling version*) as a function of the path's signaling score.

(b-e) Examples of paths shown to participants with corresponding datapoints marked on (a).

Each round, participants watched a video showing how one player moved to their target fruit. Then, participants had to indicate which version of the game they thought that player was playing on a scale of 0 (*definitely the alone version*) to 1 (*definitely the signaling version*).

After rating all 24 paths, participants were also asked whether they found any part of the task confusing and to explain their strategy.

Results and Discussion

First, we calculated a *signaling score* for each path by dividing the number of times that the path was produced in the signaling condition of Experiment 1 by the total number of times the path was produced. Therefore, a signaling score of 1 meant that a path was exclusively produced to signal, while a signaling score of 0 meant that the path was never produced in the signaling condition. We then examined the correlation between paths' signaling scores and the average participant ratings in Experiment 2. We found a significant positive correlation ($r=.65, p<.001$; Figure 3), suggesting that people were more likely to think that a path was produced to signal when that path was actually produced more often to signal in Experiment 1.

Importantly, this initial result shows that people do not rely on inefficiency to infer signaling intent. Instead, people have intuitions about when an efficient action may be driven by the intention to signal and those intuitions align with the actual signaling actions of others.

This analysis, however, is only a rough measure of how people's intuitions track with real signaling behavior. First, our analysis does not account for how frequently people produced the paths in Experiment 1; a path may have only been produced in the signaling condition and thus have a high signaling score, but it may have only been produced once. Additionally, we did not directly examine how these judgements correlate with how well each path revealed its instrumental goal. Although the paths in the signaling condition of Experiment 1 generally revealed their goal more quickly than the paths in the alone condition, future analysis could use the inferences of our observer model to derive a more continuous score for how well each path reveals its goal. It is possible that how quickly a path reveals its goal

may be even more strongly predictive of observers' inferences about signaling intent.

General Discussion

Human social behavior relies not only on our ability to understand other people's actions, but also on our ability to ensure that our own actions are easily understood. Consistent with this, past work has found that people routinely add inefficiencies into their behavior with the goal of making the movement easier to interpret (McEllin et al., 2018; Pezzulo & Dindo, 2011; Royka et al., 2021; Sacheli et al., 2013; Trujillo et al., 2018; Vesper & Richardson, 2014) and these inefficiencies help observers to recognize when someone is moving to signal (Royka et al., 2022; Trujillo et al., 2018). As such, inefficient action is central to signaling behavior (Dockendorff et al., 2019; Pezzulo et al., 2013, 2019; Royka et al., 2022). However, here we found evidence that inefficiency is not intrinsic to signaling behavior. Using a paradigm where goals can be revealed without the need to be inefficient, people overwhelmingly preferred efficient signaling. Moreover, when observers saw these efficient signals in a subsequent task, they inferred signaling intent for some of the actions even though they all moved from their start to their goal in as few moves as possible.

While our findings show that inefficiency is not intrinsic to creating or recognizing signaling behavior, our findings do not diminish the importance of signaling through inefficiency. Instead, our work suggests that inefficiency is a critical tool for the cases where it is impossible to quickly reveal one's goal through efficient action (McEllin et al., 2018; Royka et al., 2021; Sacheli et al., 2013; Vesper & Richardson, 2014). Our work therefore suggests that people's inclusion of inefficiencies in their signaling movement might be an intentional decision that reflects some (potentially implicit) awareness that efficient behavior will not suffice.

One factor that may motivate people to engage in inefficient signaling could be a need to make your observer recognize that you are shaping your actions to convey a message to them. In our first task, participants had to reveal their goal to an observer, but they did not need to make the observer aware that the movement was intentionally signaling the goal. This distinction often separates signaling (where some intentional behavior helps an agent disclose information to an observer) from communication (where the observer must be aware that the movement was generated with the explicit purpose of sharing information). It is therefore possible that inefficiency might become more important in communicative cases, where observers might find it difficult to identify efficient behavior that is intentionally communicative (although people may also be able to reveal their communicative intent via ostensive cues such as eye contact; Behne, Carpenter & Tomasello, 2005; Csibra & Gergely, 2009; Senju & Csibra, 2008).

There are, however, situations in which we want to convey information to someone, but we do not need or even want people to know whether the signaling was intentional or not. Imagine, for instance, walking late at night at a faster pace

than someone who's ahead of you. If you need to cross the street at some point, you might consider crossing the street sooner than you otherwise would, simply so that the other person does not have to wonder if they are being followed. In this case, it might be less critical that the observer infers that you wanted to signal you were going a different way, and the person might not even become aware that you were walking behind them, but the signal is low-cost enough that it is worth generating. The use of efficient signals may be shaped by these and related pragmatic concerns not typically explored in signal production contexts.

In our observer inference task, however, we explicitly told participants that the agents would sometimes be moving to reveal their instrumental goal. As such, participants did not need to spontaneously infer that signaling was occurring. It remains an open question whether people can spontaneously infer signaling intent for efficient action without any additional context or cues, and what kind of behavioral cues (e.g., action repetitions) may trigger such spontaneous inferences about an intention to signal. On the other hand, consider a situation where you repeatedly see someone act in a manner that clearly reveals their goals across a wide range of contexts. Would observing multiple actions that each disambiguate their own goal trigger spontaneous inferences about an intention to signal? Future work should examine what observations lead people to spontaneously infer that someone else is signaling through efficient action.

At a broader level, our work advances the idea that humans are able to flexibly convey their mental states to observers through action. While the capacity to use both efficient and inefficient strategies to reveal goals and intentions is a powerful tool that helps humans to solve social coordination problems, there are likely many other mechanisms through which people can quickly reveal a wide array of mental states in order to ease observers' inferences. For example, some interjections convey information about knowledge (saying "Oh!" indicates that you did not previously know about something) and intention ("Oops!" indicates that you did something unintentionally). Similarly, people will spontaneously explicitly inform others of their intentions in some situations, but not others (e.g., when getting up from a table at a restaurant, people will usually tell the rest of their party what they are getting up to do). It's likely that these spontaneous expressions of intention are somewhat dictated by the inferability of the person's goal (is the person getting up to use the restroom? Going to the bar? Making a hasty exit?), suggesting that even seemingly straight-forward uses of mental state language may be driven by what an observer can infer.

The signaling behaviors explored here may constitute just one instance of what is actually a suite of cognitive tools that help us to help others make sense of our behavior. In this way, Theory of Mind can be understood as an interactive process in which mental state inference does not solely rely on behaviors performed by a disinterested actor. Instead, people navigate their social world in a way that makes their own minds more legible to others.

Acknowledgments

This work was supported by NSF award BCS-2045778.

References

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*.

Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*.

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of memory and language*.

Behne, T., Carpenter, M., & Tomasello, M. (2005). One-year-olds comprehend the communicative intentions behind gestures in a hiding game. *Developmental Science*.

Brand, R. J., Baldwin, D. A., & Ashburn, L. A. (2002). Evidence for 'motionese': modifications in mothers' infant-directed action. *Developmental science*.

Csibra, G., & Gergely, G. (2009). Natural pedagogy. *TiCS*.

De Ruiter, J. P., Noordzij, M. L., Newman-Norlund, S., Newman-Norlund, R., Hagoort, P., Levinson, S. C., & Toni, I. (2010). Exploring the cognitive infrastructure of communication. *Interaction Studies*.

Dockendorff, M., Sebanz, N., & Knoblich, G. (2019). Deviations From Optimality Should Be an Integral Part of a Working Definition of SMC. *Physics of life reviews*.

Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naive theory of rational action. *TiCS*.

Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *TiCS*.

Gopnik, A., & Meltzoff, A. N. (1994). *Minds, bodies, and persons: Young children's understanding of the self and others as reflected in imitation and theory of mind research*.

Gweon, H. (2021). Inferential social learning: Cognitive foundations of human social learning and teaching. *TiCS*.

Ho, M. K., Littman, M., MacGlashan, J., Cushman, F., & Austerweil, J. L. (2016). Showing versus doing: Teaching by demonstration. *NeurIPS*.

Jara-Ettinger, J., & Rubio-Fernandez, P. (2021). Quantitative mental state attributions in language understanding. *Science advances*.

Jara-Ettinger, J., Schulz, L. E., & Tenenbaum, J. B. (2020). The naive utility calculus as a unified, quantitative framework for action understanding. *Cognitive Psychology*.

Liu, S., Ullman, T. D., Tenenbaum, J. B., & Spelke, E. S. (2017). Ten-month-old infants infer the value of goals from the costs of actions. *Science*.

McEllin, L., Knoblich, G., & Sebanz, N. (2018). Distinct kinematic markers of demonstration and joint action coordination? Evidence from virtual xylophone playing. *JEP: Human Perception and Performance*.

Pezzulo, G., & Dindo, H. (2011). What should I do next? Using shared representations to solve interaction problems. *Experimental Brain Research*.

Pezzulo, G., Donnarumma, F., & Dindo, H. (2013). Human sensorimotor communication: A theory of signaling in online social interactions. *PloS one*.

Pezzulo, G., Donnarumma, F., Dindo, H., D'Ausilio, A., Konvalinka, I., & Castelfranchi, C. (2019). The body talks: Sensorimotor communication and its brain and kinematic signatures. *Physics of life reviews*.

Royka, A., Chen, A., Aboody, R., Huanca, T., & Jara-Ettinger, J. (2022). People infer communicative action through an expectation for efficient communication. *Nature Communications*.

Royka, A., Schouwstra, M., Kirby, S., & Jara-Ettinger, J. (2021). I Know You Know I'm Signaling: Novel gestures are designed to guide observers' inferences about communicative goals. *In CogSci Proceedings*.

Sacheli, L. M., Tidoni, E., Pavone, E. F., Aglioti, S. M., & Candidi, M. (2013). Kinematics fingerprints of leader and follower role-taking during cooperative joint actions. *Experimental brain research*.

Schachner, A., & Carey, S. (2013). Reasoning about 'irrational' actions: When intentional movements cannot be explained, the movements themselves are seen as the goal. *Cognition*.

Scott, R. M., & Baillargeon, R. (2013). Do infants really expect agents to act efficiently? A critical test of the rationality principle. *Psychological science*.

Scott-Phillips, T. C., Kirby, S., & Ritchie, G. R. (2009). Signalling signalhood and the emergence of communication. *Cognition*.

Senju, A., & Csibra, G. (2008). Gaze following in human infants depends on communicative signals. *Current Biology*.

Trujillo, J. P., Simanova, I., Bekkering, H., & Özyürek, A. (2018). Communicative intent modulates production and comprehension of actions and gestures: A Kinect study. *Cognition*.

Vesper, C., & Richardson, M. J. (2014). Strategic communication and behavioral coupling in asymmetric joint action. *Experimental brain research*.

Wilson, D., & Sperber, D. (2004) Relevance Theory. In G. Ward, L. Horn. *Handbook of Pragmatics*, Oxford: Blackwell.

Young, L., Cushman, F., Hauser, M., & Saxe, R. (2007). The neural basis of the interaction between theory of mind and moral judgment. *PNAS*.