

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Individual vs. Joint Perception: a Pragmatic Model of Pointing as Smithian Helping

Permalink

<https://escholarship.org/uc/item/3fp0v20n>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

ISSN

1069-7977

Authors

Jiang, Kaiwen
Stacy, Stephanie
Chan, Adelpha
et al.

Publication Date

2021

Peer reviewed

Individual vs. Joint Perception: a Pragmatic Model of Pointing as Communicative Smithian Helping

Kaiwen Jiang¹
kaiwenj@g.ucla.edu

Stephanie Stacy¹
stephaniestacy@g.ucla.edu

Chuyu Wei³
chuyuwei@g.ucla.edu

Adelpha Chan⁴
adelchan07@g.ucla.edu

Federico Rossano⁵
frossano@ucsd.edu

Yixin Zhu¹
yixin.zhu@ucla.edu

Tao Gao^{1,2}
tao.gao@stat.ucla.edu

¹ Department of Statistics, UCLA ² Department of Communication, UCLA ³ Department of Psychology, UCLA
⁴ Department of Linguistics, UCLA ⁵ Department of Cognitive Science, UCSD

Abstract

The simple gesture of pointing can greatly augment one’s ability to comprehend states of the world based on observations. It triggers additional inferences relevant to one’s task at hand. We model an agent’s update to its belief of the world based on individual observations using a partially observable Markov decision process (POMDP), a mainstream artificial intelligence (AI) model of how to act rationally according to beliefs formed through observation. On top of that, we model pointing as a communicative act between agents who have a mutual understanding that the pointed observation must be relevant and interpretable. Our model measures “relevance” by defining a Smithian Value of Information (SVI) as the utility improvement of the POMDP agent before and after receiving the pointing. We model that agents calculate SVI by using the cognitive theory of Smithian helping as a principle of coordinating separate beliefs for action prediction and action evaluation. We then import SVI into rational speech act (RSA) as the utility function of an utterance. These lead us to a pragmatic model of pointing allowing for contextually flexible interpretations. We demonstrate the power of our Smithian pointing model by extending the Wumpus world, a classic AI task where a hunter hunts a monster with only partial observability of the world. We add another agent as a guide who can only help by marking an observation already perceived by the hunter with a pointing or not, without providing new observations or offering any instrumental help. Our results show that this severely limited and overloaded communication nevertheless significantly improves the hunters’ performance. The advantage of pointing is indeed due to a computation of relevance based on Smithian helping, as it disappears completely when the task is too difficult or too easy for the guide to help.

Keywords: pointing; pragmatics; joint attention; Smithian helping; cooperation; rational speech act

Introduction

Like all animals, we understand the world by collecting observations through our individual perception, which we call “individually perceived observations.” Being social creatures, however, we also get observations that are pointed out to us by others. When someone points, the addressee knows that they and the person who initiated the pointing gesture both get an observation, propagating a belief update based on the observation. We call this a “jointly perceived” observation. In this modeling paper, we use the term “observation” to refer to the raw sensory input as in the field of artificial intelligence (AI) (Kaelbling, Littman, & Cassandra, 1998). We use the term “perception” to refer to the inference of the most likely world state that generates the observation following the Bayesian perspective of perception (Knill & Richards, 1996).

This paper aims to demonstrate from a computational perspective that the joint perception enabled by pointing is more potent than the individual perception of the same observation. We aim to prove that insights of paternalistic helping from developmental psychology can inspire the development of socially capable AI systems. Also, our mathematical modeling grounded in AI algorithms can scaffold future cognitive science researches on Smithian (paternalistic) helping.

Pointing Gesture in Human Communication

“Point to a piece of paper. And now point to its shape—now to its color—now to its number. . . . How did you do it?”
—Wittgenstein and Anscombe (1953/2001)

Until Wittgenstein called attention to its underlying complexity, pointing had generally been perceived as an intuitive, unremarkable communicative gesture. It is among the most conspicuous and common forms of human communication. Children point to help adults retrieve objects they were looking for, foraging partners point out the potential locations of food to help each other, and customers point to their empty glasses to request assistance from the server. Pointing is among the first communicative gestures human infants learn to use (Butterworth, Simion, et al., 2013): at as young as the age of one, infants use pointing to communicate information (Liszkowski, Carpenter, Striano, & Tomasello, 2006). Although pointing is pervasive in everyday human life, it is rarely observed in wild animals. Human-raised great apes can produce pointing-like gestures to invite humans to cooperate with them in obtaining food (Leavens & Hopkins, 1998), yet pointing in great apes lacks the cooperative properties of human pointing: apes are not bothered when the partner is distracted or non-responding (Van der Goot, Tomasello, & Liszkowski, 2014). Chimpanzees’ failure to achieve a deep understanding of pointing suggests that human pointing may reveal surprising intricacy of human communication.

The same pointing act can be interpreted differently in various contexts, which reveals two properties of pointing. First, pointing is **overloaded**. As Wittgenstein and Anscombe (1953/2001) pointed out, the same pointing gesture has many interpretations, making the referent of pointing ambiguous when considered in isolation. Second, pointing is **indirect**; a big gap can exist between the referent and the meaning of the pointing. The receiver must infer what to do with the referent beyond looking at it. It has been shown that by following

an adult's pointing to a toy, young children could adaptively decide what to do with the toy, put it away or examine it, based on the context of the pointing (Liebal, Carpenter, & Tomasello, 2011).

The overloadedness and indirectness of pointing enable it to express manifold meanings with the same observation, no new observations provided. The meaning of the pointing can be interpreted depending on the context. This makes pointing a powerful communicative act that can significantly facilitate human cooperation. For example, when two hunters walk in a forest together, the young hunter perceives a broken stick on the ground but does not think it is relevant to the hunt. Just the moment he is about to move on, the experienced hunter grabs his attention and points to the same broken stick he has already perceived. The young hunter immediately realizes that the broken stick is a trace of their prey. This example highlights that joint perception enabled by pointing can evoke richer inferences than the individual perception of the same observation (Sperber & Wilson, 1986).

As in the above example, a crucial function of pointing is helping. Crucially, it is a particular type of helping with two unique characteristics. First, the helper is in a position to help because her belief is closer to reality, not because she has any physical advantage. Therefore, pointing must involve diverging beliefs in which the helper knows better how to improve the helpee's well-being. This type of helping is called paternalistic helping or Smithian helping in developmental psychology (Martin, Lin, & Olson, 2016), which we will introduce later. Second, unlike instrumental actions, pointing does not change the physical states at all. Its only function is to change the helpee's mind. Therefore, models of instrumental helping would fail to apply (Ullman et al., 2009). Instead, we argue that it should be understood as an "utterance" in rational speech act (RSA), a pragmatic model of language which also views language as cooperative (Frank & Goodman, 2012; Goodman & Frank, 2016). RSA treats an utterance as an action with a utility function. The generation and interpretation of an utterance can be modeled with the principle of maximizing expected utility from decision theory.

Due to the above two unique characteristics, we propose to model pointing as an utterance with a utility derived from Smithian helping for coordinating diverging beliefs.

Smithian Empathy and Helping

The concept of Smithian helping is based on Adam Smith's discussion of empathy. During his discussion, he compared two types of empathy. First, he addressed Hume's definition, which proposes that empathy is a simple resonance of other's feelings. In contrast to this conventional definition, Smith proposed that true empathy involves the coordination of mindsets between the empathizer and the agent being empathized with. To better understand the two competing views addressed in Smith's argument, one can imagine a theoretical example involving you and your friend. Both of you are backstage preparing for your friend's performance in the school talent show. While your friend is excited to perform, you

dread the performance because you know that he is objectively bad at singing.

Hume's conventional definition of empathy proposes that empathy is the **direct mirroring** of another's mindset (Hume, 1751/2018). In other words, the agent and subject involved in a particular act of empathy should share the same mindset. In the talent show example, you and your friend each have a different perspective in regards to the same action of your friend performing. However, according to Hume's definition, to successfully empathize with them, you must abandon your personal opinion and take on your friend's perspective.

On the other hand, Smith's own proposed definition of empathy is the act of investigating how an agent would feel if they, in their current state of consciousness, were placed into the target individual's situation (Smith, 1759/2010). In the talent show example, you and your friend each hold a distinct perspective regarding the action of your friend performing in the talent show. When executing the Smithian empathy, you maintain your own mindset when evaluating the action of interest. As a result, empathizing involves applying your own mindset to evaluate the situation of your friend performing. Because you know your friend is bad at singing, your evaluation of how you would feel in their situation leads you to be worried that your friend may embarrass himself on stage. As illustrated in the example, Smithian empathy involves **coordinating diverging mindsets** of two agents. Your act of worrying is an attempt to coordinate your friend's belief of excitement about the upcoming performance with your opposing belief that their performance will have a bad outcome.

Smithian empathy has been explicitly extended to the well-studied phenomenon of human behavior known as paternalistic helping. Paternalistic helping involves a helper doing what she thinks to be good to the helpee, even if that is not what the helpee wants (Martin et al., 2016). This act involves the coordination of two mindsets as the helper must balance the helpee's desire with what she personally judges to be the best for the helpee. The helper acts to optimize the well-being of the helpee. In the talent show example, you believe that your friend will utterly embarrass himself if he performs. With this belief in mind, you become anxious and feel a strong urge to help your friend by convincing him not to perform. This action of stopping your friend from performing is an example of paternalistic helping as it opposes their desire to perform. You stopping your friend is driven by a sense of worry that stems from an opinion formed by applying your unique perspective to your friend's situation. On the surface, paternalistic seems deeply related to the ability to reason about other's beliefs, but it is arguably more complicated. In the famous false belief task (Wimmer & Perner, 1983), the child only needs to select one belief to predict other's actions. In paternalistic helping, one not only needs to predict action with others' beliefs but also evaluate the actions with their own belief.

Paternalistic helping is a behavior prevalent in children that has been well studied. It has been shown that children will override a request if they recognize that following the request




Figure 1: **Modeling pointing using Smithian coordination of beliefs:** action prediction using receiver’s belief, action evaluation using signaler’s belief. Numbered boxes represent key components of the model.

might harm the requester. When children interact with another child who expressed a preference for chocolate over fruit snacks but would be sick after eating chocolate, most override the request for chocolate and offer the fruit snacks (Martin et al., 2016). In addition, when the wrong tool is requested, children offer the tool they believe the requester needs, not the one they are asked for (Hepach, Benziad, & Tomasello, 2020). These studies provide a solid theoretical foundation on how to coordinate two beliefs in helping. We build upon these insights and focus on cases in which helping behavior is executed to provide information through communication, as in pointing.

Modeling Individual vs. Joint Perception

Our modeling study is directly inspired by the perception of the broken stick in the hunting example. The ultimate goal here is to demonstrate that an intelligent agent with joint perception enabled by pointing can outperform an agent only with individual perception in a hunting task.

We start by outlining the model of an intelligent agent that acts based on individual perception. On top of that, we formulate a model of pointing for joint perception. The components of the models are shown in Fig. 1.

Modeling Agent with Individual Perception Using POMDP

For modeling an intelligent agent with individual perception, we use the partially observable Markov decision process (POMDP) (Kaelbling et al., 1998). POMDP provides a generic formulation of how an agent takes rational actions in an uncertain environment with only limited observations. It models two processes in the agent-environment interaction.

The first process is how an agent updates its belief of the world state s with observations from the environment using Bayesian inference (resulting in Box 1 in Fig. 1). A belief is defined as a probabilistic distribution over the set of possible states:

$$b(s) = P(s|b). \quad (1)$$

This definition of belief is used in both AI (Kaelbling et al., 1998) and cognitive modeling (Baker, Saxe, & Tenenbaum,

2011). The agent updates its belief when it gets an observation o after taking action a . Let b' be the updated belief and s' the next state after taking a , we have

$$b'(s') = P(s'|b') = P(s'|b, o) \propto P(o|s', a) \sum_{s \in \mathcal{S}} P(s'|a, s) P(s|b). \quad (2)$$

This belief update only involves individual perception because the agent treats the observation only as being generated by its own interaction with the environment. It does not view it as a referent of any communicative intention as in Grice (1975), even with the presence of a second agent.

The second process is the decision model of how an agent takes rational actions based on its belief. Planning over beliefs involves the calculation of expected utility of each action. The conventional expected utility of an action a can be defined as the expectation of utility of the outcomes s' of the action (Russell, Norvig, & Davis, 2010):

$$\mathbb{E}U(a|b) = \int_{s' \in \mathcal{S}} U(s')P(s'|a, b), \quad (3)$$

where the probability of the outcomes can be calculated with a transition model:

$$P(s'|a, b) = \int_{s \in \mathcal{S}} P(s'|s, a) P(s|b). \quad (4)$$

The agent then selects an action expected to maximize its utility:

$$a^* = \operatorname{argmax}_a \mathbb{E}U(a). \quad (5)$$

In a POMDP, the actions of a rational agent is determined by its belief. With knowledge of an agent’s belief, the agent’s actions (Box 3 in Fig. 1) can be predicted. This action prediction process is represented in Fig. 1 as Box 2.

Here we only outline the general principle of utility calculation in belief space. Planning rational actions in belief space to optimize long-term accumulated rewards is a challenging problem. In our study, we use the point-based value iteration (PBVI) algorithm (Pineau, Gordon, & Thrun, 2003) as the solver for POMDP.

Modeling Agent with Joint Perception: Smithian Pointing

For an individual agent, the expected utility of action, belief, and the value of information provided by an observation can be calculated based on the agent’s individual belief. We refer to them as the “conventional” utilities, which can be easily derived based on the classic theories (Russell et al., 2010). Our Smithian model of pointing augments these conventional utilities to reflect the coordination of two beliefs in Smithian helping: The signaler takes on the role of helper, while the receiver is the helpee. Specifically, the signaler uses the receiver’s belief to predict the receiver’s actions, and then uses her own belief to evaluate the receiver’s actions.

Smithian Utility of Action We define Smithian utility of actions by augmenting the conventional utility of actions defined in Eq. (3). There are no subscripts in Eq. (3), which implies that the action and belief are from the same agent. Eq. (3) can apply when one agent evaluates its own action and belief, or when one agent (A) uses theory of mind to take the perspective of another agent (B) to evaluate the agent B’s action based on agent B’s belief.

With Smithian empathy, the signaler should use her own utility function to evaluate the outcomes (Box 4 in Fig. 1). In many scenarios, the utility functions of the signaler and the receiver are consistent. In some cases, the utility functions may be different; but constrained by Smithian helping, the signaler’s utility should always be aligned with the receiver’s physical well-being. Here we are using subscripts to represent the source of belief or action, with *Sig* for signaler and *Rec* for receiver. We write down the formulation of Smithian utility of action by changing the belief used for action evaluation in Eq. (3) to the belief of the signaler:

$$\mathbb{E}U_{Smith}(a_{Rec}|b_{Sig}) = \int_{s' \in S} U_{Sig}(s') \int_{s \in S} P(s'|s, a_{Rec}) P(s|b_{Sig}). \quad (6)$$

Smithian Utility of Belief The effect of pointing in our study is to change the receiver’s belief. Therefore, to evaluate the effect of pointing, we need to define the utility of a belief, especially the receiver’s belief without the pointing (Box 1 in Fig. 1) and after the pointing (Box 5 in Fig. 1). It can be derived from the expected utility of actions, as an agent’s distribution (Boxes 3 and 7 in Fig. 1) of its action $P(a|b)$ can be predicted from its belief.

Given $P(a|b)$, we can derive the conventional utility of a belief based on conventional expected utility of action. Let \mathcal{A} be the set of possible actions. The utility of a belief can be defined as:

$$U(b) = \int_{a \in \mathcal{A}} P(a|b) \mathbb{E}U(a|b). \quad (7)$$

Eq. (7) can represent an agent’s evaluation of its own beliefs. It can also represent an agent A’s evaluation of another agent B’s beliefs, when A takes B’s perspective using Hume’s definition of empathy. Using theory of mind, A can predict B’s action based on B’s belief (Wellman, 2014), hence evaluate the belief by integrating out the evaluation of actions.

We can derive Smithian utility of belief by replacing the expected utility of actions in Eq. (7) with the Smithian utility of actions:

$$U_{Smith}(b_{Rec}|b_{Sig}) = \int_{a_{Rec} \in \mathcal{A}_{Rec}} P(a_{Rec}|b_{Rec}) \mathbb{E}U_{Smith}(a_{Rec}|b_{Sig}). \quad (8)$$

Smithian utility of belief represents the signaler’s evaluation of the receiver’s well-being (Boxes 5 and 8 in Fig. 1), which the signaler tries to improve.

Smithian Value of Information (SVI) Pointing carries information. In the field of AI, the value of information is “the difference in expected value between the best actions before and after the information is obtained” (Russell et al., 2010), with which an agent can decide how much the information is worth. Here we leverage the insight that the value of information should be calculated as the change in expected utility before and after the information is obtained. Smithian Value of Information (SVI) adopts this formalization with two significant differences. First, we calculate the value of information from the perspective of the signaler, not the receiver of the information. Second, the utility we use is the Smithian utility of belief, which is the signaler’s estimate of the receiver’s well-being. Therefore, SVI measures the improvement of the signaler’s estimate of the receiver’s well-being before and after pointing, which is the difference between Box 5 and Box 8 in Fig. 1. In other words, **SVI is the utility of pointing**.

To later incorporate our formulation of pointing into RSA as the utterance, here we use notation u for pointing. Let b_{Rec} be the receiver’s belief before receiving the pointing signal, b'_{Rec} be receiver’s belief after receiving the pointing signal:

$$b'_{Rec} = P_{Rec}(s|u) = P(s|b_{Rec}, u). \quad (9)$$

We can write SVI as:

$$SVI(u|b_{Sig}) = U_{Smith}(b'_{Rec}|b_{Sig}) - U_{Smith}(b_{Rec}|b_{Sig}). \quad (10)$$

Pointing as a Rational Speech Act With the utility of pointing clearly defined as SVI, we can treat pointing as a special type of utterance and model its use and interpretation using RSA (Box 9 in Fig. 1). We can write down how a signaler generates pointing u :

$$P_{Sig}(u|b_{Sig}) \propto \exp\{\alpha[SVI(u|b_{Sig}) - c(u)]\}. \quad (11)$$

For simplicity of the model, we can set $c(u) = 0$ as the cost of pointing in real world is small.

We still need to adapt RSA to consider the entire POMDP challenge as the “context” of communication. We first outline the RSA in language games where the environment is fully observable to both the signaler and receiver. A pragmatic receiver updates its belief upon receiving an utterance u with Bayesian inference (Goodman & Frank, 2016)

$$b'_{Rec}(s) = P_{Rec}(s|b_{Rec}, u) \propto P_{Sig}(u|s) P_{Rec}(s|b_{Rec}). \quad (12)$$

Then we adapt RSA to the case in which neither the signaler nor the receiver can fully observe the environment. In

this case, both signaler and receiver maintain a belief based on their observations. Based on existing common ground knowledge, the receiver knows the probability that the signaler has a belief given the physical state $P(b_{Sig}|s)$ (Goodman & Stuhlmüller, 2013). The receiver can infer the state of the world with the utterance from the signaler:

$$P_{Rec}(s|u) \propto \int_{b_{Sig}} P_{Sig}(u|b_{Sig})P(b_{Sig}|s)P(s|b_{Rec}). \quad (13)$$

Eq. (13) can be simplified when the signaler has full knowledge of the world. In this case, the signaler knows a state s for sure without uncertainty. Of course it can still be represented as a belief, a probability distribution with probability 1 on the true state s^* , denoted as b_{Sig}^* . The receiver only needs to consider all possible s and their corresponding belief b_{Sig}^s . Since the signaler knows the true state, $P(b_{Sig}^{s_i}|s_j) = 1$ if $s_i = s_j$, $P(b_{Sig}^{s_i}|s_j) = 0$ if $s_i \neq s_j$. The integration in Eq. (13) has only one nonzero entry for each state. Eq. (13) reduces to:

$$P_{Rec}(s|u) \propto P_{Sig}(u|b_{Sig}^s)P(s|b_{Rec}). \quad (14)$$

Assuming b_{Sig}^s is true is equivalent to assuming the signaler knows the true state is s . Therefore, using s to denote the true state is s , Eq. (14) further reduces to:

$$b'_{Rec}(s) \propto P_{Sig}(u|s)P_{Rec}(s|b_{Rec}). \quad (15)$$

In our experiment, we only consider the case that the signaler is certain about the state, so Eq. (15) is used in the experiment.

Signal generation and interpretation are modeled through recursive reasoning: For each signal, the signaler estimates the change of receiver’s belief after receiving the signal using Eq. (15) and selects the signal that maximizes Smithian value of information in Eq. (11). This model of signal generation can then be passed to the next level of receiver for computing the probability of a signal given a world state. This allows the receiver to update its belief as the posterior of the world state given a signal using Bayes’ rule. Note that this recursive social reasoning is based on RSA. However, in RSA, the recursion begins with a literal receiver who infers the reference given an utterance. In our model, the literal receiver is the POMDP agent who takes rational actions based on observations. In Eq. (8), to compute $U_{Smith}(b_{Rec}|b_{Sig})$, one needs to compute $P(a_{Rec}|b_{Rec})$ and then integrate out a . Here $P(a_{Rec}|b_{Rec})$ is the policy solved by POMDP.

Modeling Experiment

Task: Guided Wumpus Hunting

To highlight the strengths of the Smithian pointing model, we need a single-agent partially observable task as a baseline. We pick a simplified version of the classic AI problem: *the Wumpus world* (Russell et al., 2010). We augment the task for Smithian pointing by adding an additional helper agent. We call this task the *guided Wumpus hunting game*, directly inspired by the hunting example described in the introduction.




Figure 2: **The environment of the guided Wumpus hunting game.** Wumpus can only show up in one of three shaded tiles. Starting from (0, 0), the hunter tries to infer Wumpus’s location from the stench and shoot it.

In the guided Wumpus hunting game, a hunter navigates through a set of tiles to shoot a stationary monster, the Wumpus, without knowing its exact location. The Wumpus emits a “stench” when it is nearby, which the hunter can smell. The hunter navigates to collect observations, which help him infer the Wumpus’s location; however, he cannot go near the Wumpus’s location. He has one arrow that he can shoot to kill the Wumpus from a distance, based on where he believes the Wumpus to be. In addition, for a second agent—the guide—the environment is fully observable. However, her communication to the hunter is minimal: she can only decide whether to point to a stench after the hunter observes it, without specifying what she hopes the hunter to do with that stench. This setting is directly inspired by the “broken stick” example, to capture the overloadedness and indirectedness of pointing. Here, although the state space is smaller than the classic Wumpus world, the inference is in fact computationally more expensive as the POMDP solver is recursively called by RSA.

Game

The guided Wumpus hunting game has an environment shown in Fig. 2. An agent plays this game with a POMDP solver. Components of the POMDP formulation are introduced below.

State Space The Wumpus can show up in one of the three possible locations: (0, 2), (1, 1), and (2, 0). Hunters will always start from (0, 0) and explore 3 possible locations for collecting observations: (0, 0), (0, 1), and (1, 0).

Action Space The hunter can choose from 4 possible actions: move vertically, move horizontally, shoot to the upper tile, or shoot to the right tile, which are the only possible directions given the map.

Transition Function The outcomes of the hunter’s actions are deterministic. At each step, he can decide to move or shoot. With moving, he will always move one tile in the direction of the action. If the hunter moves outside of the map, he will return to (0, 0). For shooting, the arrow will hit or miss the Wumpus depending on whether the Wumpus is in

the shoot direction. The game ends after shooting.

Reward Function For moving one step, there is an action cost, which is manipulated in our experiment. When shooting, the reward is set to 100 for hitting the Wumpus, -100 for missing.

Observation Space and Observation Function There are only two possible observations: a stench or nothing. The observation function is stochastic. If a tile is near the Wumpus, the hunter’s probability of observing a stench in that tile is 0.85. If a tile is not nearby the Wumpus, the probability of observing a stench is 0.15. In the classic Wumpus world, the observation is deterministic; however, we add stochasticity to make the belief inference more interesting, as in other POMDP tasks. We do not place pointing in the observation space as it is modeled as an communicative act, instead of an observation of the physical world (Kaelbling et al., 1998).

Experiment

We designed an experiment with the guided Wumpus hunting game to compare the models of individually perceived observations and jointly perceived observations. We manipulate the hunter’s model of interpreting the observation and the moving cost in the game.

We use the POMDP as the baseline model of a hunter who individually perceives the observation. In this condition, the hunter, as a receiver, will completely ignore the pointing signal sent by the guide.

We use Smithian pointing as the model of a hunter who pragmatically perceives the observation pointed to by the guide. The guide will use the Smithian pointing model to generate signals. Since the purpose of pointing is to help the receiver, we predict that hunters who use the Smithian pointing model will outperform those who use POMDPs.

It is possible that this predicted improvement in performance is simply caused by the amount of information provided by pointing but not the pragmatic inference process. To test this possibility, we add a third condition. In this condition, the hunter uses POMDPs to individually perceive observations, but when receiving a pointing signal, he receives an additional observation to which the pointing is directed. We call this condition POMDPs with “double observations”.

We also manipulate the cost of each step in the environment. The purpose of the guide pointing is to help the hunter achieve higher performance, but when the cost of moving is too high or low, help becomes unnecessary. If the moving cost is too high, for example -9, the hunter will shoot in a rush without considering the effect of observations. If the moving cost is too low, for example -1, the hunter will move around the environment more actively to collect observations no matter whether the guide points or not. The effect of communication as a function of moving cost may not be linear; therefore, in our experiment, we test 100 trials under each model using a moving cost of -9, -7, -5, -3, and -1. We expect to see the effect of pointing when the moving cost is moderate, but not when it is too high (-9) or too low (-1).




Figure 3: **Experimental results.** Black line denotes the ideal performance upper bound if the game was fully observable. Shaded areas represent 95% bootstrap confidence interval.

Result

The average reward across trials for each model under various moving cost is depicted in Fig. 3. Overall, the proposed Smithian pointing model achieves better performance compared to the classic POMDP model or the POMDP with double observations. The main effect of model type is significant ($F(2, 1485) = 9.602, p < 0.001$), and the main effect of moving cost is also significant ($F(4, 1485) = 19.535, p < 0.001$). However, the interaction between models and moving costs is not significant ($F(8, 1485) = 1.035, p = 0.407$). A post-hoc test with bonferroni correction shows that the Smithian pointing model achieves higher performance than “Double observation” condition ($F(1, 998) = 8.875, p = 0.009$). As hypothesized, our experiment also shows that the advantage stemmed from the pragmatic inference of pointing disappears when the task is too hard/easy for the guide to help. Specifically, when the moving cost is -1 or -9, the effect of model type is not significant ($F(2, 297) = 0.163, p = 0.850$; $F(2, 297) = 0.228, p = 0.796$). Taken together, these results demonstrate that pointing is relevant only when the signaler could offer help. Our computational model captures this relevance (Sperber & Wilson, 1986) and highlights how joint perception can be more powerful than individual perception of the same observation, demonstrated by the improved hunting performance.

Discussion

In this paper, we devise a computational model for pointing by defining the Smithian Value of Information (SVI), applying it to RSA to define the utility of pointing. In an example task, our pointing model shows a significant performance improvement compared to a single-agent POMDP or a single-agent POMDP with “double observation”. This improvement indicates that the advantage of pointing does not come from providing a new observation for individual perception. Instead, it comes from the pragmatic inference of how the jointly perceived observation is relevant to the task. Supporting the argument, our experiment also shows that the advantage of the Smithian pointing model works the best only when the receiver is in a position to be helped.

Our results suggest that Smithian coordination of beliefs is

necessary for modeling pointing. Seemingly simple, pointing requires the intelligent social capacity of mind coordination and altruistic motivation of helping. As a type of sophisticated social cognition, pointing involves identifying receiver's belief to predict action, evaluating actions with signaler's own belief, generating pointing to help change receiver's mind, and interpreting pointing with the assumption that the signaler is trying to help. The possible lack of some of these components, especially the cooperative motivation in pointing generation and interpretation, may explain the rarity of pointing in wild non-human primates. Indeed, the ability of generating and interpreting pointing is a milestone in human unique communication (Tomasello, 2010).

Due to the complexity of recursive reasoning of RSA, POMDP is solved multiple times in the recursion, which is computationally expensive and limits the current experiment setting to have a small state space. In follow up studies, we can use a faster POMDP solver or approximate the social recursion using computationally cheaper approaches (Kaelbling & Lozano-Pérez, 2013). In this way, we can expand our model to more complex tasks with larger state space and more observations. As many insights of this paper directly come from studies of child development, we hope this work will foster further interdisciplinary studies between developmental psychology and AI.

Acknowledgments

This research was supported by DARPA PA 19-03-01 and ONR MURI project N00014-16-1-2007 to TG.

References

Baker, C., Saxe, R., & Tenenbaum, J. (2011). Bayesian theory of mind: Modeling joint belief-desire attribution. In *Proceedings of the annual meeting of the cognitive science society* (pp. 2469–2474).

Butterworth, G., Simion, F., et al. (2013). *The development of sensory, motor and cognitive capacities in early infancy: From sensation to cognition*. Psychology Press.

Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998–998.

Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 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.

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Speech acts* (pp. 41–58). Brill.

Hepach, R., Benziad, L., & Tomasello, M. (2020). Chimpanzees help others with what they want; children help them with what they need. *Developmental Science*, 23(3), Article e12922.

Hume, D. (2018). *An enquiry concerning the principles of morals*. Yale University Press. (Original work published 1751)

Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1-2), 99–134.

Kaelbling, L. P., & Lozano-Pérez, T. (2013). Integrated task and motion planning in belief space. *The International Journal of Robotics Research*, 32(9-10), 1194–1227.

Knill, D. C., & Richards, W. (1996). *Perception as bayesian inference*. Cambridge University Press.

Leavens, D. A., & Hopkins, W. D. (1998). Intentional communication by chimpanzees: a cross-sectional study of the use of referential gestures. *Developmental Psychology*, 34(5), 813.

Liebal, K., Carpenter, M., & Tomasello, M. (2011). Young children's understanding of markedness in non-verbal communication. *Journal of Child Language*, 38(4), 888.

Liszkowski, U., Carpenter, M., Striano, T., & Tomasello, M. (2006). 12- and 18-month-olds point to provide information for others. *Journal of cognition and development*, 7(2), 173–187.

Martin, A., Lin, K., & Olson, K. R. (2016). What you want versus what's good for you: Paternalistic motivation in children's helping behavior. *Child development*, 87(6), 1739–1746.

Pineau, J., Gordon, G., & Thrun, S. (2003). Point-based value iteration: An anytime algorithm for pomdps. In *Proceedings of international joint conference on artificial intelligence* (pp. 1025–1032).

Russell, S. J., Norvig, P., & Davis, E. (2010). *Artificial intelligence: a modern approach* (3rd ed.). Prentice Hall.

Smith, A. (2010). *The theory of moral sentiments*. Penguin. (Original work published 1759)

Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Harvard University Press Cambridge, MA.

Tomasello, M. (2010). *Origins of human communication*. MIT press.

Ullman, T. D., Baker, C. L., Macindoe, O., Evans, O., Goodman, N. D., & Tenenbaum, J. B. (2009). Help or hinder: Bayesian models of social goal inference. In *Proceedings of the 22nd international conference on neural information processing systems* (pp. 1874–1882).

Van der Goot, M. H., Tomasello, M., & Liszkowski, U. (2014). Differences in the nonverbal requests of great apes and human infants. *Child Development*, 85(2), 444–455.

Wellman, H. M. (2014). *Making minds: How theory of mind develops*. Oxford University Press.

Wimmer, H., & Perner, J. (1983). Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition*, 13(1), 103–128.

Wittgenstein, L., & Anscombe, G. (2001). *Philosophical investigations: the german text, with a revised english translation* (3rd ed.). Blackwell. (Original work published 1953)