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

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). 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 from 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 evaluation. We then import SVI into rational speech act (RSA) as the utility of an utterance. These lead us to a pragmatic model of pointing allowing for contextually flexible interpretations. We test Smithian pointing model by extending the Wumpus world, adding another agent as a guide who can only help by pointing to an observation already perceived by the agent. Our results show that this severely limited and overloaded communication nevertheless significantly improves the hunters' performance.**

I. INTRODUCTION

Like all animals, we understand the world by collecting observations through our individual attention, 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 point 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) [1]. 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 [2]. This paper aims to computationally demonstrate that the joint perception from pointing is more potent than the individual perception of the same observation.

A. 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 [3]

Until Wittgenstein called attention to its underlying complexity, pointing had generally been perceived as an intuitive, unremarkable communicative gesture. It is among the most common forms of human communication and the first communicative gestures human infants learn to use [4], as early as one year old [5]. 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 [6], yet it lacks the cooperative properties of human pointing: apes are not bothered when the partner is distracted or non-responding [7]. 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, it is **overloaded**. As Wittgenstein [3] 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.

The overloadedness and indirectness of pointing enable it to express manifold meanings with the same observation. The meaning of the pointing can be interpreted depending on the context. Based on the context, young children could adaptively decide what to do with the toy pointed to by an adult, put it away or examine it [8]. Context-dependent interpretation 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 through pointing can evoke richer inferences than the individual perception of the same observation [9].

As in the above example, a crucial function of pointing is helping, but with two unique characteristics. First, the helper can 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 [10], which we will introduce later. Second, pointing only changes the helpee's mind. Unlike instrumental actions, it does not change the physical states at all, so models of instrumental helping would fail to apply [11]. 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 [12], [13]. 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.

B. Smithian Empathy and Helping

The concept of Smithian helping is based on Adam Smith's discussion of two types of empathy. First, he addressed Hume's definition that empathy is a **direct mirroring** of another's mindset [14]. One can imagine a theoretical example involving you and your friend preparing for your friend's performance in the school talent show. While your friend is excited to perform, you know that he is objectively bad at singing. When executing Hume's empathy, you must abandon your personal opinion and take on his perspective, which is the excitement of performing.

In contrast to this conventional definition, Smith proposed that true empathy involves **coordinating mindsets** between the empathizer and the agent being empathized with [15]. In the talent show example, when executing the Smithian empathy, you maintain your own mindset and use it to evaluate the situation of your friend performing. Because you know your friend is bad at singing, your evaluation based on your belief leads you to worry that your friend's performance on stage may not end up well, which is against his belief.

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 [10]. In the talent show example, if you use Smithian empathy, you should help your friend by convincing him not to perform. This is an example of paternalistic helping as it opposes his desire to perform. Paternalistic helping is arguably more complicated than the ability to reason about other's beliefs or to use other's belief to predict actions as in the famous false belief task [16]. In paternalistic helping, one needs to coordinate two mindsets: **predict** action with others' beliefs and **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 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 [10]. In addition, when the wrong tool is requested, children offer the tool they believe the requester needed, not the one they are asked for [17]. These studies provide a solid theoretical foundation on how to coordinate two beliefs in helping. We build upon these insights to focus on cases in which helping behavior is executed to provide information through communication, as in pointing.

II. Modeling Individual vs. Joint Perception

Our study is directly inspired by the perception of the broken stick in the hunting example. We aim to demonstrate that an intelligent agent with joint perception through pointing can outperform an intelligent agent only with individual perception in a hunting task. We start by outlining the model of an intelligent agent with individual perception. On top of that, we formulate a model of pointing for joint perception.

A. Modeling Agent with Individual Perception: POMDP

We use the partially observable Markov decision process (POMDP) [1] to model an agent with individual perception. POMDP models two processes in how an agent interacts with an uncertain environment with only limited observations.

The first process is how an agent uses Bayesian inference to update its belief of the world state s, treating the observation o received as data. A belief is defined as a probabilistic distribution over the set of possible states [1], [18]:

$$b\left(s\right) = P\left(s|b\right). \tag{1}$$

The Bayesian belief update only involves individual perception because the agent treats the observation only as being generated by its own interaction with the environment, not a referent of any communicative intention as in [19].

The second process is the decision model of how an agent takes rational actions based on its belief, which involves the calculation of expected utility of each action. The expected utility of an action a can be defined as the expectation of utility of the outcomes s' of the action [20]:

$$\mathbb{E}\, U(a|b) = \int_{s' \in \mathcal{S}} U(s') \int_{s \in \mathcal{S}} P\left(s'|s,a\right) P\left(s|b\right), \tag{2}$$

The agent chooses its action distribution $P(a|b)$ to maximize its expected utility. 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 [21] as the solver for POMDP.

B. Modeling Agent with Joint Perception: Smithian Pointing

For an individual agent, the expected utilities 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 [20]. There are no subscripts in (2), which implies that the action and belief are from the same agent. Equation (2) 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. Our Smithian model of pointing augments the 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.

a) Smithian Utility of Action: We define Smithian utility of actions by augmenting the conventional utility of actions defined in (2). With Smithian empathy, the signaler should use her own utility function to evaluate the outcomes.

The utility function of the signaler and the receiver may be different in some cases; but constrained by Smithian helping, the signaler's utility should always be aligned with the receiver's physical well-being. Here we use 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 (2) to the belief of the signaler:

$$\mathbb{E}U_{Smith}\left(a_{Rec}|b_{Sig}\right)=\int_{s'\in\mathcal{S}}U_{Sig}\left(s'\right)\int_{s\in\mathcal{S}}P\left(s'|s,a_{Rec}\right)P\left(s|b_{Sig}\right). \quad (3)$$

b) 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. It can be derived from the expected utility of actions as an agent's belief determines the distribution of its action $P(a|b)$. Using theory of mind, A can predict B's distribution of action $P(a|b)$ based on B's belief [22], hence evaluate the belief by integrating out the evaluation of actions. Let $\mathcal{A}$ be the set of possible actions of the receiver. We can derive Smithian utility of belief by integrating over the Smithian utility of actions in (3):

$$U_{Smith}\left(b_{Rec}|b_{Sig}\right)=\int_{a_{Rec}\in\mathcal{A}_{Rec}}P\left(a_{Rec}|b_{Rec}\right)\mathbb{E}U_{Smith}\left(a_{Rec}|b_{Sig}\right). \quad (4)$$

c) Smithian Value of Information (SVI): Pointing carries information. In the field of AI, the value of information is "the difference in expected value between best actions before and after the information is obtained" [20]. Here we adopt this insight to define Smithian Value of Information (SVI), but with two significant differences. First, we calculate the value of information from the perspective of the signaler, not the receiver of the information in the definition in [20]. 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 from the pointing. In other words, **SVI is the utility of pointing**.

To later incorporate 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:

$$b'_{Rec}=P_{Rec}(s|u)=P\left(s|b_{Rec},u\right). \quad (5)$$

We can write SVI as:

$$\mathrm{SVI}\left(u|b_{Sig}\right)=U_{Smith}\left(b'_{Rec}|b_{Sig}\right)-U_{Smith}\left(b_{Rec}|b_{Sig}\right). \quad (6)$$

d) Pointing as a Rational Speech Act: With the utility of pointing clearly defined as SVI, we can treat pointing as an utterance and model its use and interpretation using RSA. We can write down how a signaler generates pointing u:

$$P_{Sig}\left(u|b_{Sig}\right)\propto\exp\{\alpha\left[\mathrm{SVI}\left(u|b_{Sig}\right)-c\left(u\right)\right]\}. \quad (7)$$

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

We can also consider the receiver in RSA with the POMDP challenge as the "context" of communication. If the signaler knows the true state is s, her belief b_{Sig} can be reduced to a single state s, representing 1 on the state s, 0 on other states. A pragmatic receiver updates its belief upon receiving an utterance u with Bayesian inference [13], [23]:

$$b'_{Rec}\left(s\right)\propto P_{Sig}(u|s)P_{Rec}(s|b_{Rec}). \quad (8)$$

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 (8) and selects the signal that maximizes SVI in (7). 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. enabling the receiver to update its belief using Bayes' rule. This recursive social reasoning is based on RSA but uses the POMDP agent as the literal receiver. In (4), $P(a_{Rec}|b_{Rec})$ is the policy solved by POMDP.

III. Modeling Experiment

A. Task: Guided Wumpus Hunting

To highlight the strengths of the Smithian pointing model, we pick a simplified version of the classic single-agent partially observable task in AI: *the Wumpus world* [20]. We augment the task 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.

The guided Wumpus hunting game has an environment shown in Fig. 1. A hunter navigates through a set of tiles to shoot a stationary monster, the Wumpus, without knowing its exact location. 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, (0, 0), (0, 1), and (1, 0), for collecting observations to help him infer the Wumpus's location. When the hunter moves, 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 moving one step, there is an action cost, which is manipulated in our experiment.

The hunter can collect two kinds of observation: one is the "stench" that the Wumpus emits when it is nearby; the other is nothing, which means no "stench" is observed. 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. We add observation stochasticity to the Wumpus world to make the belief inference more interesting, as in other POMDP tasks [1]. Here we do not include pointing as an observation of the physical world, but model it as a communicative act.

The hunter has one arrow to shoot the Wumpus from a distance, based on where he believes the Wumpus to be, but he cannot go near the Wumpus's location. If he shoots, the arrow will hit or miss the Wumpus depending on whether the Wumpus is in the shoot direction. The shooting reward is set to 100 for hitting the Wumpus, -100 for missing. The game ends after shooting.

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. Directly inspired by the "broken stick" example, this setting captures

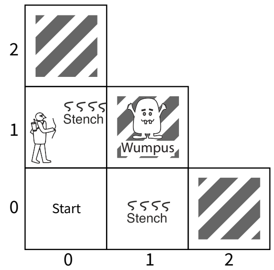

Fig. 1: **The environment of the guided Wumpus hunting game.** Wumpus can only show up in one of three shaded tiles. Starting from (0, 0), the hunter tries to infer Wumpus's location from the stench and shoot it.

the overloadedness and indirectedness of pointing. Here, although the state space is smaller than the classic Wumpus world, the inference is computationally more expensive as the POMDP solver is recursively called by RSA.

B. 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. The hunter, as a receiver, will 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 uses the Smithian pointing model to generate signals. We predict that hunters using Smithian pointing model will outperform those using 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 where 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 guide points to help the hunter, but when the cost of moving is too high or low, help becomes unnecessary. If the moving cost is too high, the hunter will shoot in a rush without considering the effect of observations. If the moving cost is too low, the hunter will move around more actively to collect observations whether the guide points or not. The effect of communication as a function of moving cost may not be linear; therefore, 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).

IV. Result

The average reward across trials for each model under various moving cost is depicted in Fig. 2. 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 ($F(2, 1485) = 9.602$, $p < 0.001$). The main effect of moving cost is 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 shows that the Smithian pointing model achieves higher performance than "Double observation" condition ($F(1, 998) = 8.875$, $p = 0.003$). As hypothesized, our experiment also shows that the advantage of Smithian pointing model disappears when the task is too hard/easy. Specifically, the effect of model type is not significant when the moving cost is -1 ($F(2, 297) = 0.163$, $p = 0.850$) or -9 $F(2, 297) = 0.228$, $p = 0.796$). Taken together, these results demonstrate that pointing is relevant only when the signaler could offer help. Our computational model captures this relevance [9] and highlights how joint perception can be more powerful than individual perception of the same observation, demonstrated by the improved hunting performance.

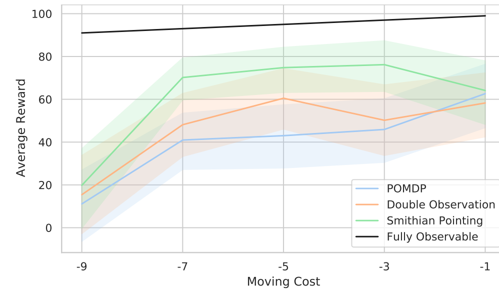

Fig. 2: **Experimental results.** Black line denotes the ideal performance upper bound if the game was fully observable. Shaded areas represent 95% bootstrap confidence interval

V. Discussion

In this paper, we devise a computational model for pointing by defining the Smithian Value of Information (SVI), applying it to RSA as the utility of pointing. In an example task, our pointing model shows a significant performance improvement compared to a single-agent POMDP. 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.

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 [24]. In this way, we can expand our model to a more complex task with more states and 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.

REFERENCES

[1] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, "Planning and acting in partially observable stochastic domains," *Artificial Intelligence*, vol. 101, no. 1-2, pp. 99–134, 1998.

[2] D. C. Knill and W. Richards, *Perception as Bayesian inference*. Cambridge University Press, 1996.

[3] L. Wittgenstein and G. Anscombe, *Philosophical investigations*, 1953.

[4] G. Butterworth, F. Simion, and G. Butterworth, "What is special about pointing in babies," *The development of sensory, motor and cognitive capacities in early infancy: From perception to cognition*, pp. 171–190, 2013.

[5] U. Liszkowski, M. Carpenter, T. Striano, and M. Tomasello, "12- and 18-month-olds point to provide information for others," *Journal of Cognition and Development*, vol. 7, pp. 173–187, 2006.

[6] D. A. Leavens and W. D. Hopkins, "Intentional communication by chimpanzees: a cross-sectional study of the use of referential gestures." *Developmental Psychology*, vol. 34, no. 5, p. 813, 1998.

[7] M. H. Van der Goot, M. Tomasello, and U. Liszkowski, "Differences in the nonverbal requests of great apes and human infants," *Child Development*, vol. 85, no. 2, pp. 444–455, 2014.

[8] K. Liebal, M. Carpenter, and M. Tomasello, "Young children's understanding of markedness in non-verbal communication," *Journal of Child Language*, vol. 38, no. 4, p. 888, 2011.

[9] D. Sperber and D. Wilson, *Relevance: Communication and cognition*. Harvard University Press Cambridge, MA, 1986.

[10] A. Martin, K. Lin, and K. R. Olson, "What you want versus what's good for you: Paternalistic motivation in children's helping behavior," *Child development*, vol. 87, no. 6, pp. 1739–1746, 2016.

[11] T. Ullman, C. L. Baker, O. Macindoe, O. Evans, N. D. Goodman, and J. B. Tenenbaum, "Help or hinder: Bayesian models of social goal inference," in *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2009.

[12] M. C. Frank and N. D. Goodman, "Predicting pragmatic reasoning in language games," *Science*, vol. 336, no. 6084, pp. 998–998, 2012.

[13] N. D. Goodman and M. C. Frank, "Pragmatic language interpretation as probabilistic inference," *Trends in Cognitive Sciences*, vol. 20, no. 11, pp. 818–829, 2016.

[14] D. Hume, "An enquiry concerning the principles of morals." 1751.

[15] A. Smith, *The Theory of Moral Sentiments*, 1759.

[16] H. Wimmer and J. Perner, "Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception," *Cognition*, vol. 13, no. 1, pp. 103–128, 1983.

[17] R. Hepach, L. Benziad, and M. Tomasello, "Chimpanzees help others with what they want; children help them with what they need," *Developmental Science*, vol. 23, no. 3, p. e12922, 2020.

[18] C. Baker, R. Saxe, and J. Tenenbaum, "Bayesian theory of mind: Modeling joint belief-desire attribution," in *Proceedings of the Annual Meeting of the Cognitive Science Society (CogSci)*, 2011.

[19] H. P. Grice, "Logic and conversation," in *Speech Acts*. Brill, 1975, pp. 41–58.

[20] S. Russell and P. Norvig, *Artificial intelligence: a modern approach*, 2002.

[21] J. Pineau, G. Gordon, and S. Thrun, "Point-based value iteration: An anytime algorithm for pomdps," in *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, 2003.

[22] H. M. Wellman, *Making minds: How theory of mind develops*. Oxford University Press, 2014.

[23] N. D. Goodman and A. Stuhlmüller, "Knowledge and implicature: Modeling language understanding as social cognition," *Topics in Cognitive Science*, vol. 5, no. 1, pp. 173–184, 2013.

[24] L. P. Kaelbling and T. Lozano-Pérez, "Integrated task and motion planning in belief space," *International Journal of Robotics Research (IJRR)*, vol. 32, no. 9-10, pp. 1194–1227, 2013.