

Emergence of Pragmatics from Referential Game between Theory of Mind Agents

Paper #928

ABSTRACT

Pragmatics studies how context can contribute to language meanings [40]. In human communication, language is never interpreted out of context, and sentences can usually convey more information than their literal meanings [31]. However, this mechanism is missing in most multi-agent systems [4, 26, 36, 45], restricting the communication efficiency and the capability of human-agent interaction. In this paper, we propose an algorithm, using which agents can spontaneously learn the ability to “read between lines” without any explicit hand-designed rules. We integrate theory of mind (ToM) [2, 42] in a cooperative multi-agent pedagogical situation and propose an adaptive reinforcement learning (RL) algorithm to develop a communication protocol. ToM is a profound cognitive science concept, claiming that people regularly reason about other’s mental states, including beliefs, goals, and intentions, to obtain performance advantage in competition, cooperation or coalition. With this ability, agents consider language as not only messages but also rational acts reflecting others hidden states. Our experiments demonstrate the advantage of pragmatic protocols over non-pragmatic protocols. We also show the teaching complexity following the pragmatic protocol empirically approximates to recursive teaching dimension (RTD).

KEYWORDS

Emergent Behavior; Pragmatics; Cooperative Games; Theory of Mind

1 INTRODUCTION

The study of emergent languages has become an important topic in cognitive science and artificial intelligence for years, as effective communication is the prerequisite for the successful cooperation in both human society [10, 33] and multi-agent systems (MAS) [20, 26, 37]. Communication, as Grice characterized in his original model of pragmatics [30], should follow cooperative principles, where listeners make exquisitely sensitive inferences about what utterances mean given their knowledge of the speaker, the language, and the context, and the speaker collaboratively makes contributions to the conversational goals [29, 46]. For example, if there are three boys in a class and teacher *A* told teacher *B* that “**some** of the boys went to the party”, then teacher *B* will legitimately infer that one or two of the boys went to the party [30, 34]. Although according to the literal meaning, it could be that all three of the boys went to the party, teacher *B* is less likely to interpret the utterance that way, because teacher *A*, a cooperative speaker, would

have said “**all** of the boys went to the party” in that case to prevent ambiguity. This seemingly trivial example (called scalar implicature in pragmatics) illustrates that pragmatic conversation is, in most of the time, taken for granted in human communication, and shows how significant hidden information can be acquired from literal meanings. However, current MAS tend to model communication merely as information exchange between agents, among which messages are deciphered only by their literal meanings [4, 26, 36, 45]. Even with perfect understanding among each other, this type of communication cannot achieve optimal efficiency, as the intention of the communicator implicitly suggested by the message is ignored.

To perform pragmatic communication, the listener needs to not only comprehend the speaker’s utterance but also infer his mental states. It has been shown that humans, during an interaction, can reason about others’ beliefs, goals, intentions and predict opponent/partner’s behaviors [2, 42, 51], a capability called ToM. In some cases, people can even use ToM recursively, and form beliefs about the way others reason about themselves [15]. Thus, to collaborate and communicate with people smoothly, artificial agents must also bear similar potentially recursive mutual reasoning capability. Despite the recent surge of multi-agent collaboration modeling [12, 22, 36, 45], integrating ToM is still a nontrivial challenge. A few approaches attempted to model the nested belief of other agents in general multi-agent systems, but extensive computation restricts the scale of the solvable problems [17, 32]. When an agent has an incomplete observation of the environment, it needs to form a belief, a distribution over the actual state of the environment, to take actions [32, 51]. ToM agents, besides their own beliefs about the state, or 0-th level beliefs, also model other agents’ beliefs, forming 1-st level beliefs. They can further have beliefs about others’ 1-st level beliefs about their 0-th level belief, so on and so forth [13–15, 17, 51]. The intractability of distribution over distribution makes exact solving for ToM agents’ nested beliefs extremely complicated [17].

Therefore, an approach to acquire the sophistication of high-level recursions without getting entangled into the curse of intractability is needed. In this paper, we propose an adaptive training process, following which pragmatic communication protocol can emerge between cooperative ToM agents modeling only the 1-st level belief over belief. The complexity of higher level recursions can be preserved by the dynamic evolving of agents’ tractable belief estimation functions. We don’t assume agents have a certain level of recursions, which requires modeling nested beliefs from the 0-th level up to the desired level [13–15, 17]. Instead, we directly learn a function to approximate partners’ actual beliefs and how to react accordingly. In cooperative games, this learning becomes mutual adaptation, with a controlled exploration rate, improving the performance of the multi-agent system [11]. Intuitively, for a pair of agents, we update them alternatively, namely, fixing one




Figure 1: An example referential game. There are three objects, a blue sphere, a red sphere, and a blue cone. If a student hears “blue” from the teacher, he should be able to identify the blue sphere instead of the blue cone.

while training the other, simulating the iterative best response (IBR) model, which proves to converge to a fixed point Nash equilibrium in strong and weak interpretation game [34]. We demonstrate the effectiveness and advantage of the pragmatic protocol with referential games, a communication game widely used in linguistic and cognitive studies in the context of language evolution [7, 37, 38]. It provides a good playground for pragmatic and pedagogical interactions between a teacher and a student and can easily generalize to a comprehensive teaching task with a large concept space.

We tested our algorithm and evaluated the pragmatic teaching protocol with both symbolic and pixel data, and achieved significant performance gain over previous algorithms in terms of referential accuracy. Also, we found that if messages are grounded with human expertise before interaction, the emerged protocol achieves teaching complexity empirically approximates recursive teaching dimension (RTD) [16], the worst-case number of examples required for cooperative agents concept learning [8].

This paper makes two major contributions: 1) we showed that a pragmatic communication protocol can emerge between ToM agents through adaptive reinforcement learning. This protocol significantly improves the cooperative multi-agent communication by enabling the agents to extract hidden meanings from the context to enrich the literal information; 2) we proposed an algorithm to develop protocols empirically approximating the teaching complexity bound between cooperative agents provided by RTD. Extensive experiments on both the symbolic and 3D object datasets demonstrate the effectiveness of our proposed protocol.

2 BACKGROUND

2.1 Referential Game

There are a teacher and a student in a referential game. The teacher has a target in mind and aims to send a message to the student so that the student can identify the target out of a set of distractors after receiving this message. The motivation of our algorithm is the rational speech act (RSA) model between ToM agents (also termed as bilateral optimality [5]): in order to establish proper communication, the speaker has to take into account the perspective of the listener, while the listener has to take into account the perspective of the teacher [29]. Figure 1 shows an example. There are three objects, a blue sphere, a red sphere, and a blue cone. Suppose the target is the blue sphere. If the only allowed messages are colors and shapes, then, for a literal student, there is no unique identifier for

the blue sphere, because both “blue” and “sphere” have more than one consistent candidates. Nonetheless, a pragmatic student, after hearing “blue” from the teacher, should be able to do counterfactual reasoning and identify the blue sphere instead of the blue cone, because he knows the teacher is helpful and would have used cone to refer to the blue cone unambiguously. During pragmatic communication, a message conveys more information than the message itself. The usage of that message can usually suggest the intention of the teacher.

The referential game can be formally defined by a tuple $\langle A, B, \Omega, \mathcal{M}, \mathcal{A} \rangle$, where A and B stand for a teacher and a student. Ω is the instance space, where the distractors and targets are sampled from. \mathcal{M} is the message space and \mathcal{A} is the student’s action space. In a specific game, a set of instances $O \subseteq \Omega$ is sampled from Ω as candidates, and one of the candidates $o^* \in O$ is designated as the target, while the rest, $O/\{o^*\}$, are distractors. The candidates O are available to both of the agents, while only the teacher knows the target, o^* . Agents take turns in this game. In every round, the teacher first sends a message $m_t \in \mathcal{M}$ to the student, followed by an action $a_t \in \mathcal{A} = \{1, 2, \dots, |O|, \Xi\}$ taken by the student, where number 1 to $|O|$ represent “identify a certain instance as the target”, Ξ means “wait for next message” and t stamps the t -th round. Every message comes with a message cost, c_m , and the total gain for both the teacher and the student, given a_T the first non Ξ action, is $R = \sum_{t=1}^T -c_{m_t} + 1(o^* = O[a_T])$. Notice that the game ends when the student performs a non Ξ action. We define a protocol between A and B as a set of policies

$$\begin{aligned} \Pi &= \langle \pi_A : \mathbb{P}(\Omega) \times O \times \mathcal{M}^* \times \mathcal{M} \rightarrow [0, 1], \\ \pi_B &: \mathbb{P}(\Omega) \times \mathcal{M}^* \times \mathcal{M} \times \mathcal{A} \rightarrow [0, 1] \rangle \end{aligned}$$

$\mathbb{P}(\Omega)$ is the power set of Ω , where O sampled from, and $*$ is the kleene star, standing for the history of message. Intuitively, the teacher selects a message based on the distractors, the target and communication history. The student chooses an action according to candidates, history and the latest message. The goal for both of the agents is to maximize the expected gain:

$$\mathbb{E}_{\substack{O \sim \mathbb{P}(\Omega), o^* \sim O, \\ m_{1:T} \sim \pi_A, a_{1:T} \sim \pi_B}} \left[- \sum_{t=1}^T c_{m_t} + 1(o^* = O[a_T]) \right] \quad (1)$$

2.2 Recursive Teaching Dimension

In computational learning theory, the teaching dimension (TD) measures the minimum number of examples required to uniquely identify all concepts within a concept class. We draw a comparison between identifying a concept using examples and identifying the target using messages. Formally, given an instance space $X = \{1, \dots, n\}$ and a concept class $C \subseteq \{0, 1\}^n$, a teaching set for a concept $c \in C$ with respect to C is a subset $S \subseteq X$ such that $\forall c' \neq c \in C \quad \exists x \in S \text{ s.t. } c(x) \neq c'(x)$. Intuitively, the teaching set is a subset of instance space that can uniquely identify c in C . TD of c in C , denoted by $TD(c, C)$, is the size of the smallest teaching sets. TD of the whole concept class is $TD(C) = \max_{c \in C} TD(c, C)$. However, this definition of TD often overestimates the number of examples needed for cooperative agents, who teach and learn using “helpful” examples. For instance, the game in figure 1 has $TD = 2$, but all concepts can be taught with one message.

The recursive teaching dimension (RTD), a variation of the classic TD, can model the behavior between cooperative agents. Define the teaching hierarchy for C as the sequence $((C_1, d_1), \dots, (C_h, d_h))$ such that for all $j \in \{1, \dots, h\}$, $C_j = \{c \in \bar{C}_j | d_j = TD(c, \bar{C}_j) \leq TD(c', \bar{C}_j), \forall c' \in \bar{C}_j\}$ where $\bar{C}_1 = C$ and for all $i \in \{1, \dots, h-1\}$, $\bar{C}_{i+1} = C \setminus (C_1 \cup \dots \cup C_i)$. RTD of c in C is defined as $RTD(c, C) = d_j$. RTD of C is $RTD(C) = \max\{d_j | 1 \leq j \leq h\}$. Intuitively, one can first learn simplest concepts, C_i , and remove them from the concept class. Then, for the remaining set of concepts, \bar{C}_{i+1} , recursively learn the simplest concepts and so on. RTD lower bounds the classic teaching dimension, i.e. $RTD(C) \leq TD(C)$ [16]. Also, $RTD(C)$ measures the worst-case number of labeled examples needed to learn any target concept in C , and the teaching hierarchy can be derived by the teacher and student separately without communication [8].

3 RELATED WORK

Pragmatics [30] has profound linguistic and cognitive science origin dating back to 1970s. However, the integration of this topic with multi-agent games only draws people's attention in this decade. The RSA model was raised by Golland et al. and developed in [1, 25, 29, 44] in later works. Yet, all of these works model only single utterance and require hand designed production rules for the agents, while in our algorithm, all policies are learned by the agents and multi-round of communication is allowed. [24, 35, 46] model pragmatic reasoning for human action prediction, but they all require domain-specific engineering, and pragmatically annotated training data. In our work, all models are trained by self-playing with RL. [46] needs human data to initialize the literal speaker and uses RL to train listener with pragmatic reasoning, but they assume fixed literal speaker. On the contrary, we endow both the teacher and the student ToM and allow them to do mutual pragmatic reasoning. The discriminative best response algorithm in [47] is also inspired by the IBR model [34], but they used supervised learning given language data instead of RL. Also, our game states are more complicated than their toy examples. Integration of ToM provides new insights to MAS, with most usage focusing on predicting agents' actions from a 3rd person view [41, 43, 50]. Yet, our agents combine ToM into their policy making.

Emergence of language in communication games. Emergence of language is a topic about a group of agents developing a communication protocol to complete a task, which can either be cooperative or competitive. In most recent studies, agents start with a set of ungrounded symbols and first learn to ground these symbols using reinforcement learning approaches supervised by the task rewards. The major novelty of our work comparing with previous methods is that pragmatic reasoning can emerge from our protocols without explicit rules coded.

In [37], two agents try to develop a protocol through playing the referential game, in which the teacher sees only the target but no distractors, eliminating the possibility of taking advantage of ToM, as no counterfactual reasoning can happen at the student side. In [38] the teacher sends a message according to the context of candidates, but no student reaction is simulated before the selection. In [6, 9], the teacher can simulate the reaction of a fixed student, who does not model teacher's mind and cannot benefit from counterfactual reasoning about teacher's intention behind messages, terminating the recursive process after only one cycle.

A variation of the referential game was played in [18], where the teacher sees an image and the student sees a set of descriptions. The goal is for the student to identify a suitable description for the teacher's image by multi-round communication. A similar multi-round communication game using natural language was also played in [12]. In these games, the student can ask further questions after receiving the first message and is in charge of the final decision making, while the teacher will answer the question based on the target and the communication history. These works focus on learning shared embedding between objects and messages instead of grounding messages to attributes that reappear across the object space. Also, neither of these papers include agents who model their partner's minds.

Multiagent communication. Modelling multiagent communication dates back to [1998] when Kinney and Tsatsoulis proposed adaptive learning of multiagent communication strategy as a pre-defined rule-based control system. To scale up from rule-based systems, decentralized partially observable Markov decision process (DEC-PODMP) was used to model multiagent interaction with communication as a special type of action among agents [4, 27]. Solving DEC-POMDP exactly is a NEXP-complete problem [4] requiring agents to remember the complete observation history for their policy. In recent works, a more compact representation of the history is often used for action selection. In [20, 45], agents maintain a memory variable and refer to it when act and speak. However, the centralized training process in these work needs channels with large bandwidth to pass gradients across different agents. Also, the messages in [45] are not discrete symbols, but continuous outputs of neural networks. Other extension of single agent value based algorithms to multiagent problems [20, 39] usually suffers from non-stationarity induced by simultaneous updates of agents.

While DEC-POMDP considers other agents as part of the environment and learns the policy as a mapping from local observation to action, LOLA in [21] learns the best response to evolving opponents. Yet, opponent/partners' real-time belief is not considered into policy. Interactive-POMDP (I-POMDP) [17, 32] moves one more step forward by actually modeling opponents' mental states at the current moment and integrates other's belief into the agent's own policy. However, I-POMDP requires extensive sampling to approximate the nested integration over the belief space, action space and observation space, limiting its scalability. Because whenever the teacher sends a message, she thinks about not only the student's current belief but also current distractor sets, the value iteration process in I-POMDP needs to be repeated for every single game. In our algorithm, training only needs to be performed once for one pair of agents and reusable to all games. The Bayesian action decoder (BAD)-MDP proposed by Foerster et al. also yields counterfactual reasoning in their belief update, but their method is more centralized in the testing process than ours. The BAD-agent is a super-agent controlling all other agents collectively. Deterministic partial policies can easily reveal agents' private information to the BAD-agent and make it public. Instead, our model doesn't depend on any implicit information flowing between agents during testing.

4 METHOD

4.1 Adaptive Emergence of Pragmatic Protocol

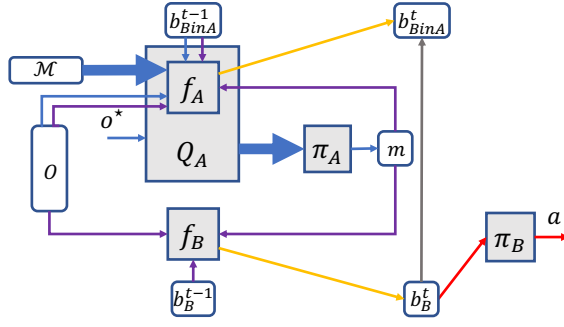


Figure 2: ToM Agents Interaction Pipeline. First, the teacher chooses a message according to the context and her prediction of the student's reaction (blue arrows). After a message is sent, the student updates his belief and the teacher updates her estimation of student's belief (purple and orange arrows). Then, the student either waits or selects a candidate (red arrows). Only in the training phase, the actual student belief will be returned to the teacher (gray arrow). Bold arrows stand for the whole message space being passed. Notice f_A is part of the Q_A . O and b_{BinA}^{t-1} are passed to f_A twice, for message selection and teacher's new belief estimation. Empty boxes are game and time variants; shadowed boxes are agents' mental structures. Notations are introduced in section 4.1 with θ omitted from subscripts.

Emergence of Pragmatic Protocol: Our goal is to learn a protocol for agent A and B so that they can communicate with the contextual information being considered. To avoid tracking the message history, which scales exponentially with the time, we use beliefs as sufficient statistics for the past. Hence, ToM can be embodied as estimating partner's current and future belief, then choose the most ideal action to manipulate them as needed. In the referential game, since the teacher knows the target, only the student holds a belief, b_B , about the target. Inspired by the overver technique, a communication philosophy, motivating an agent to simulate others' reactions to its own messages and then select the one leading to the best outcome [6, 9], we let the teacher holds a belief b_{BinA} as her estimation of student's belief, which enables her to considerate student's future actions. As b_{BinA} is an estimation of the student's belief, it should be a belief over belief, i.e. a distribution over a distribution over the candidates. Since a distribution is a continuous random variable, distribution over a continuous variable can be represented as a set of particles. However, to avoid the complexity we only use one particle to approximate. That is, b_{BinA} is still a distribution over the candidates. This is reasonable because the belief update process is deterministic for rational agents following the Bayesian rule [19, 46]. Given b_B^0 a uniform distribution over candidates, $P(b_B^t)$ is a single mode distribution and can be approximated with a particle.

Before speaking, teacher traverses all messages and predicts the student's new belief after receiving each message. She then sends the message leading to the most optimal student's new belief. Hearing the message, student updates his belief and takes action. This process is visualized in figure 2 and formalized in algorithm 1 line 11 to 20. The recursive mutual modeling in ToM is integrated

Algorithm 1 Iterative Adaption Protocol Emergence

```

1: Randomly initialize  $\theta_A, \theta'_A, \theta_B$ 
2: No. candidates  $K$ 
3: Learning rate  $\eta$ , Batch size  $N$ 
4: for each phase do
5:   for  $i \in \{A, B\}$  do
6:     Initialize replay buffer  $\mathcal{D} \leftarrow \emptyset$ 
7:   while train agent  $i$  do
8:      $t = 1$ 
9:     Initialize  $\mathcal{E} \leftarrow \emptyset$ 
10:    repeat
11:      if  $t = 1$  then
12:        Sample  $O = \{\omega_1, \dots, \omega_K\}$ 
13:        Random select  $o^* = \omega_j$ 
14:        Initialize  $b_B^0, b_{BinA}^0$  as uniform distribution
15:      end if
16:       $m_t \sim \pi_{\theta_A}(\cdot | O, o^*, b_{BinA}^{t-1})$ 
17:       $b_B^t = f_{\theta_B}(O, b_B^{t-1}, m_t)$ 
18:       $a_t \sim \pi_{\theta_B}(b_B^t)$ 
19:       $r_t = -c m_t + 1(O[a_t] = o^*)$ 
20:       $b_{BinA}^t = b_B^t$ 
21:      if  $i = A$  then
22:         $\mathcal{D} \leftarrow \mathcal{D} \cup \{(O, o^*, b_{BinA}^{t-1}, m_t, b_B^t, r)\}$ 
23:      else
24:         $\mathcal{E} \leftarrow \mathcal{E} \cup \{(O, b_B^{t-1}, m_t, a_t, r)\}$ 
25:      end if
26:       $t \leftarrow t + 1$ 
27:    until  $a_t \neq \Xi$ 
28:    if  $i = A$  then
29:      Sample  $\{(O, o^*, b_{BinA}^{t-1}, m_t, b_{BinA}^t, r)\}_N \sim \mathcal{D}$ 
30:       $\xi = r + \gamma \arg \max_m Q_{\theta'_A}(O, o^*, b_{BinA}^t, m)$ 
31:       $L^Q = \frac{1}{N} \sum_N \|\xi - Q_{\theta_A}(O, o^*, b_{BinA}^{t-1}, m_t)\|^2$ 
32:       $L^{Obv} = \frac{1}{N} \sum_N H(b_{BinA}^t, f_{\theta_A}(O, b_{BinA}^{t-1}, m_t))$ 
33:       $\theta_A \leftarrow \theta_A - \eta \nabla_{\theta_A}(L^Q + \lambda L^{Obv})$ 
34:      Update  $\theta'_A \leftarrow \theta_A$  periodically
35:    else
36:      Compute  $R_t = \sum_{k=t}^{|\mathcal{E}|} \gamma^{k-t} r_k$  for  $t$  in  $1, \dots, |\mathcal{E}|$ 
37:       $J = \frac{1}{|\mathcal{E}|} \sum_{t=1}^{|\mathcal{E}|} \log \pi_{\theta_B}(a_t | f_{\theta_B}(O, b_B^{t-1}, m_t)) R_t$ 
38:       $\theta_B \leftarrow \theta_B + \eta \nabla_{\theta_B} J$ 
39:    end if
40:  end while
41: end for
42: end for

```

within the belief update process. $f_{\theta_i}, i \in \{A, B\}$ are belief update functions parameterized by θ_i , taking in candidates, current belief, message and returning a new belief. The beliefs in our model are semantically meaningful hidden variables in teacher's Q-function and student's policy network, as the student directly samples an action according to his belief. The evolving of the belief update function reflects the protocol dynamics between the agents. Within f , we code in the Bayesian rule with the likelihood function varying across different training phases. Our implementation detail can be found in section 4.2. In each phase, we first train the teacher for a fixed student, then adapt the student to the teacher.

Difference from multiagent Q-learning: Our algorithm considers both the physical state and agent's mental state in the value function, and has a dynamic belief update function. Moreover, since

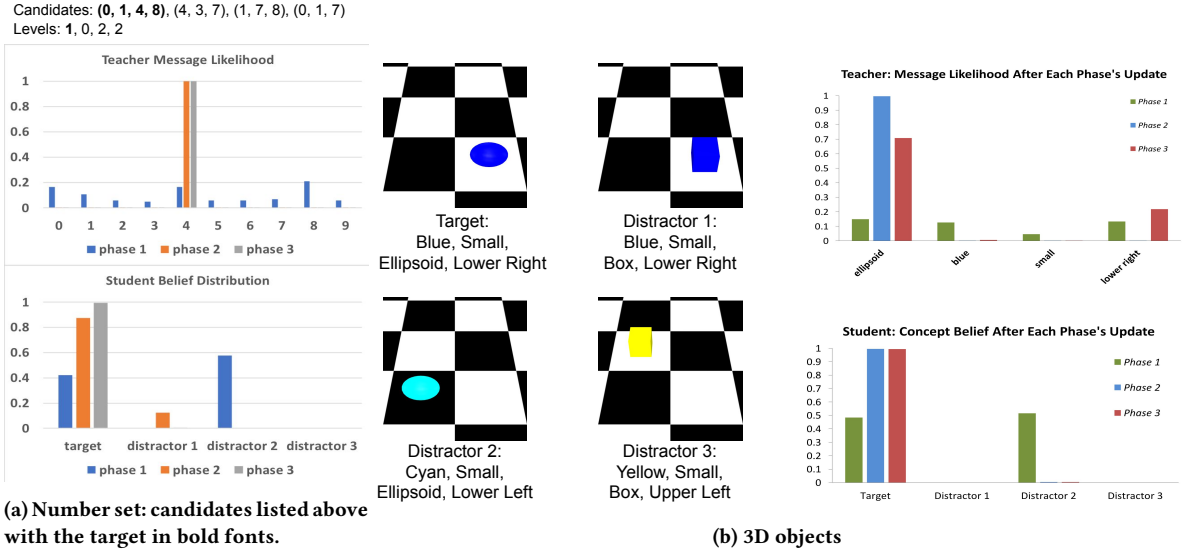


Figure 4: 4 distractors referential game examples. Due to the space limit, we only show the message distribution for the target and student’s new belief after receiving the most probable message. As for the teacher’s message distribution for distractors, all probability weights concentrate on the unique identifiers after the first phase of training. Student’s belief illustrates that teacher’s most probable message, though consistent with multiple candidates, can successfully indicate the target with more confidence as training goes. In general, both agents’ behavior becomes more certain, and the certainty coordinates.

Table 1: The first three rows are the total accuracy, while the second three rows measure the accuracy for difficult games (numbers in percentage). We define game difficulty with average cosine similarity between the target and the distractors. The larger the cosine similarity, the harder the game is. We report the accuracy for top 10% hard games. L.A. stands for Lazaridou et al. To verify the generalization of our algorithm, we form training sets with 70% of all instances and test using the rest 30% unseen instances. The performance doesn’t exhibit noticeable degradation: $99.1 \pm 0.1\%$ and $93.0 \pm 1.7\%$ for 4 and 7 candidates number set respectively (comparable with the No. Set results in the third line). Mean and std calculated using 3 different random splits, 2 experiments per split.

No. Set	3D objects	$\leftarrow 4 \mid 7 \rightarrow$	No. Set	3D Objects
79.1 ± 3.3	86.9 ± 4.1	L. A. [2018]	64.2 ± 6.0	77.3 ± 2.8
96.8 ± 0.2	97.0 ± 0.5	L. A. [2017]	80.8 ± 3.1	88.2 ± 1.7
98.9 ± 0.1	99.6 ± 0.2	Pragmatics	93.2 ± 1.0	97.4 ± 1.2
79.3 ± 3.1	86.9 ± 4.5	L. A. [2018]	67.2 ± 5.8	77.2 ± 2.6
91.5 ± 0.4	88.0 ± 1.9	L. A. [2017]	66.2 ± 3.1	68.2 ± 2.7
98.1 ± 0.3	98.8 ± 0.3	Pragmatics	88.3 ± 0.6	94.1 ± 2.3

encoded by multi-hot encoding. There are 385 and 1585 different possible number sets, consisting up to 9.0×10^9 and 4.9×10^{18} different games with four and seven candidates. Number sets make a generic referential game prototype, where each instance can be disentangled into independent attributes perfectly. To verify the generality of our algorithm on more complicated candidates, we used MoJoCo physical engine to synthesize RGB images of resolution 128×128 depicting single 3D object scenes. For each object, we pick one of six colors (blue, red, yellow, green, cyan, magenta), six

shapes (box, sphere, cylinder, pyramid, cone, ellipsoid), two sizes and four locations, resulting in 288 combinations. In every game, candidates are uniformly sampled from the instances space. We use a message space with the same size as the number of attributes appeared in the dataset, i.e., 10 or 12 for number set, and 18 for 3D objects. In every game, we only allow one round of communication with one message. To prevent collusion using trivial position indicator, candidates are presented to the agents in different orders.

5.1 Referential Game with Symbolic and Pixel Input

For number set, we encode candidates with multi-hot vectors and messages with one-hot vectors. For 3D objects, we used a convolutional neural network (CNN) to extract features of the candidates. For all datasets, we manually generated 600k and 100k mutually exclusive games for training and testing. Namely, the same instances can appear in both datasets, but not any identical candidates combinations. To win a game, the same instance needed to be handled differently given different contexts, so, as long as the games are exclusive between training and testing sets, sharing instances won’t cause over-fitting. To test the robustness, we also report results using exclusive testing instances in table 1. We compared the pragmatic protocol developed using our algorithm against previous works on referential game [37, 38]. Both utilized RL to train a protocol, but neither modeled recursive mind reasoning between agents. In [38], only the teacher considers the context, while no context is included in [37]. We trained our model for 3 phases, with 20k iterations for each phase and switch the training agent in the middle of every phase. Both benchmarks were trained for 100k iterations. Since there is no official code released by the authors, we implemented their model by ourselves and did thorough hyper-parameter grid search. Results shown in table 1. Our experimental results in all

settings are significantly better than both of the benchmarks. Using the paired T-test, the one-tail p -value is smaller than 0.001 for all settings in table 1. We found that even with simpler representation, number set games are more difficult than 3D objects, because we don't have any limitations generating the instances in number sets. 3D objects, on the other hand, form special cases of number sets, as some attributes can never coexist. E.g. a shape cannot be a sphere and a cone simultaneously.

5.2 Emergence of Pragmatics and Connection with RTD

The pragmatic goal of our algorithm is related to the definition of RTD, which measures the number of examples needed for concept learning between a pair of cooperative and rational agents [8]. We included the formal definition of RTD in section 2.2. Intuitively, in a concept class, there is a subset of concepts which are the simplest to learn i.e. has the minimum sized teaching set among all concepts. One can first learn those concepts and remove them from the concept class. Now, for the remaining set of concepts, one can recursively learn the simplest concepts and so on. The teaching complexity of this learning schema lower bounds classic teaching dimension [16]. In every phase of our iterative training, the agent learns to identify the optimal teaching set for the "simplest" remaining candidates. In our experiments, candidates identifiable with a unique message are the simplest. If a candidate becomes the simplest after k times of removal, then we call it a level k candidate.¹ To better illustrate the connection, we show two example referential games in figure 4 and the accuracy improvement after each phase of training in figure 5. We can see from figure 5 that after one phase of training all level 0 targets can be perfectly identified. Thus, the student, if shown the four 3D objects in figure 4b, will know that the teacher will send "Box" for distractor 1 (also a level 1 candidate), "Cyan" for distractor 2 and "Yellow" for distractor 3. Hence, "Ellipsoid", though consistent with two objects, must indicate the target.² The accuracy for higher-level targets in figure 5 keeps increasing as they become uniquely identifiable after lower-level targets are pruned out. We can observe the emergence of pragmatics from these results. From the student's perspective, the messages from the teacher are no longer merely comprehended by their literal meanings, and from the teacher's perspective, she selects the most helpful message to teach. In the 4-candidate scenario, most questions with level 0 and 1 are correctly answered, a similar capability shown in human one-shot referential game study [3].

Notice that RTD is derived under the assumption that both agents only decipher the candidates as sets of discrete attributes. To eliminate the usage of other hidden patterns in the candidates, we need to pretrain the agents to ground messages to instances attributes. This can be easily achieved by initializing agents' belief update function as Bayesian belief update. That is, before running algorithm 1, we train the belief neural network f with cross-entropy

¹There is no inference necessary for level 0 candidates as they can be uniquely identified between literal agents. On the contrary, in figure 4b, the message "ellipsoid" could literally refer to both the target and the distractor 2. A pragmatic inference is required to conclude that the message refers to the target.

²Notice the change of the student belief after hearing "Ellipsoid" after different phases of training, in figure 4b.




Figure 5: Finishing one phase of adaptation, agents can correctly identify level 0 targets. Namely, the teacher always sends a unique identifier as long as there is one. After phase 2, the student knows how the teacher will teach level 0 targets. Thus, he can prune out all level 0 candidates if he doesn't receive their unique identifiers, leaving level 1 candidates the "simplest". The decrease of level 2 accuracy among 4 candidates might be caused by the lack of level 3 targets. Only 3% of the number set and 0.02% 3D Objects games have level 2 targets, so agents have no motivation to explore for a higher level of reasoning.

loss between the generated new belief and ground truth Bayesian belief. Afterward, every message grounds to an attribute.

Bayesian pretraining provides human decipherable examples for failure cases. Most of our model's mistakes are on targets with $RTD > 1$ or $RTD = 1$ but requiring high (≥ 2) level unique identifiers (hard games). Since we only allow one message, targets with $RTD > 1$ are theoretically impossible to be certainly identified, even between cooperative agents with ToM. As for the hard games, their relatively low frequency in the training set may impinge the acquiring of high-level best response. Failure to handle a certain type of scenarios is a common empirical defect of the current RL algorithms, in our case, the hard games. Another benefit of Bayesian pretrain is that the initial message grounding is decipherable to human. In the next section, we show this interpretability can be preserved by our algorithm.

5.3 Stability of the Protocol

We also explored how much the communication changes after the mutual adaptation process. Suppose the agents are initialized with human-understandable message groundings, ideally, we want the emerged protocol preserves its human interpretability. To test this property, we give the teacher and student different but equivalent candidates. Namely, we randomly generate a one-to-one mapping from attributes to attributes, such as replacing all 1 with 7 in number sets or replacing all red with blue in 3D objects. After the teacher sends a message to the student, both the message and candidates are converted with the same mapping before presented to the student. The converted candidates form an equivalent game for the original one. For example, if the teacher sees [(0, 1, 5, 8), (0, 4, 6), (4, 8, 9), (0, 1, 4, 5)] and sends 4. Converted by a mapping which adds 1 to all attributes, the student gets [(1, 2, 6, 9), (1, 5, 7), (5, 9, 0), (1, 2, 5, 6)] and 5.

Table 2: All models are pretrained with Bayesian belief update. The three rows above are tested with agents seeing different but equivalent candidates, while the three rows below is tested with agents sees identical candidates. The larger the performance difference is, the more the protocol diverges from the initial grounding.

No. Set	3D Objects	← 4 7 →	No. Set	3D objects
59.5 ± 0.4	53.3 ± 1.8	L. A. [2018]	35.0 ± 0.2	37.7 ± 1.8
52.3 ± 2.9	60.8 ± 2.9	L. A. [2017]	33.3 ± 2.0	46.0 ± 1.4
97.5 ± 0.4	98.7 ± 0.5	Pragmatics	73.9 ± 0.3	84.2 ± 0.1
81.5 ± 2.6	83.7 ± 3.4	L. A. [2018]	74.9 ± 0.5	75.3 ± 2.9
98.6 ± 1.0	98.1 ± 0.2	L. A. [2017]	93.0 ± 0.5	91.2 ± 0.5
98.1 ± 0.3	99.7 ± 0.3	Pragmatics	88.2 ± 0.8	91.2 ± 0.1

Table 3: A message is valid if it corresponds to one of the attributes of a novel target. Std calculated with 3 different random splits.

	Valid %		Accuracy	
	4	7	4	7
Pragmatics	98.8 ± 0.29	98.1 ± 0.36	98.6 ± 0.08	87.9 ± 0.61
L.A. 2017	74.5 ± 5.69	79.1 ± 4.73	97.0 ± 0.02	88.8 ± 0.25

We pretrain agents with Bayesian belief update separately, then train them together without equivalent games but test them with equivalent games. See table 2 for the results. We can see that our algorithm preserves human interpretability the most. The iterative adaptation contributes to stability because, in every phase, one agent is fixed while optimizing the other. Hence, the evolving of the protocol is not arbitrary and will maintain the effective part of the existing protocol while improving the rest. This property can facilitate human-robot communication, as we only need to provide natural language grounding to robots, and they can self-evolve to take the best advantage of this grounding without developing human undecipherable protocols.

5.4 Global Mapping and Local Selection

The core of pragmatics is the consideration of the context while comprehending the language. In this experiment, we want to show that the teacher using a pragmatic protocol can learn global mappings from instances to messages and select messages dynamically according to the context. First, we test global mapping by checking if messages used by the teacher for a target are consistent with this target’s attributes. Then we evaluate message selection through seeing whether the same target yields different messages given different distractors. To make the message grounding easier to understand, we still pretrain agents with Bayesian belief update. We took 80% of the instances to generate training data. In the testing phase, the rest 20% of images are used as targets with 3/6 distractors randomly selected from all images. Results in table 3 justify that the pragmatic protocols achieve the best balance between message validity and referential accuracy. Then we calculated covariance between messages and distractors attributes given the same target but changing distractors. For a target, the covariance is calculated using 100 games. In figure 6, we visualized mean covariance for 58 targets. Since the teacher doesn’t have access to distractors in [37], we only compared our model and [38].

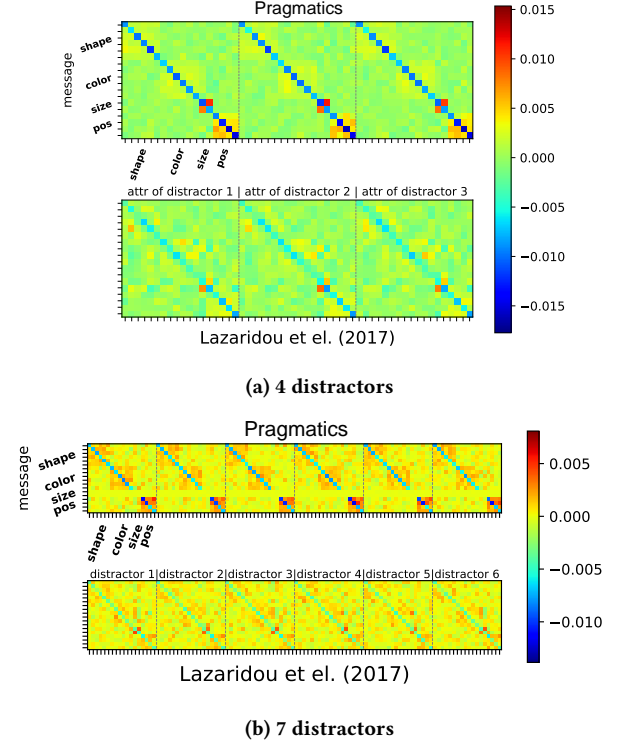


Figure 6: Covariance between messages and distractor attributes given the same target but different distractors. In both experiments, our algorithm has a more clear pattern. The blue diagonals show that if an attribute appears among distractors the teacher tends to avoid the corresponding message. The brighter sub-squares illustrate that messages are significantly influenced by distractors’ attributes in the same category. Namely, color messages have stronger covariance with color attributes than with shapes, sizes or positions. Notice that size messages are seldom used in 7-distractor game, thus their covariance with distractors is close to 0. Better be viewed in high resolution, e.g. using Adobe Acrobat Reader.

6 CONCLUSION

In this paper we propose an end-to-end trainable adaptive training algorithm integrating ToM in multi-agent communication protocol emergence. The pragmatic protocol developed by our algorithm yields significant performance gain over non-pragmatic protocols. With human knowledge pre-grounding, the teaching complexity using the pragmatic protocol approximates RTD. Our algorithm incorporates a global mapping from instances to messages with a local message selection mechanism sensitive to the context. In future research, we hope to generalize the referential game to a new communicative learning framework, where students, instead of learning from data or an oracle, learn from a helpful teacher with ToM. We also plan to apply our algorithm to more generic communication settings, where agents have more symmetric roles. Namely, we have agents each holding some information unknown to the group and need communication to accomplish a task. We also want to relax the need of exchanging beliefs directly in the training phase, replacing it with discrete feedback requiring smaller channel bandwidth.

REFERENCES

- [1] Jacob Andreas and Dan Klein. 2016. Reasoning about pragmatics with neural listeners and speakers. *arXiv preprint arXiv:1604.00562* (2016).
- [2] Chris L Baker, Julian Jara-Ettinger, Rebecca Saxe, and Joshua B Tenenbaum. 2017. Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour* 1, 4 (2017), 0064.
- [3] Leon Bergen, Noah Goodman, and Roger Levy. 2012. That's what she (could have) said: How alternative utterances affect language use. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 34.
- [4] Daniel S Bernstein, Robert Givan, Neil Immerman, and Shlomo Zilberstein. 2002. The complexity of decentralized control of Markov decision processes. *Mathematics of operations research* 27, 4 (2002), 819–840.
- [5] Reinhard Blutner. 2000. Some aspects of optimality in natural language interpretation. *Journal of semantics* 17, 3 (2000), 189–216.
- [6] Ben Bogin, Mor Geva, and Jonathan Berant. 2018. Emergence of communication in an interactive world with consistent speakers. *arXiv preprint arXiv:1809.00549* (2018).
- [7] Kris Cao, Angeliki Lazaridou, Marc Lanctot, Joel Z Leibo, Karl Tuyls, and Stephen Clark. 2018. Emergent communication through negotiation. *arXiv preprint arXiv:1804.03980* (2018).
- [8] Xi Chen, Yu Cheng, and Bo Tang. 2016. On the recursive teaching dimension of vc classes. In *Advances in Neural Information Processing Systems*. 2164–2171.
- [9] Edward Choi, Angeliki Lazaridou, and Nando de Freitas. 2018. Compositional oververber communication learning from raw visual input. *arXiv preprint arXiv:1804.02341* (2018).
- [10] Morten H Christiansen and Simon Kirby. 2003. Language evolution: Consensus and controversies. *Trends in cognitive sciences* 7, 7 (2003).
- [11] Caroline Claus and Craig Boutilier. 1998. The dynamics of reinforcement learning in cooperative multiagent systems. *AAAI/IAAI* 1998 (1998).
- [12] Abhishek Das, Satwik Kottur, José MF Moura, Stefan Lee, and Dhruv Batra. 2017. Learning cooperative visual dialog agents with deep reinforcement learning. In *Proceedings of the IEEE International Conference on Computer Vision*. 2951–2960.
- [13] Harmen de Weerd, Denny Diegmond, and Rineke Verbrugge. 2017. Estimating the use of higher-order theory of mind using computational agents. *The BE Journal of Theoretical Economics* 18, 2 (2017).
- [14] Harmen De Weerd, Rineke Verbrugge, and Bart Verheij. 2014. Theory of Mind in the Mod Game: An Agent-Based Model of Strategic Reasoning. In *ECSI*. 128–136.
- [15] Harmen De Weerd, Rineke Verbrugge, and Bart Verheij. 2015. Higher-order theory of mind in the tacit communication game. *Biologically Inspired Cognitive Architectures* 11 (2015), 10–21.
- [16] Thorsten Doliwa, Gaojian Fan, Hans Ulrich Simon, and Sandra Zilles. 2014. Recursive teaching dimension, VC-dimension and sample compression. *The Journal of Machine Learning Research* 15, 1 (2014), 3107–3131.
- [17] Prashant Doshi and Piotr J Gmytrasiewicz. 2009. Monte Carlo sampling methods for approximating interactive POMDPs. *Journal of Artificial Intelligence Research* 34 (2009), 297–337.
- [18] Katrina Evtimova, Andrew Drozdov, Douwe Kiela, and Kyunghyun Cho. 2017. Emergent communication in a multi-modal, multi-step referential game. *arXiv preprint arXiv:1705.10369* (2017).
- [19] Jaime F Fisac, Monica A Gates, Jessica B Hamrick, Chang Liu, Dylan Hadfield-Menell, Malayandi Palaniappan, Dhruv Malik, S Shankar Sastry, Thomas L Griffiths, and Anca D Dragan. 2017. Pragmatic-pedagogic value alignment. *arXiv preprint arXiv:1707.06354* (2017).
- [20] Jakob Foerster, Ioannis Alexandros Assael, Nando de Freitas, and Shimon Whiteson. 2016. Learning to communicate with deep multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems*.
- [21] Jakob Foerster, Richard Y Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. 2018. Learning with opponent-learning awareness. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 122–130.
- [22] Jakob N Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. 2018. Counterfactual multi-agent policy gradients. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [23] Jakob N Foerster, Francis Song, Edward Hughes, Neil Burch, Iain Dunning, Shimon Whiteson, Matthew Botvinick, and Michael Bowling. 2018. Bayesian action decoder for deep multi-agent reinforcement learning. *arXiv preprint arXiv:1811.01458* (2018).
- [24] Michael Frank, Noah Goodman, Peter Lai, and Joshua Tenenbaum. 2009. Informative communication in word production and word learning. In *Proceedings of the annual meeting of the cognitive science society*, Vol. 31.
- [25] Michael C Frank and Noah D Goodman. 2012. Predicting pragmatic reasoning in language games. *Science* 336, 6084 (2012), 998–998.
- [26] Claudia V Goldman, Martin Allen, and Shlomo Zilberstein. 2007. Learning to communicate in a decentralized environment. *Autonomous agents and multi-agent systems* 15, 1 (2007), 47–90.
- [27] Claudia V Goldman and Shlomo Zilberstein. 2003. Optimizing information exchange in cooperative multi-agent systems. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*. ACM, 137–144.
- [28] Dave Golland, Percy Liang, and Dan Klein. 2010. A game-theoretic approach to generating spatial descriptions. In *Proceedings of the 2010 conference on empirical methods in natural language processing*. Association for Computational Linguistics, 410–419.
- [29] Noah D Goodman and Michael C Frank. 2016. Pragmatic language interpretation as probabilistic inference. *Trends in cognitive sciences* 20, 11 (2016), 818–829.
- [30] H Paul Grice, Peter Cole, Jerry Morgan, et al. 1975. Logic and conversation. 1975 (1975), 41–58.
- [31] Paul Grouchy, Gabriele MT D'Eleuterio, Morten H Christiansen, and Hod Lipson. 2016. On the evolutionary origin of symbolic communication. *Scientific reports* 6 (2016), 34615.
- [32] Yanlin Han and Piotr Gmytrasiewicz. 2018. Learning Others' Intentional Models in Multi-Agent Settings Using Interactive POMDPs. In *Advances in Neural Information Processing Systems*. 5634–5642.
- [33] Rasmus Ibsen-Jensen, Josef Tkadlec, Krishnendu Chatterjee, and Martin A Nowak. 2018. Language acquisition with communication between learners. *Journal of The Royal Society Interface* 15, 140 (2018), 20180073.
- [34] Gerhard Jäger. 2012. Game theory in semantics and pragmatics. *Semantics: An international handbook of natural language meaning* 3 (2012), 2487–2425.
- [35] Fereshte Khani, Noah D Goodman, and Percy Liang. 2018. Planning, Inference and Pragmatics in Sequential Language Games. *Transactions of the Association for Computational Linguistics* 6 (2018), 543–555.
- [36] Michael Kinney and Costas Tsatsoulis. 1998. Learning communication strategies in multiagent systems. *Applied intelligence* 9, 1 (1998), 71–91.
- [37] Angeliki Lazaridou, Karl Moritz Hermann, Karl Tuyls, and Stephen Clark. 2018. Emergence of Linguistic Communication from Referential Games with Symbolic and Pixel Input. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=HJGv1Z-AW>
- [38] Angeliki Lazaridou, Alexander Peysakhovich, and Marco Baroni. 2017. Multi-agent cooperation and the emergence of (natural) language. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=Hk8N3ScIg>
- [39] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems*. 6379–6390.
- [40] Jacob L Mey and Chenghuan Xu. 2001. Pragmatics: an introduction. (2001).
- [41] Jan Pöppel and Stefan Kopp. 2018. Satisficing Models of Bayesian Theory of Mind for Explaining Behavior of Differently Uncertain Agents: Socially Interactive Agents Track. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 470–478.
- [42] David Premack and Guy Woodruff. 1978. Does the chimpanzee have a theory of mind? *Behavioral and brain sciences* 1, 4 (1978), 515–526.
- [43] Neil C Rabinowitz, Frank Perbet, H Francis Song, Chiyuan Zhang, SM Eslami, and Matthew Botvinick. 2018. Machine theory of mind. *arXiv preprint arXiv:1802.07740* (2018).
- [44] Patrick Shafto, Noah D Goodman, and Thomas L Griffiths. 2014. A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive psychology* 71 (2014), 55–89.
- [45] Sainbayar Sukhbaatar, Rob Fergus, et al. 2016. Learning multiagent communication with backpropagation. In *Advances in Neural Information Processing Systems*. 2244–2252.
- [46] Adam Vogel, Max Bodoia, Christopher Potts, and Daniel Jurafsky. 2013. Emergence of Gricean maxims from multi-agent decision theory. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 1072–1081.
- [47] Adam Vogel, Andreas Gomez Emilsson, Michael C Frank, Dan Jurafsky, and Christopher Potts. 2014. Learning to reason pragmatically with cognitive limitations. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 36.
- [48] Christopher John Cornish Hellaby Watkins. 1989. *Learning from delayed rewards*. Ph.D. Dissertation. King's College, Cambridge.
- [49] Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8, 3-4 (1992), 229–256.
- [50] James R Wright and Kevin Leyton-Brown. 2010. Beyond equilibrium: Predicting human behavior in normal-form games. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*.
- [51] Wako Yoshida, Ray J Dolan, and Karl J Friston. 2008. Game theory of mind. *PLoS computational biology* 4, 12 (2008), e1000254.