

---

# Few-shot Language Coordination by Modeling Theory of Mind

---

Hao Zhu<sup>1</sup> Graham Neubig<sup>1</sup> Yonatan Bisk<sup>1</sup>

## Abstract

*No man is an island.* Humans communicate with a large community by coordinating with different interlocutors within short conversations. This ability has been understudied by the research on building neural communicative agents. We study the task of few-shot *language coordination*: agents quickly adapting to their conversational partners’ language abilities. Different from current communicative agents trained with self-play, we require the lead agent to coordinate with a *population* of agents with different linguistic abilities, quickly adapting to communicate with unseen agents in the population. This requires the ability to model the partner’s beliefs, a vital component of human communication. Drawing inspiration from theory-of-mind (ToM; Premack & Woodruff (1978)), we study the effect of the speaker explicitly modeling the listeners’ mental states. The speakers, as shown in our experiments, acquire the ability to predict the reactions of their partner, which helps it generate instructions that concisely express its communicative goal. We examine our hypothesis that the instructions generated with ToM modeling yield better communication performance in both a referential game and a language navigation task. Positive results from our experiments hint at the importance of explicitly modeling communication as a socio-pragmatic progress. Code can be found at <https://github.com/CLAW-Lab/ToM>.

## 1. Introduction

Natural language is an ubiquitous communication medium between human interlocutors, and is shaped by the desire to efficiently cooperate and achieve communicative goals (Gibson et al., 2019). Because of this, there has been interest in creating artificial agents that mimic this communication

process, with a wide variety of works examining communication between agents via either completely artificial *emergent* language (Wagner et al., 2003; Bouchacourt & Baroni, 2018; Li & Bowling, 2019; Kharitonov et al., 2020), or through *natural* language such as English (Lazaridou et al., 2016; Lowe et al., 2019a). In general, these methods model interaction between a pair of agents, a *speaker* and a *listener* that attempt to jointly achieve a goal where the language use is learned to optimize success with respect to a collaborative task (for example, in Fig. 1 the speaker instructs an embodied agent to perform a task in an environment).

However, in contrast to this setup, human speakers interact with not a single listener, but many different conversational partners. In doing so, they also adapt to each other’s language within short conversations. One of the representative phenomena is entrainment, in which interlocutors align their language on both acoustic-prosodic and lexical dimensions during communication (Brennan & Clark, 1996; Levitan et al., 2018). These issues remain mostly unattested in previous work on multi-agent communication — it remains an open question how to train an agent which can adapt to communicating with novel partners quickly.

We define this adaptive communication problem as *few-shot language coordination*. As shown in Fig. 1, within a few rounds of the game, the speaker must adapt its language based on the responses from the listener to achieve the *communicative goal* of helping the listener perform the correct actions as many times as possible. This few-shot coordination setting provides the agents with the pressure to adapt on-the-fly – something that current models generally cannot achieve, even those that model pragmatics, e.g. the rational speech act model (Frank & Goodman, 2012).

Developmental psychology argues for the importance of Theory of Mind (ToM), the understanding of others’ mental states, and also the ability of interlocutors to act upon others mental states to achieve desired effects (Tomasello, 2018). In this paper, we study the importance of modeling the beliefs of one’s conversational partner to achieve few-shot language coordination. In particular, we train a model that quickly adapts to predict the actions of the listeners in real time. At each time step, the speaker predicts the listener’s likely next action given all possible instructions and the listener’s previous actions. This is essentially a few-shot

---

<sup>1</sup>Language Technologies Institute, Carnegie Mellon University. Correspondence to: Hao Zhu <zuhao@cmu.edu>.

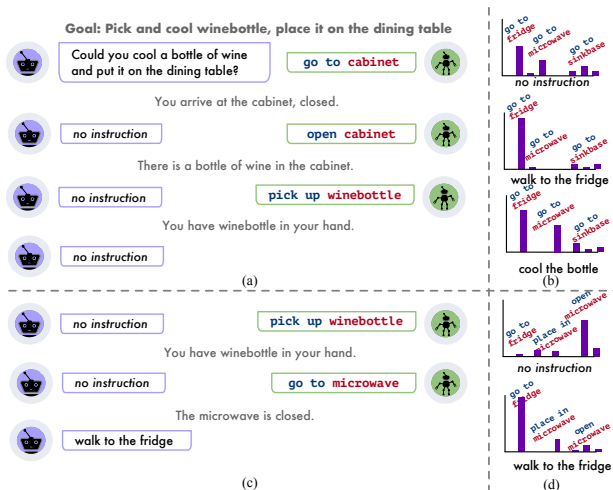


Figure 1. A conversation between a speaker and a listener collaboratively solving a navigation task. (a) At the start of the task, a goal (bold font) is given to the speaker (purple robot head). The speaker first gives a task-level instruction. Without previous knowledge of the listener, the speaker thinks the listener (green robot) will proceed to the fridge after three correct actions (monospace font) in a row. Grey observations are given by the environment after each action. (b) shows the belief of the speaker about the listener’s action after a few instruction candidates. Note that to keep instructions concise the speaker chooses “no instruction” over “walk to fridge” despite the higher probability of listener taking correct action given the latter instruction. (c) After the listener makes a mistake by going to the microwave, the speaker figures out that the listener cannot understand “cool” in the high-level instruction given, and gives low-level instruction “walk to the fridge”. (d) shows the belief of speaker at this time step. Note that the probability of action “go to fridge” without instruction decreases due to the wrong action of the listener.

learning problem, which we attack with model-agnostic meta-learning (MAML) (Finn et al., 2017). In order to achieve the communicative goal, the speaker chooses the best instruction to give so that the probability of the listener performing the correct actions is maximised.

We expect the resulting agent to not only mimic humans’ ability to model a listener’s mental state but also to leverage this estimate to choose better instructions. Through empirical evaluation, we aim to answer the question:

*Can an agent, equipped with a model of theory of mind, quickly adapt to a listener’s language in a few-shot language coordination game?*

Our experiments answer in the affirmative in the both a referential game and a vision-language navigation setting.

## 2. Few-shot Language Coordination

Consider again the example depicted in Fig. 1. A speaker model observes the goal of the current task (“pick and cool wine bottle, place it on the dining table”), and sends messages to a listener model, in an attempt to help it finish the task. The listener model chooses actions to take given the latest and previous instructions. If the listener makes a mistake, such as going to the microwave instead of the refrigerator, the speaker should realize that the listener misunderstood some aspect of the high-level instructions in the first message. This misstep informs the speaker’s model of the listener’s ability, leading them to give lower-level instructions to help correct the trajectory. Through *several games*, the speaker gathers enough data to provide customized messages for individual listeners. This is only possible if a proper prior is provided to the speaker. The simplest prior can be hand-coded rules, e.g. if listener cannot understand abstract instructions, try simpler ones. However, to pursue a more general and powerful communication model (e.g. knowing when to simplify vs. rephrase), we study whether this kind of few-shot language coordination can be *learned* by playing with a *population* of listeners. This section proposes a method to construct meaningful populations.

### 2.1. Asymmetric Speaker-Listener Games

Following previous work on communicative agents (Lazari-dou et al., 2016; Cao et al., 2018; Lowe et al., 2019a), we use goal-oriented language games as the test bed for few-shot language coordination. A general goal-oriented language game provides an environment where the participants uses language to communicate with each other to achieve the given goal. We consider the most basic setting of a simplified two-player shared-goal multi-round setup:

**Environment:** The environment is defined by Observation space, Action space, Goal space and transition function  $E : \mathcal{O} \times \mathcal{A} \rightarrow \mathcal{O} \times \mathcal{G}$ . At the start of each game, the environment provides the speaker with a goal and both participants with observations after each action is taken by the listener. A new game starts after the previous one succeeds or reaches a maximum number of steps.

**Participants:** The participants consist of a speaker and a listener sending and receiving natural language messages. After observing the goal, the speaker gives an instruction to the listener, and the listener performs an action in the environment. If the game is sequential, the speaker can also give an instruction after each action until the game is solved or the maximum number of steps is reached. The speaker is a message-and-action producing model defined by the *vocabulary*  $\Sigma$ ; the space of *observations*  $\mathcal{O}$ ; the space of *actions*  $\mathcal{A}$ ; and a *model*  $f : \mathcal{O} \times \mathcal{G} \rightarrow \Sigma^* \times \mathcal{A}$ . The listener is an

instruction-follower defined by the same vocabulary  $\Sigma$ , observation space  $\mathcal{O}$ , and space of actions  $\mathcal{A}$  as the speaker; and a model  $g : \Sigma^* \times \mathcal{O} \rightarrow \mathcal{A}$ .

**Multi-round Games:** The pair of participants will play a session of  $N$  rounds of games, which are sampled independently. Different from single-round games ( $N=1$ ) used in most previous work (Lazaridou et al., 2016; Cao et al., 2018; Fried et al., 2018; Lowe et al., 2019a), the participants keep the memory of past games in the same session. Multi-round games are not only more general than single-round games, but are essential to few-shot language coordination, because participants have the opportunity to adapt to the interlocutors by learning from feedback during previous rounds.

Note that the listeners in this setting cannot directly observe the goal, so the speakers need to use instructions to inform the listeners about the final or intermediate goals of each game. Within  $N$  rounds, the speaker needs to adapt to the listener’s knowledge to provide the most effective instructions.

## 2.2. Population

Rabinowitz et al. (2018) coined the notion of “machine ToM”, which is a model for tracking agents’ behaviors. To train and evaluate its ability to adapt to different agents, they create populations of subject agents by using different neural architectures and random seeds for parameter initialization.

Our design of populations draws inspiration from their work. However, to make the similarities and differences between agents controllable, we consider a population as a distribution over parameters neural listeners with the same architecture which have been trained on different datasets. A neural listener  $f_\theta^L : \mathcal{O} \times \mathcal{I} \rightarrow \mathcal{A}$  is a mapping from Observations and Instructions to Actions with parameter  $\theta$  of the neural networks. The parameters trained on dataset  $\mathcal{D}$  are:

$$\theta_{\mathcal{D}} = \arg \min_{\theta} \mathcal{L}(f_L, \mathcal{D}) \quad (5)$$

In this way, the variation over listeners is mainly determined by the features of the dataset. By constructing datasets with different feature distributions, we control the listeners’ language abilities.

In the example of Fig. 1, the observations in language games are the items in the visual field, and the agent may perform any number of actions (e.g. “go to cabinet”). The speaker may provide natural language instructions, which can range from high-level (e.g. “cool the winebottle”) to low-level (e.g. “get the wine bottle”, “take it to the refrigerator”, “put it in the refrigerator”). A listener that has never been trained on a particular variety of (usually high-level) instruction would have trouble performing the appropriate

actions. This leads to an exponential population of listeners that are trained on datasets containing, or not containing, particular relevant instructions. Because of this, in order to effectively and concisely communicate, an effective speaker will have to judge the language abilities of its various partners in the population and adjust appropriately; we explain how we do so in the following section.

## 2.3. Theory-of-mind Model

Theory-of-mind, the ability to build a model of one’s conversational partners, is deemed to be crucial in socio-pragmatics theory (Premack & Woodruff, 1978; Tomasello, 2018). Drawing inspiration from this, we build a theory-of-mind model to learn to mimic the behavior of the listener within a short time window.

**Mental State** Modeling mental states is the central concept in building a theory-of-mind. We define the mental state of the listener as the parameters of a neural model, the *ToM model*, that produces the same output for the same inputs as the listener:  $\forall x \in \Sigma^*, o \in \mathcal{O}, g_{\text{ToM}}(x, o; \theta_{\text{mind}}) \approx g(x, o; \theta)$ . It should be noted that in the general case, particularly when different model architectures are used to represent the model itself and the ToM model, the mental state representations may not be unique or even exist. In other words, for any model  $\theta$  there may be more than one parameter setting  $\theta_{\text{mind}}$  that satisfies this condition, or there may be no  $\theta_{\text{mind}}$  that produces the exact same output.

**Building a Theory-of-mind** Learning a ToM model is reduced to inferring the mental state of the listener. For a given listener  $g$  with parameters  $\theta$  and ToM model  $g_{\text{ToM}}$ , we seek a mental state representation  $\theta_{\text{mind}}$ . In practice, we use identical neural architectures for both the listener and ToM Model. However, inferring the exact mental state is infeasible within few interactions. Therefore, we estimate  $g_{\text{ToM}}$  such that

$$\theta_{\text{mind}} = \arg \min_{\theta'} \mathbb{E}_{o,m} \mathcal{L}(g_{\text{ToM}}(o, m; \theta'), g(o, m; \theta)) \quad (6)$$

It is straightforward to apply this definition of mental state in the psychological context for which it was originally proposed. The mental state  $\theta_{\text{mind}}$  is the representation of the listener’s language abilities, which are not directly observable, and which are ultimately used for predicting the belief and behavior of the speaker (Premack & Woodruff, 1978). For example, in our first set of experiments we focus on referential game where the speaker describes the target in order to let the listener pick it out from distractors. We construct a population in which neural listeners with LSTMs and word embeddings have different language comprehension abilities for different languages. One of the possible representations controls the word embeddings in different languages: the

**Training Theory-of-Mind Model for Few-shot Language coordination**

Given

- $N$  training listeners  $L = \{l_i\}_{i=0}^{N-1} \in \mathcal{O} \times \mathcal{I} \rightarrow \mathcal{A}$  ( $i = 0, 1, \dots, N - 1$ ) sampled from  $\mathcal{D}_{\text{listener}}$
- Language game environment  $E: \mathcal{O} \times \mathcal{A} \rightarrow \mathcal{O}$
- Speaker  $S: \mathcal{O} \times \mathcal{G} \rightarrow \mathcal{I}^+ \times \mathcal{A}$
- Message cost function  $C: \mathcal{I} \rightarrow \mathbb{R}$
- Constants cost coefficient  $\kappa \in \mathbb{R}$ , distribution coefficient  $\sigma \in [0, 1]$ , maximum number of interactions  $K \in \mathbb{N}$

While not converged:

1. Define dataset  $\mathcal{D}_{\theta_{\text{mind}}}(l_i) = \{(o_j, m_j, a_j)\}$  for each training listener  $l_i$  and game. For a given game, the goal is  $g$ ; the first observation is  $o_1$ ; the message and action are

$$M, a_j^g = S(o_j, g) \quad (1)$$

$$\mathcal{Q}(M) = \underset{m \in \mathcal{M}}{\text{normalize}}(\mathcal{P}_{\text{ToM}}(a_j^g \mid o_j, m, \{(o_k, m_k, a_k)_{k=1}^{j-1}\}; \theta_{\text{mind}}) \exp(-\kappa C(m))) \quad (2)$$

$$m_j \sim \sigma \mathcal{Q}(M) + (1 - \sigma) \mathcal{U}(M) \quad a_j = l_i(o_j, m_j) \quad o_{j+1} = E(o_j, a_j) \quad (3)$$

where  $a_j^g$  is the planned action of the speaker;  $\underset{m \in \mathcal{M}}{\text{normalize}}$  represents normalizing unnormalized probabilities.

2. Compute prediction loss

$$\mathcal{L}^{\text{pred}}(\mathcal{D}_{\theta_{\text{mind}}}) = -\mathbb{E}_{i \sim \mathcal{U}([N]), k \sim \mathcal{U}([K]), \mathcal{D}_{\text{supp}} \sim \mathcal{U}(\mathcal{D}_{\theta_{\text{mind}}}^k(l_i)), (o, m, a) \sim \mathcal{U}(\mathcal{D}_{\theta_{\text{mind}}}(l_i))} \log \mathcal{P}_{\text{ToM}}(a \mid o, m, \mathcal{D}_{\text{supp}}; \theta_{\text{mind}}) \quad (4)$$

where  $i$  is the index of the listener,  $k$  is the size of the support set which are uniformly sampled from  $\{0, 1, \dots, N - 1\}$  and  $\{0, 1, \dots, K - 1\}$ , the support set  $\mathcal{D}_{\text{supp}}$  and target sample  $(o, m, a)$  are sampled from  $\mathcal{D}_{\theta_{\text{mind}}}$  uniformly.

3. Update the ToM parameters:  $\theta_{\text{mind}} \leftarrow \arg \min_{\theta} \mathcal{L}^{\text{pred}}(\mathcal{D}_{\theta_{\text{mind}}})$

Procedure 1. General Theory-of-Mind (ToM) model training procedure.

mental state of a good language listener should have more meaningful word embeddings, while the one which cannot understand the language should have more random ones. Given that the speaker can acquire an accurate mental state for the listener, it can be used for predicting the probability of listener choosing the correct image when hearing descriptions in different languages. By choosing the one that yields the correct image with the highest probability, the speaker generates the descriptions which improve the referential game. On the other hand, high quality descriptions help the speaker better narrow down the language abilities of the listener. This is similar to the two-way interrelation between language and ToM in humans (De Villiers, 2007).

Following this direction, we present a dynamic view of ToM by putting the observer inside the conversation, instead of the static view of Rabinowitz et al. (2018), which uses ToM for tracking the behavior of the agent without interfering in the games. Our training procedure is presented in Proc. 1. We aggregate a dataset  $\mathcal{D}_{\theta_{\text{mind}}}$  at each epoch, and update the parameters by optimizing the ToM model on the dataset. To aggregate the dataset for each training listener, we randomly sample from the posteriors of the ToM model and uniform distributions over the candidates, which keeps a certain degree of exploration, modulated by distribution coefficient  $\sigma$  (through the paper, we use  $\sigma = 0.5$ ). In practice, parameters are updated with stochastic gradient descent by sampling listeners and using the history of each listener at each time step as a support set for predicting the next actions of the

listener. Following the literature on speech acts, e.g. Monroe & Potts (2015), we also add exponential cost penalty  $\exp(-\kappa C(m))$  as a prior to penalize long instructions. (We have not explored the space of penalty functions in this paper, but the exponential function is widely used in the pragmatics literature, e.g. (Monroe & Potts, 2015), (Morris & Yurovsky, 2019).) In Fig. 1 (a&b), although “go to fridge” yields the highest probability of gold action, no instruction is given in order to express the goal concisely.

Similarly to the imitation learning algorithm DAgger (Ross et al., 2011), the dataset is collected using expert actions. However, there is a major difference between Proc. 1 and DAgger — we optimize the prediction of actions conditioned on the observations and instructions instead of the instruction probability directly. The following theorem shows that our model will improve the instruction generation quality:

**Theorem 1** (informal). *Given a small enough distribution coefficient  $\sigma$  and good enough bounded candidate pools, the instruction distribution produced by the ToM model becomes optimal as prediction loss goes to zero.*

**Discussion** The conditions of Theorem 1 mean that the speaker model  $S$  must be a well-trained model to produce good enough candidates pools. In practice, this condition is not hard to meet: for instance, in our language navigation experiment the listeners can at least understand the lowest-level instructions, and the speaker generates four

levels of instructions by rule-based experts. Therefore, the practical implication of this theorem is helpful – our method reduces to DAgger without expert instructions. Different from DAgger, our training method doesn’t directly optimize the instruction distribution against expert’s instructions, but optimizes the action prediction loss instead, which upper-bounds the instruction loss.

## 2.4. Meta-learning ToM Model

To acquire an estimate of the mental state from very few interactions, the ToM model needs to quickly learn from a small support set. In theory, any model for parameterizing  $\mathcal{P}_{\text{ToM}}(a \mid o, m, \mathcal{D}_{\text{supp}}; \theta_{\text{mind}})$  could work in our framework. As a general method applicable to all ToM models, we apply model-agnostic meta-learning (MAML; Finn et al. (2017)), a method that explicitly trains models to be easily adaptable in few-shot settings. Given support dataset  $\mathcal{D}_{\text{supp}}$ , the inner loop updates parameters for  $N_{\text{inner}}$  steps:

$$\begin{aligned} \theta^{(0)} &= \theta_{\text{mind}} \\ \theta^{(i+1)} &= \theta^{(i)} - \eta \nabla_{\theta} \mathcal{L}(\mathcal{D}_{\text{supp}}; \theta) \Big|_{\theta=\theta^{(i)}} \end{aligned} \quad (7)$$

where  $\eta$  is the inner loop learning rate, which, in practice, is not share across modules following Antoniou et al. (2019). The loss function on the support set is the negative log-likelihood of listener’s action given observations  $o$  and the instructions  $m$

$$\mathcal{L}(\mathcal{D}_{\text{supp}}; \theta) = -\mathbb{E}_{(o,m,a) \sim \mathcal{D}_{\text{supp}}} \log p_{\theta}(a \mid o, m) \quad (8)$$

After  $N_{\text{inner}}$  steps, we get the prediction on the target observation  $o$  and instruction  $m$

$$\mathcal{P}_{\text{ToM}}(a \mid o, m, \mathcal{D}_{\text{supp}}; \theta_{\text{mind}}) = p_{\theta^{N_{\text{inner}}}}(a \mid o, m) \quad (9)$$

Outer loop optimize  $\theta_{\text{mind}}$  by mini-batch stochastic gradient descent

$$\theta_{\text{mind}} \leftarrow \theta_{\text{mind}} - \eta_{\text{outer}} \nabla_{\theta} \mathcal{L}^{\text{pred}}(\mathcal{D}_{\theta_{\text{mind}}}) \quad (10)$$

The outer loop also runs for a given  $N_{\text{outer}}$  epochs.

## 2.5. Deploying the ToM Model

Similar to Proc. 1, we evaluate ToM by using it to measure the probability of the gold action given the instructions. However, here we choose the best one instead of sampling from the posterior. Alg. 1 shows the evaluation procedure.

## 2.6. Connection with Other Pragmatics Models

It should be noted that using a listener model to help choose best utterance has been studied for almost a decade under the rational speech act model (RSA, Frank & Goodman

---

### Algorithm 1 Evaluate ToM Model

---

**Require:** Testing Listeners  $L_{\text{test}}, E, S, C, \kappa, K$  as in Proc. 1  
 $pt \leftarrow 0$   
**for all**  $l_i$  in  $L_{\text{test}}$  **do**  
 $\mathcal{D} \leftarrow \emptyset$   
 $o, a \leftarrow \text{RESTART}$   
**for j** in  $1..K$  **do**  
 $o, g \leftarrow E(o, a)$   
**if DONE then**  
 $pt \leftarrow pt + 1[\text{SUCCESS}]$   
 $o, g \leftarrow E(\text{RESTART})$   
**end if**  
 $M, a^g \leftarrow S(o, g)$   
 $m \leftarrow \arg \max_m \mathcal{P}_{\text{ToM}}(a^g \mid o, m, \mathcal{D}); \theta_{\text{mind}} e^{-\kappa C(m)}$   
 $a \leftarrow l_i(o, m)$   
 $\mathcal{D} \leftarrow \mathcal{D} \cup \{(o, m, a)\}$   
**end for**  
**end for**  
**Return**  $pt$

---

(2012); including recent more general models, e.g. Wang et al. (2020)), a Bayesian framework that takes listener’s choices in to account by

$$\begin{aligned} P_{S^n}(m \mid a, o) &= \frac{P_{L^{n-1}}(a \mid m, o) P(m \mid o)}{\sum_{m' \in M} P_{L^{n-1}}(a \mid m', o) P(m' \mid o)} \\ P_{L^n}(a \mid m, o) &= \frac{P_{S^{n-1}}(m \mid a, o) P(a \mid o)}{\sum_{a' \in A} P_{L^{n-1}}(m \mid a', o) P(a' \mid o)} \end{aligned} \quad (11)$$

where  $S^n$  denotes the  $n$ -level speaker and  $L^{n-1}$  denotes the  $(n-1)$ -level listener,  $M, A, o$  denotes the space of instructions, actions, and the observation shared by the speaker and listener respectively,  $P(m \mid o)$  and  $P(a \mid o)$  are the priors over instructions and actions. The base speaker  $S^0$  and listener  $L^0$  are often parameterized using neural networks directly (Fried et al., 2018).

As a general framework for computational pragmatics, RSA models both language production and language comprehension in a recursive fashion, although the first and the second levels are predominantly used. In this paper, we focus on language production, while improving the listeners with more layers of reasoning is left for future work.

However, the most notable difference between our model and neural RSAs is the notion of few-shot coordination. RSA base speaker and listener models are often fixed after training, making them unable to adapt to new partners during testing. While our model has a similar formulation (Eq. 2) to the first level speaker of RSA, our ToM listener’s action probability conditions on the listener’s previous behavior.

## 3. Multilingual Referential Games

We test the ability of the proposed ToM model to perform few-shot language coordination in two settings: the running example of vision-language navigation, and also in a simpler

setting of *referential games*, which we discuss first in this section. In a referential game, the speaker gives a description for the target image as its instruction, and the listener’s action is to choose the target from distractors, after which the listener either wins the game and gets one point or loses it.

Following Lazaridou et al. (2016); Lowe et al. (2019a), we use 30k image-caption pairs from MSCOCO dataset (Lin et al., 2014). In each game, a target image sampled from the dataset uniformly, and nine distractors are sampled from 1,000 nearest images in terms of cosine similarity of outputs of second last layer of pretrained ResNet (He et al., 2016). In contrast to previous work, which mainly deals with a pair of one speaker and one listener, we are interested in learning with a population of listeners. In order to achieve this, we propose a setting of *multilingual* referential games, where each listener has the ability to understand different languages at different levels of ability.

**Listener distribution** We first translate MSCOCO captions into nine languages, German, Lithuanian, Chinese, Italian, French, Portuguese, Spanish, Japanese and Greek, from English, using Google Translate<sup>1</sup>. For each listener, we sample a vocabulary distribution  $v_1, v_2, \dots, v_{10}$  from 10-dimensional Dirichlet distribution  $Dir(0.5, 0.5, \dots, 0.5)$ . The listener’s vocabulary is built up with 5,000 words, where for each language  $i$  we select the most frequent  $5,000 * v_i$  words in MSCOCO captions in that language to be added to the listener’s vocabulary. The reason behind this design is cognitively motivated; word frequency has high correlation with age of acquisition (AoA) of words (Juhasz, 2005). The dataset used to train the listener is finally created by filtering out sentences with more than one word outside the vocabulary. Given target image  $x^*$ , instruction  $m$ , and distractors  $x_i, i = 1, 2, \dots, 9$ , the listener computes

$$\begin{aligned} z_i &= \text{ResNet}(x_i) \quad \text{for } i = 1, 2, \dots, 9, * \\ z &= \text{LSTM}(m) \\ \hat{y} &= \text{softmax}(z^\top \{z_1, z_2, \dots, z_9, z^*\}) \end{aligned} \tag{12}$$

The listener is trained to minimize the expected negative log-likelihood  $-\log \hat{y}^*$  by stochastic gradient descent.

Following Lowe et al. (2019a), we train the listeners by randomly<sup>2</sup> interleaving between self-play (training with a companion speaker) and supervised training (with MSCOCO annotations or their translations). The companion speaker

<sup>1</sup><https://translate.google.com>

<sup>2</sup>We have also tried other schemes in their paper, but those do not yield significantly better performance.

takes the representation of the target image as input:

$$\begin{aligned} z^* &= \text{ResNet}(x^*) \\ l &= \text{teacher-forcing}(\text{LSTM}(z^*), m) \\ \hat{m} &= \text{gumbel-softmax}(\text{LSTM}(z^*)) \end{aligned} \tag{13}$$

During supervised training, the model is trained to minimize the teacher-forcing NLL loss, while during self-play the sampled instruction is fed to the listener with Gumbel-softmax (Jang et al., 2017). This procedure produces 120 listeners, for which the average success rate with MSCOCO captions within the listener’s vocabulary is 81.6% and the average success rate with companion speakers is 83.3%. These listeners are randomly divided into training, validation, and testing listeners (80/20/20).

**Speaker training** Using the setup in Eqs. 12 and 13, we equip the speaker with a vocabulary of 20K words equally distributed in ten languages. We use the same data filtering method and training scheme as described above. To produce a pool of candidates in all languages, we add a language marker at the front of each training caption, so that the languages of instructions are controllable. Using beam search (size of 10), we generate five instructions per language (i.e.  $N_M = 50$ ). The speaker achieves an 87% success rate with the listeners used to train the speaker and a caption PPL of 23.7.

**ToM Model** The ToM models uses the same architecture as Eq. 12. We use penalty  $\kappa = 0$ . In the referential game, the action space  $\mathcal{A} = \{1, 2, \dots, 9, *\}$  and observation  $o = (x_1, x_2, \dots, x_9, x^*)$ , we have

$$p_\theta(a | o, m) = \hat{y}_a. \tag{14}$$

The MAML hyper-parameters are  $\eta = 0.01, N_{\text{inner}} = 5, \eta_{\text{outer}} = 0.0001, N_{\text{outer}} = 500$ , and batch size is 2.

**Evaluation** We evaluate the ToM-assisted speaker and other baselines with the same set of testing listeners. For each pair of speaker and listener, we calculate the average success rate of 500  $K = 20$ -game sessions.

Model	Ave success
Gold-standard speaker	91.20%
Non-ToM speaker	37.38%
RSA speaker	42.83%
ToM-assisted speaker	<b>58.19%</b>

Table 1. Models and their respective referential game accuracy.

The gold-standard speaker denotes the success rate of using the testing listener in place of the ToM listener. The score




Figure 2. Average prediction accuracy of ToM model at each time step during evaluation. (95% confidence interval)

of over 90% indicates that the candidate pool is of high quality, so a speaker with a well-modeled ToM listener has ample room for achieving high accuracy. The non-ToM speaker uses the instruction with the highest probability in the speaker model; the RSA speaker uses the listener for training the speaker in place of the ToM listener. Our model achieves a significantly higher success rate, demonstrating that the ToM model could help produce better instructions for this referential game.

However, does ToM model truly learn to adapt to individual listeners? We compute the accuracy of predicting the listener’s behavior during the same session. Fig. 2 shows that the prediction accuracy of listener’s actions is significantly improved within sessions, which shows ToM indeed learns to adapt to individual test listeners.

#### 4. ALFWorld Task Navigation

In the previous section, we have shown that ToM could help games with simple dynamics, and learn to adapt to listeners. This section will show its application in a more complex game: language navigation.

We use the Alfworld (Shridhar et al., 2021) platform, which creates a natural language command action space upon the Alfred environment for vision-language navigation (Shridhar et al., 2020). In each game, a household task is given by the environment. We generate expert trajectories for all tasks in Alfworld, and manually create four levels of instructions, from task-level to action-level, which are denoted as  $\mathcal{I}_i, i = 1, 2, 3, 4$ . The differences between these four levels are the levels of abstraction. An action-level instruction corresponds to a single action in the Alfworld, while a task-level instruction corresponds to a whole trajectory which consists of more than eight commands. The two other levels are in between. The candidate pool at each time step consists of four instructions from each level ( $N_M = 4$ ). To create each listener, we draw an instruction distribution from  $Dir(0.6, 0.4, 0.3, 0.2)$  for each of the six types of task in the environment. Listeners are of the same neural model

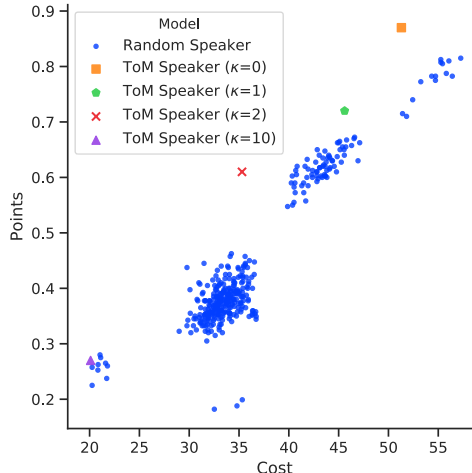


Figure 3. Experimental results for the language navigation setting with average instruction length on the horizontal-axis and game points on the vertical. Colors represent different models.

as in Shridhar et al. (2021). While training the listeners, the instructions are randomly drawn from the instruction distributions according to the task type. This procedure produces 50 listeners. The average success rate of listeners is 83.6%. These listeners are randomly divided into training, validation, and testing listeners (30/10/10). We define the cost as the average total length of the set of instructions, i.e. repetitive instructions are only calculated once. The cost function is defined as  $C(m) = 2^i$  if  $m \in \mathcal{I}_i$ . Within one session, the maximum number of interactions between speaker and listener is  $K = 100$ , and maximum number of interactions in a game is 20. Listeners’ hyper-parameters are the same as the ones in Shridhar et al. (2021), while MAML hyper-parameters are the same as referential game.

In Fig. 3, we compare ToM-assisted speakers and random speakers. We didn’t compare with an RSA speaker, because differently from the multilingual referential games, the speaker is rule-based and no listener is used for training the speaker. A random speaker draws an instruction distribution from  $Dir(0.7, 0.7, 0.7, 0.7)$ , and sample instructions from the candidate pools using the instruction distribution. The ToM speaker with  $\kappa = 0$  predominantly uses action-level instructions, while the ToM speaker with  $\kappa = 10$  uses task-level instructions most of the time. Comparing ToM speakers and random speakers, we find that for  $\kappa = \{0, 1, 2\}$ , ToM speakers achieve higher game points and lower cost than random ones; for  $\kappa = 10$ , the ToM speaker does not have significant improvement over a random one, since only the listeners that have been trained on sufficient action-level instructions can succeed.

## 5. Related Work

### 5.1. Language Games

Language games have been the proving ground for various linguistics theories since their conception by Wittgenstein (1953). Recently, the most widely used language game is the referential game, in which the speaker observes the target and distractors and uses language to instruct the listener on how to pick out the target.

**Emergent Communication** Without natural language annotations, this pressure for the speaker and listener enables language emergence. Batali (1998) first uses the same recurrent neural networks as the speaker and the listener to conduct emergent communication in referential game. Following this lead, Lazaridou et al. (2016) study how emergent languages are grounded to the input images. Cao et al. (2018) studies multi-turn communication via negotiation. Chaabouni et al. (2020); Gupta et al. (2020) study the compositionally and systematicity of emergent languages.

**Learning Common Languages** By using natural language annotations, agents learn a common language so that agents that are never trained together can be expected to communicate. Lazaridou et al. (2016) studies using MSCOCO (Lin et al., 2014) annotations as gold labels for both speakers and listeners. Lowe et al. (2019a) found that alternating between self-playing and supervised learning benefits communication performance. Wang et al. (2016) show that humans have the ability to adapt to the machine’s capability in language games. Bullard et al. (2020) found that when language follows Zipf law, zero-shot communication is possible.

**Language Games with Community** Tieleman et al. (2019) learns representations by training with a community of encoders and decoders. The difference between our work and theirs is that our MAML listener learns to adapt to different listeners in the population in a few games. The performance of their model should be equivalent to our model’s result at time step 1. Lowe et al. (2019b) considers the adaptation problem, which is definitely relevant. However, adapting their model to our settings is non-trivial, which is beyond the scope of this paper.

This referential game setting used in most previous work can be seen as a special case of our few-shot coordination formulation, where the number of games is one and the partners are mostly the same ones as in the training phase. These two differences prevent the previous models from learning to adapt due to the lack of pressure to do so.

### 5.2. Machine Theory of Mind

Computational and neural models of theory-of-mind have been studied for decades (Siegal & Varley, 2002; Rescorla, 2015). Rabinowitz et al. (2018) are the first to present a

successful modeling of the mental state of various species of agent. While we also train a ToM model with meta learning, we put the ToM model into use. The predictions provided by the ToM model serve as reranker in the speaker’s model. The variety of models is also more diverse than the species used in this paper. Importantly, Nematzadeh et al. (2018); Le et al. (2019) find that neural models for question answering fail to keep track of inconsistent states of the world. Moreno et al. (2021) extends machine ToM to neural recursive belief states. We expect improvement over our current model by modeling higher-order recursive belief, which is left for future work. Yuan et al. (2020) explicitly trains belief state prediction with supervised learning, while our model’s belief state is latent.

### 5.3. Similar Topics in Reinforcement Learning

**Model-based Reinforcement Learning** Model-based reinforcement learning focuses on building a model of the environment to improve data efficiency (Kaelbling et al., 1996), which could be applied to the zero-shot coordination problem by treating the listener as a part of the environment. Recently, neural networks have been used widely for model-based RL (Gal et al., 2016; Depeweg et al., 2016; Nagabandi et al., 2018; Chua et al., 2018; Janner et al., 2019). We should point out that despite their similarities to our model, we focus on modeling different and unseen agents in the population within a few interactions.

**Alternatives to Self-play** Zero-shot coordination has attracted much attention recently. Hu et al. (2020) propose other-play which maximizes the expected reward working with random partners. In contrast to their approach, we explicitly model the partner’s ToM and focus on language coordination, which required more complicated modeling than the environments in their experiments.

## 6. Implications and Future Work

We have introduced few-shot language coordination task and proposed ToM model for tracking the listener’s mental state. Different from previous work using single-round games and self-play training, we consider more general multi-round games and playing with novel listeners. ToM model shows its ability to adapt to novel listeners and assist speakers in choosing the best instructions in both multilingual referential games and the language navigation task. We attribute the success of ToM model to modeling socio-pragmatics process in an explicit way. Many interesting questions about modeling ToM in few-shot language coordination remain open. The most immediate is how to model ToM from the listener’s perspective, leading to a dialog agent that can acquire new knowledge through conversation. One step further, similar to RSA, ToM can also be modeled in a recursive manner, which may further improve language games.



## Acknowledgements

This work was supported by the DARPA GAILA project (award HR00111990063). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

## References

- Antoniou, A., Edwards, H., and Storkey, A. How to train your MAML. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=HJGven05Y7>.
- Batali, J. Computational simulations of the emergence of grammar. *Approach to the Evolution of Language*, pp. 405–426, 1998.
- Bouchacourt, D. and Baroni, M. How agents see things: On visual representations in an emergent language game. *arXiv preprint arXiv:1808.10696*, 2018.
- Brennan, S. E. and Clark, H. H. Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(6): 1482, 1996.
- Bullard, K., Meier, F., Kiela, D., Pineau, J., and Foerster, J. Exploring zero-shot emergent communication in embodied multi-agent populations. *arXiv preprint arXiv:2010.15896*, 2020.
- Cao, K., Lazaridou, A., Lanctot, M., Leibo, J. Z., Tuyls, K., and Clark, S. Emergent communication through negotiation. In *International Conference on Learning Representations*, 2018.
- Chaabouni, R., Kharitonov, E., Bouchacourt, D., Dupoux, E., and Baroni, M. Compositionality and generalization in emergent languages. *arXiv preprint arXiv:2004.09124*, 2020.
- Chua, K., Calandra, R., McAllister, R., and Levine, S. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*, pp. 4754–4765, 2018.
- De Villiers, J. The interface of language and theory of mind. *Lingua*, 117(11):1858–1878, 2007.
- Depeweg, S., Hernández-Lobato, J. M., Doshi-Velez, F., and Udluft, S. Learning and policy search in stochastic dynamical systems with bayesian neural networks. *arXiv preprint arXiv:1605.07127*, 2016.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pp. 1126–1135. PMLR, 2017.
- Frank, M. C. and Goodman, N. D. Predicting pragmatic reasoning in language games. *Science*, 336(6084):998–998, 2012.
- Fried, D., Andreas, J., and Klein, D. Unified pragmatic models for generating and following instructions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 1951–1963, 2018.
- Gal, Y., McAllister, R., and Rasmussen, C. E. Improving pilco with bayesian neural network dynamics models. In *Data-Efficient Machine Learning workshop, ICML*, volume 4, pp. 34, 2016.
- Gibson, E., Futrell, R., Piantadosi, S. P., Dautriche, I., Mahowald, K., Bergen, L., and Levy, R. How efficiency shapes human language. *Trends in cognitive sciences*, 23(5):389–407, 2019.
- Gupta, A., Resnick, C., Foerster, J., Dai, A., and Cho, K. Compositionality and capacity in emergent languages. In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pp. 34–38, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.repl4nlp-1.5. URL <https://www.aclweb.org/anthology/2020.repl4nlp-1.5>.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hu, H., Lerer, A., Peysakhovich, A., and Foerster, J. “Other-Play” for Zero-Shot Coordination. In *International Conference on Machine Learning*, 2020.
- Jang, E., Gu, S., and Poole, B. Categorical Reparameterization with Gumbel-Softmax. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017. URL <https://openreview.net/forum?id=rkE3y85ee>.
- Janner, M., Fu, J., Zhang, M., and Levine, S. When to trust your model: Model-based policy optimization. In *Advances in Neural Information Processing Systems*, pp. 12519–12530, 2019.
- Juhasz, B. J. Age-of-acquisition effects in word and picture identification. *Psychological bulletin*, 131(5):684, 2005.

- Kaelbling, L. P., Littman, M. L., and Moore, A. W. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4:237–285, 1996.
- Kharitonov, E., Chaabouni, R., Bouchacourt, D., and Baroni, M. Entropy minimization in emergent languages. In *International Conference on Machine Learning*, pp. 5220–5230. PMLR, 2020.
- Lazaridou, A., Peysakhovich, A., and Baroni, M. Multi-agent cooperation and the emergence of (natural) language. In *International Conference on Learning Representations*, 2016.
- Le, M., Boureau, Y.-L., and Nickel, M. Revisiting the evaluation of theory of mind through question answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 5875–5880, 2019.
- Levitan, S. I., Xiang, J., and Hirschberg, J. Acoustic-prosodic and lexical entrainment in deceptive dialogue. In *Proc. 9th International Conference on Speech Prosody*, pp. 532–536, 2018.
- Li, F. and Bowling, M. Ease-of-teaching and language structure from emergent communication. In *Advances in Neural Information Processing Systems*, pp. 15851–15861, 2019.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft coco: Common objects in context. In *European conference on computer vision*, pp. 740–755. Springer, 2014.
- Lowe, R., Gupta, A., Foerster, J., Kiela, D., and Pineau, J. On the interaction between supervision and self-play in emergent communication. In *International Conference on Learning Representations*, 2019a.
- Lowe, R., Gupta, A., Foerster, J., Kiela, D., and Pineau, J. Learning to learn to communicate. In *Proceedings of the 1st Adaptive & Multitask Learning Workshop*, 2019b.
- Monroe, W. and Potts, C. Learning in the rational speech acts model. *arXiv preprint arXiv:1510.06807*, 2015.
- Moreno, P., Hughes, E., McKee, K. R., Pires, B. A., and Weber, T. Neural recursive belief states in multi-agent reinforcement learning. *arXiv:2102.02274*, 2021.
- Morris, B. and Yurovsky, D. Pressure to communicate across knowledge asymmetries leads to pedagogically supportive language input. In *CogSci*, pp. 2399–2405, 2019.
- Nagabandi, A., Kahn, G., Fearing, R. S., and Levine, S. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 7559–7566. IEEE, 2018.
- Nematzadeh, A., Burns, K., Grant, E., Gopnik, A., and Griffiths, T. Evaluating theory of mind in question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2392–2400, 2018.
- Premack, D. and Woodruff, G. Does the chimpanzee have a theory of mind? *Behavioral and brain sciences*, 1(4): 515–526, 1978.
- Rabinowitz, N., Perbet, F., Song, F., Zhang, C., Eslami, S. A., and Botvinick, M. Machine theory of mind. In *International Conference on Machine Learning*, pp. 4218–4227, 2018.
- Rescorla, M. The computational theory of mind. 2015.
- Ross, S., Gordon, G., and Bagnell, D. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 627–635. JMLR Workshop and Conference Proceedings, 2011.
- Shridhar, M., Thomason, J., Gordon, D., Bisk, Y., Han, W., Mottaghi, R., Zettlemoyer, L., and Fox, D. ALFRED: A Benchmark for Interpreting Grounded Instructions for Everyday Tasks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. URL <https://arxiv.org/abs/1912.01734>.
- Shridhar, M., Yuan, X., Côté, M.-A., Bisk, Y., Trischler, A., and Hausknecht, M. ALFWorld: Aligning Text and Embodied Environments for Interactive Learning. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021. URL <https://arxiv.org/abs/2010.03768>.
- Siegal, M. and Varley, R. Neural systems involved in ‘theory of mind’. *Nature Reviews Neuroscience*, 3(6):463–471, 2002.
- Tieleman, O., Lazaridou, A., Mourad, S., Blundell, C., and Precup, D. Shaping representations through communication: community size effect in artificial learning systems. *arXiv:1912.06208*, 2019.
- Tomasello, M. How children come to understand false beliefs: A shared intentionality account. *Proceedings of the National Academy of Sciences*, 115(34):8491–8498, 2018.

Wagner, K., Reggia, J. A., Uriagereka, J., and Wilkinson, G. S. Progress in the simulation of emergent communication and language. *Adaptive Behavior*, 11(1):37–69, 2003.

Wang, P., Wang, J., Paranamana, P., and Shafto, P. A mathematical theory of cooperative communication. *ArXiv*, abs/1910.02822, 2020.

Wang, S. I., Liang, P., and Manning, C. Learning language games through interaction. In *Association for Computational Linguistics (ACL)*, 2016.

Wittgenstein, L. *Philosophical Investigations*. 1953.

Yuan, L., Fu, Z., Shen, J., Xu, L., Shen, J., and Zhu, S.-C. Emergence of pragmatics from referential game between theory of mind agents. *arXiv:2001.07752*, 2020.