
Modeling Others using Oneself in Multi-Agent Reinforcement Learning

Roberta Raileanu¹ Emily Denton¹ Arthur Szlam² Rob Fergus^{1,2}

Abstract

We consider the multi-agent reinforcement learning setting with imperfect information. The reward function depends on the hidden goals of both agents, so the agents must infer the other players’ goals from their observed behavior in order to maximize their returns. We propose a new approach for learning in these domains: Self Other-Modeling (SOM), in which an agent uses its own policy to predict the other agent’s actions and update its belief of their hidden goal in an online manner. We evaluate this approach on three different tasks and show that the agents are able to learn better policies using their estimate of the other players’ goals, in both cooperative and competitive settings.

1. Introduction

Reasoning about other agents’ intentions and being able to predict their behavior is important in multi-agent systems, in which the agents might have different, and sometimes competing, goals. In this paper, we introduce a new approach for estimating other agents’ unknown goals from their behavior and using those estimates to choose actions. We demonstrate that in the proposed tasks, using an explicit model of the other player leads to better performance than simply considering the other agent as part of the environment.

We frame the problem as a two-player stochastic game (Shapley, 1953), in which each agent is randomly assigned a different goal from a fixed set, which is shared between the agents. Players have full visibility of the environment, but no direct knowledge of the other’s goal and no communication channel. The reward obtained by each agent at the end of an episode depends on the goals of both agents, so an optimal policy must take into account both of their goals.

The key idea of this work is that as a first approximation

¹New York University, New York City, USA ²Facebook AI Research, New York City, USA. Correspondence to: Roberta Raileanu <raileanu@cs.nyu.edu>.

of understanding what the other player is trying to achieve, an agent should ask itself “what would be my goal if I had acted as the other player had?”. We instantiate this idea by parametrizing the agent’s action and value functions with a neural network that takes as input the observation state and a goal. As the agent plays the game, it uses its own policy (with the input expressed in the other agent’s frame of reference) to maximize the likelihood of the other’s observed actions and optimize directly over the goal representation to infer the other agent’s unknown goal. In contrast with the current literature, our approach does not require building any model of the other agent in order to infer its intention and predict its behavior.

2. Approach

Background: A two-player Markov game is defined by a set of states \mathcal{S} describing the possible configurations of all agents, a set of actions $\mathcal{A}^1, \mathcal{A}^2$ and a set of observations $\mathcal{S}^1, \mathcal{S}^2$ for each agent, and a transition function $\mathcal{T} : \mathcal{S} \times \mathcal{A}^1 \times \mathcal{A}^2 \rightarrow \mathcal{S}$ which gives the probability distribution on the next state as a function of current state and actions. Each agent i chooses actions by sampling from a stochastic policy $\pi^i : \mathcal{S} \times \mathcal{A}^i \rightarrow [0, 1]$. The reward function of each agent is: $r^i : \mathcal{S} \times \mathcal{A}^1 \times \mathcal{A}^2 \rightarrow \mathbb{R}$. Each agent i aims to maximize its discounted return from time t onward: $R_t^i = \sum_{t=0}^{\infty} \gamma^t r_t^i$, where r_t^i is the reward obtained by agent i at time t and $\gamma \in (0, 1]$ is the discount factor. In this work, we consider both cooperative and adversarial settings. In cooperative games, the agents have the same reward function: $r^1 = r^2$.

We now describe Self Other-Modeling (SOM), a new approach for inferring other agents’ goals in an online fashion and using these estimates to choose actions. To decide an action and to estimate the value of a state, we use a neural network f that takes as input its own goal z_{self} , an estimate of the other player’s goal \tilde{z}_{other} , and the observation state s_{self} , and outputs a probability distribution over actions π and a value estimate V , such that for each agent i playing the game we have:

$$\begin{bmatrix} \pi^i \\ V^i \end{bmatrix} = f(s_{self}^i, z_{self}^i, \tilde{z}_{other}^i; \theta^i).$$

Here θ^i are agent i ’s parameters for f , which has one softmax output for the policy, one linear output for the value function, and all the non-output layers shared. The actions

are sampled from policy π^i . The state s_{self}^i contains the observation features from agent i 's viewpoint.

We propose that each agent models the behavior of the other player using its own policy. Thus, each agent uses its own network f in two ways: acting mode, in which the agent uses f to choose its actions and inference mode, in which the agent uses f to infer the other agent's goal. For notation purposes, whenever f is used in acting mode (inference mode) we will refer to it as f_{self} (f_{other}):

$$\text{acting mode: } f_{self}(s_{self}, z_{self}, \tilde{z}_{other}; \theta) \quad (1)$$

$$\text{inference mode: } f_{other}(s_{other}, \tilde{z}_{other}, z_{self}; \theta). \quad (2)$$

The two modes have different relative placements of the network's inputs z_{self} and \tilde{z}_{other} . Additionally, since the environment is fully observed, the observation state of the two agents differs only by the specification of the agent's identity on the map (i.e. each agent will be able to distinguish between its own location and the other's location). Hence, in acting mode, the network f_{self} will take as input s_{self} (with the identity of the acting agent at the location of the *self*) and in inference mode, the network f_{other} will take as input s_{other} (with the identity of the acting agent at the location of the *other*).

At each step, the agent uses equation (2) to output an estimate of the probability distribution over the other agent's actions. Then, the agent uses supervision of the other's true action to backpropagate through f_{other} (without updating its parameters) and directly optimize over its input \tilde{z}_{other} , the estimate of the other agent's goal. The number of optimization steps used to update \tilde{z}_{other} is a hyperparameter that can vary with the game. The new estimate \tilde{z}_{other} is used as input to f_{self} in (1) for choosing the self agent's next action. Figure 1 illustrates this technique.

Note that the network f is never updated during inference mode (i.e. using supervision of the other agent's actions), f 's parameters θ are updated only at the end of each episode using Asynchronous Advantage Actor-Critic (A3C) (Mnih et al., 2016) with reward signal obtained by the self agent. In contrast, \tilde{z}_{other} is updated (multiple times) at each step in the game.

Algorithm 1 represents the pseudo-code for training a SOM agent for one episode. The procedure is formulated from the viewpoint of a single agent. Since the goals are discrete in all the tasks considered here, the agent's goal z_{self} is encoded as a one-hot vector of dimension equal to the total number of possible goals in the game. In line 6, s_{self}^i is the self's observation state from the perspective of agent i , which is the same as the other's observation state from the perspective of agent j , s_{other}^j .

We consider a continuous vector \tilde{z}_{other} of the same dimension as z_{self} , such that the estimate of the other agent's

Algorithm 1 SOM training for one episode

```

1: procedure SELF OTHER-MODELING
2:   for k := 1, num_players do
3:      $\tilde{z}_{other}^k \leftarrow \frac{1}{n_{goals}} \mathbf{1}_{n_{goals}}$ 
4:   game.reset()
5:   for step := 1, episode_length do
6:      $s_{self}^i = s_{other}^j \leftarrow \text{game.get\_state}()$ 
7:      $\tilde{z}_{other}^{OH,i} = \text{one\_hot}(\text{argmax}(\text{softmax}(\tilde{z}_{other}^i)))$ 
8:      $\pi_{self}^i, V_{self}^i \leftarrow f_{self}^i(s_{self}^i, z_{self}^i, \tilde{z}_{other}^i; \theta^i)$ 
9:      $a_{self}^i \sim \pi_{self}^i$ 
10:    game.action( $a_{self}^i$ )
11:    for k := 1, num_inference_steps do
12:       $\tilde{z}_{other}^{GS,j} = \text{gumbel\_soft}(\text{softmax}(\tilde{z}_{other}^j))$ 
13:       $\tilde{\pi}_{other}^j \leftarrow f_{other}^j(s_{other}^j, \tilde{z}_{other}^{GS,j}, z_{self}^j; \theta^j)$ 
14:      loss = cross_entropy_loss( $\tilde{\pi}_{other}^j, a_{self}^i$ )
15:      loss.backward()
16:      update( $\tilde{z}_{other}^j$ )
17:    for k := 1, num_players do
18:      policy.update( $\theta^k$ )
    
```

goal is a sample from the Categorical distribution with class probabilities $\text{softmax}(\tilde{z}_{other})$. Thus, the estimate of the other's goal is given by the one-hot vector \tilde{z}_{other}^{OH} , as shown in line 7. At the beginning of each game, the estimate of the other's goal \tilde{z}_{other}^{OH} is randomly initialized, as illustrated in line 3, where $\mathbf{1}_{n_{goals}}$ represents a vector of all ones with the size equal to the number of possible goals.

In inference mode, the estimate of the other agent's goal is expressed as a sample from the Gumbel-Softmax distribution (Jang et al., 2016; Maddison et al., 2016), \tilde{z}_{other}^{GS} , as shown in line 12, where $\text{gumbel_soft}(p) = \text{softmax}[g + \log(p)]$, with g sampled from the Gumbel distribution and the softmax temperature $\tau = 1$. To update the estimate of the other's goal, we directly optimize \tilde{z}_{other} by using the cross-entropy loss to backpropagate through f_{other} (lines 14, 15, 16).




Figure 1. Our Self Other-Model (SOM) architecture.

The agents' policies are parametrized by long short-term memory (LSTM) cells (Hochreiter & Schmidhuber, 1997)

with two fully-connected linear layers, and exponential linear unit (ELU) (Clevert et al., 2015) activations. The weights of the networks are initialized with semi-orthogonal matrices, as described in Saxe et al. (2013) and zero bias.

3. Related Work

Multi-Agent Learning. Recent work in deep multi-agent RL focuses on partially visible, fully cooperative settings (Foerster et al., 2016a;b; Omidshafiei et al., 2017) and emergent communication (Lazaridou et al., 2016; Foerster et al., 2016a; Sukhbaatar et al., 2016; Das et al., 2017; Mordatch & Abbeel, 2017). Lerer & Peysakhovich (2017) design RL agents that are able to maintain cooperation in complex social dilemmas by generalizing a well-known game theoretic strategy called tit-for-tat (Axelrod, 2006), to multi-agent Markov games. Leibo et al. (2017) considers semi-cooperative multi-agent environments in which the agents develop cooperative or competitive strategies depending on the task type and reward structure. Similarly, Lowe et al. (2017) proposes a centralized actor-critic architecture for efficient training in settings with such mixed strategies. Our setting is different since we do not allow communication between the agents, so the players have to indirectly reason about others’ intentions from their observed behavior.

Intent Recognition. Research on plan, activity, and intent recognition has a long history, but it usually assumes domain knowledge or a form of rationality and uses techniques such as Bayesian inference or Hidden Markov Models (Sukthankar et al., 2014). The field of inverse reinforcement learning (IRL) (Russell, 1998; Ng et al., 2000; Abbeel & Ng, 2004) is also related to the problem considered here. IRL’s aim is to infer the reward function of an agent by observing its behavior, which is assumed to be nearly optimal. In contrast, our approach uses the observed actions of the other player to directly infer its goal in an online manner, which is then used by the agent when acting in the environment. This avoids the need for collecting offline samples of the other’s (state, action) pairs in order to estimate its reward function and use it to learn a policy. The more recent papers by Hadfield-Menell et al. (2016; 2017) are also concerned with the problem of inferring intentions, but their focus is on human-robot interaction and value alignment. Motivated by similar goals, Chandrasekaran et al. (2017) consider the problem of building a theory of AI’s mind, in order to improve human-AI interaction and the interpretability of AI systems. Recent work in cognitive science attempts to understand human decision-making by using a hierarchical model of social agency that infers human intentions for choosing a strategy (Kleiman-Weiner et al., 2016). However, none of these papers design algorithms that explicitly model other artificial agents in the environment or estimate their intentions, with the purpose of improving their decision

making.

Modeling Other Agents. Opponent modeling has been extensively studied in games of imperfect information. Yet most previous approaches focuses on developing models with domain-specific probabilistic priors or strategy parametrizations. In contrast, our work proposes a more general framework for opponent modeling. Davidson (1999) uses an MLP to predict opponent actions given a game history, but the agents cannot adapt to their opponents’ behavior online. Lockett et al. (2007) designs a neural network architecture to identify the opponent type by learning a mixture of weights over a given set of cardinal opponents, but the game does not unfold within the RL framework.

The closest work to ours is Foerster et al. (2017) and He et al. (2016). Foerster et al. (2017) designs RL agents that take into account the learning of other agents in the environment when updating their own policies. This enables the agents to discover self-interested yet collaborative strategies such as tit-for-tat in the iterated prisoner’s dilemma. While our work does not explicitly attempt to shape the learning of other agents, it has the advantage that agents can update their beliefs during an episode and change their strategies online to gain more reward. Our setting is also different in that it considers that each agent has some hidden information needed by the other player to maximize its return.

Our work is very much in line with He et al. (2016), where the authors build a general framework for modeling other agents in the reinforcement learning setting. He et al. (2016) proposes a model that jointly learns a policy and the behavior of opponents by encoding observations of the opponent into a DQN. Their Mixture of Experts architecture is able to discover different opponent strategy patterns in two competitive tasks. In our approach, rather than using hand designed features of the other agent’s behavior, the agent models others using its own policy. Another difference is that in this work, the agent runs an optimization over the input vector to infer the other agent’s hidden goal, rather than using a feed-forward network. In the experiments below, we show that SOM outperforms an adaptation of the method of He et al. (2016) to our setting.

4. Experiments

In this section, we evaluate our model SOM on three tasks:

- The *coin game*, in Section 4.2, which is a fully cooperative task where the agents’ roles are symmetric.
- The *recipe game*, in Section 4.3, which is adversarial, but with symmetric roles.
- The *door game*, in Section 4.4, which is fully cooperative but has asymmetric roles for the two players.

We compare SOM to three other baselines and to a model that has access to the ground truth of the other agent’s goal. All the tasks considered are created in the Mazebase grid-world environment (Sukhbaatar et al., 2015).

4.1. Baselines

TRUE-OTHER-GOAL (TOG): We provide an upper bound on the performance of our model given by a policy network which takes the other agent’s **true** goal as input z_{other} , as well as the state features s_{self} and its own goal z_{self} . Since this model has direct access to the true goal of the other agent, it does not need a separate network to model the behavior of the other agent. The architecture of TOG is the same as the one of SOM’s policy network, f .

NO-OTHER-MODEL (NOM): The first baseline we use only takes as inputs the observation states s_{self} and its own goal z_{self} . NOM has the same architecture as the one used for SOM’s policy network, f_{self} . This baseline does not explicitly model the other agent’s policy, goal, or actions.

INTEGRATED-POLICY-PREDICTOR (IPP): Starting with the architecture and inputs of NOM, we construct a stronger baseline, IPP, which has an additional final linear layer that outputs a probability distribution over the next action of the other agent. Besides the A3C loss used to train the policy of this network, we also add a cross-entropy loss to train the prediction of the other agent’s action, using observations of its true actions.

SEPARATE-POLICY-PREDICTOR (SPP): He et al. (2016) propose an opponent modeling framework based on DQN. In their approach, a neural network (separate from the learned Q-network) is trained to predict the opponents actions given hand crafted state information specific to the opponent. An intermediate hidden representation from this network is given as input to the Q-network.

We adapt the model of He et al. (2016) to our setting. In particular, we use A3C instead of DQN and we do not use the task-specific features used to represent the hidden goal of the opponent.

The resulting model, SPP, consists of two *separate* networks, a policy network for deciding the agent’s actions, and an opponent network for predicting the other agent’s actions. The opponent network takes as input its own state observation s_{self} and goal z_{self} , and outputs a probability distribution for the action taken by the other agent at the next step, as well as its hidden (recurrent) state. As in IPP, we train the opponent policy predictor with a cross-entropy loss using the true actions of the other agent. At each step, the hidden (recurrent) state outputted by this network is taken as input by the agent’s policy network, along with the observation state and its own goal. Both the policy network and the opponent policy predictor are LSTMs with the same

architecture as SOM.

In contrast to SOM, SPP does not explicitly infer the other agent’s goal. Rather, it builds an implicit model of the opponent by predicting the agent’s actions at each time step. In SOM, an inferred goal is given as additional input to the policy network. The analog of the inferred goal in SPP is the hidden (recurrent) state obtained from the opponent policy predictor which is given as an additional input to the policy network.

Training Details. In all our experiments, we train the agents’ policies using A3C (Mnih et al., 2016) with an entropy coefficient of 0.01, a value loss coefficient of 0.5, and a discount factor of 0.99. The parameters of the agents’ policies are optimized using Adam (Kingma & Ba, 2014) with $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1 \times 10^{-8}$, and weight decay 0. SGD with a learning rate of 0.1 was used for inferring the other agent’s goal, \tilde{z}_{other} .

The hidden layer dimension of the policy network was 64 for the Coin and Recipe Games and 128 for the Door Game. We use a learning rate of 1×10^{-4} for all games and models.

The observation state s is represented by few-hot vectors indicating the locations of all the objects in the environment (including the other player). The dimension of this input state is $1 \times n_{features}$, where the number of features is 384, 192, and 900 for the Coin, Recipe, and Door games, respectively.

For each experiment, we trained the models using 5 different random seeds. All the results shown are for 10 optimization updates of \tilde{z}_{other} at each step in the game, unless mentioned otherwise.

4.2. Coin Game.

First, we evaluate the model on a fully cooperative task, in which the agents can gain more reward when using both of their goals rather than only their own goal. So it is in the best interest of each agent to estimate the other player’s goal and use that information when taking actions. The game, shown in the left diagram of Figure 4, takes place on a 8×8 grid containing 12 coins of 3 different colors (4 coins of each color). At the beginning of each episode, the agents are randomly assigned one of the three colors. The action space consists of: go up, down, left, right, or pass. Once an agent steps on a coin, that coin disappears from the grid. The game ends after 20 steps. The reward received by both agents at the end of the game is given by the formula below:

$$R(c_{self}, c_{other}) = (n_{C_{self}}^{self} + n_{C_{self}}^{other})^2 + (n_{C_{other}}^{self} + n_{C_{other}}^{other})^2 - (n_{C_{neither}}^{self} + n_{C_{neither}}^{other})^2,$$

where $n_{C_{self}}^{other}$ is the number of coins of the self’s goal-color, which were collected by the other agents, and $n_{C_{neither}}^{self}$ is




Figure 2. Coin Strategy: Average number of collected coins per episode corresponding to the color of the Self (blue), Other (red), or Neither (green) by the agents using TOG (left), SOM (center-left), NOM (center), IPP (center-right), and SPP (right). The optimal strategy is to pick up as many Self as Other coins on average, across a number of episodes, and as few Neither coins as possible. Being able to collect more Other than Neither coins indicates that the agent is able to accurately infer the other agent’s color early enough during some of the episodes and uses this information to collect more Other, instead of Neither coins, which increases its reward. The TOG model learns to collect just as many Self as Other coins, while the baseline models only learn to collect more Self coins, but cannot distinguish between the Other and Neither coins. SOM learns to collect significantly more Other coins than Neither. This shows that SOM converges to a closer-to-optimal strategy using its guess of the other’s goal.

the number of coins corresponding to neither of the agents’ goals, collected by the self. For the example in Figure 4, agent 1 has $C_{self} = \text{orange}$ and $C_{other} = \text{cyan}$, while agent 2’s C_{self} is cyan and C_{other} is orange. $C_{neither}$ is red for both agents.

The role of the penalty for collecting coins that do not correspond to any of the agents’ goals is to avoid convergence to a greedy policy in which the agents can gain a non-negligible amount of reward by collecting all the coins in their proximity, without any regard to their color.




Figure 3. Coin Performance: Average reward in the Coin game by SOM (green), TOG (blue), NOM (red), IPP (magenta), and SPP (orange). SOM performs better than all baselines.

To maximize its return, each agent needs maximize the number of collected coins of its own or its collaborator’s color, and minimize the number of coins of the remaining color. Hence, when both agents are able to infer their collaborators’ goals with high accuracy and as early as possible in the game, they can use that information to maximize their

shared utility.

Figure 3 shows the mean and standard deviation of the reward across 5 runs with different random seeds obtained by SOM. Our model clearly outperforms all other baselines on this task. We also show the empirical upper bound on the reward using the model which takes as input the true color assigned to the other agent.

Figure 2 analyzes the strategies of the different models by looking at the proportion of coins of each type collected by the agents. The optimal strategy is for each agent to maximize $n_{C_{self}}^{self} + n_{C_{other}}^{self}$ and minimize $n_{C_{neither}}^{self}$. Due to the randomness in the initial conditions (in particular, the locations of coins in the environment), this amounts to each agent collecting an equal number of coins of its own and of the other’s color on average, across a large number of episodes (i.e. $\bar{n}_{C_{self}}^{self} = \bar{n}_{C_{other}}^{self}$).

Indeed, this is the strategy learned by the model with perfect information of the other agent’s goal (TOG). SOM also learns to collect significantly more Other than Neither coins (although not as many as Self coins), indicating its ability to distinguish between the two types, at least during some of the episodes. This means that SOM can accurately infer the other agent’s goal early enough during the episode and use that information to collect more Other Coins, thus gaining more reward than if it were only using its own goal to direct its actions.

In contrast, the agents trained with the three baseline models collect significantly more Self coins, and as many Other as Neither coins on average. This shows that they learn to use their own goal for gaining reward, but they are unable to use the hidden goal of the other agent for further increasing their returns. Even if IPP and SPP are able to predict the actions of the other player with an accuracy of about 50%, they do not learn to distinguish between the coins that would increase (Other) and those that would decrease (Neither)

their reward. This shows the weaknesses of using an implicit model of the other agent to maximize reward on certain tasks.




Figure 4. Illustration of the Coin (left), Recipe (center), and Door (right) games. Above each diagram, we show the agents’ goals (not visible to one another).

4.3. Recipe Game.

Agents in adversarial scenarios have competing goals, so the ability of inferring the opponent’s goal could better inform the agent’s actions. With this motivation in mind, we evaluate our model on a game in which the agents have to craft certain compositional recipes, each containing multiple items found in the environment. The agents are given as input the names of their goal-recipes, without the corresponding components needed to make it. The resources in the environment are scarce, so only one of the agents can craft its recipe within one episode.

As illustrated in Figure 4 (center), there are 4 types of items: {sun, star, moon, lightning} and 4 recipes: {sun, sun, star}; {star, star, moon}; {moon, moon, lightning}; {lightning, lightning, sun}. The game is played in a 4×6 grid, which contains 8 items in total, 2 of each type.

At the beginning of each episode, we randomly assign a recipe to one of the agents, and then we randomly pick a recipe for the other agent so that it has overlapping items with the recipe of the first agent. This ensures that the agents are competing for resources within each episode. At the end of the episode, each agent receives a reward of +1 for crafting its own recipe and a penalty of -0.1 for each item it picked up not needed for making its recipe.

We designed the layout of the grid so that neither agent has an initial advantage by being closer to the scarce resource. At the beginning of each episode, one of the agents starts on the left-most column of the grid, while the other one starts on the right-most column, at the same y-coordinate. Their initial y-coordinate as well as which agent starts on the left/right is randomized. Similarly, one item of each of the 4 different types is placed at random in the grid formed

by the second and third columns of the maze, from left to right. The rest of the items are placed in the forth and fifth columns, so that the symmetry with respect to the vertical axis is preserved (i.e. items of the same type are placed at the same y-coordinate, and symmetric x-coordinates).

Agents have six actions to choose from: pass, go up, down, left, right, or pick up (for picking up an item, which then disappears from the grid). The first agent to take an action is randomized. The game ends after 50 steps.

We pretrain all baselines on a version of the game which does not have overlapping recipes, in order to ensure that all the models learn to pick up the corresponding items, given a recipe as goal. All of the models learn to craft their assigned recipes $\sim 90\%$ of the time on this simpler task. Then, we continue training the models on the adversarial task in which their recipes overlap in each episode. SOM is initialized with a pretrained NOM network.

Figure 5 shows the winning fraction for different pairs played against each other in the Recipe game. For the first 100k episodes, the models are not being trained. We can see that SOM significantly outperforms NOM, IPP, and SPP, winning $\sim 75 - 80\%$ of the time, while the baselines can only win $\sim 15 - 20\%$ of the games. SPP ties against NOM, and TOG outperforms SOM by a large margin. We also played the same types of agents against each other and they all win $\sim 40 - 50\%$ of the games.

4.4. Door Game.

In this section, we show that on a collaborative task with asymmetric roles and multiple possible partners, the agents can learn to figure out what role they should be playing in each game based on their partners’ actions.

In the Door game, two agents are located in a 5×9 grid, with 5 goals behind 5 doors on the left wall, and 5 switches on the right wall of the grid. The game starts with the two players in random squares on the grid, except for the ones occupied by the goals, doors, or switches, as illustrated in Figure 4. Agents can take any of the five actions: go up, down, left, right or pass. An action is invalid if it moves the player outside of the border or to a square occupied by a block or closed door. Both agents receive +3 reward when either one of them steps on its goal and they are penalized -0.1 for each step they take. The game ends when one of them gets to its goal or after 22 steps. All the goals are behind doors which are open only as long as one of the agents sits on the corresponding switch for that door.

At the beginning of an episode, each of the two players is randomly selected from a pool of 5 agents and receives as input a random number from 1 to 5 corresponding to its goal. Each of the 5 agents has its own policy which gets updated at the end of each episode they play. Note that the agents’




Figure 5. Recipe Performance: Average fraction of success in the Recipe game by SOM-NOM (left), SOM-IPP (center-left), SOM-SPP (center-center), SOM-TOG (center-right) NOM-SPP (right). The plots show the performance of SOM with 5 optimization updates of \tilde{z}_{other} at each step in the game.




Figure 6. Door Performance: Average fraction of success on the Door game by SOM (green), TOG (blue), NOM (red), IPP (magenta), and SPP (orange). On average, SOM performs better than all baselines.

identities are not visible (i.e. there is no indication in the state features that specifies the id’s of the agents playing during a given episode). This restriction is important in order to ensure that the agents cannot gain advantage by specializing into the two roles needed to win (i.e. goal-goer and switch-puller) and identifying the specialization of the other player by simply observing its unique id.

The agents need to cooperate in order to receive reward. In contrast to our previous tasks, the two players must take different roles. In fact, the player who sits on the switch should ignore its own goal and instead infer the other’s goal, while the player who goes to its goal does not need to infer the other’s goal, but only use its own. In order to sit on the correct switch, an agent has to infer the other player’s goal from their observed actions. The only way in which an agent can use its own policy to model the other player is if each agent learns to play both roles of the game, i.e. go to its own goal and also open its collaborator’s door by sitting on the corresponding switch. Indeed, we see that the agents learn to play both roles and they are able to use their own policies to infer the other player’s goals when needed.

Fig 6 shows the mean and standard deviation of the winning fraction obtained by one of the agents on the Door game. While our model is still able to outperform the three baselines, the gap between the performance of our model and that of IPP or SPP (an approximate version of (He et al., 2016)) is smaller than in the previous tasks. However, this is a more difficult task for our model since it needs the agent to learn both roles before effectively using its own policy to infer the other agent’s goal. The plot shows that SOM actually performs worse than IPP and SPP during the initial part of training, before outperforming them. Nevertheless, we see that SOM training allows the agents to play both roles in an asymmetric cooperative game, and to infer the goal and role of the other player.

4.5. Analyzing the goal inference

In this section we further analyze the ability of the SOM models to infer other’s intended goals.




Figure 7. Inference Accuracy during Training: The mean fraction of episodes in which the agent correctly infers the other’s goal for the Coin (left), Recipe (center), and Door (right) games, as a function of training epoch. The estimate of the other’s goal is considered correct if it remains accurate during all the following steps in the game.

Figure 7 shows the fraction of episodes in which the goal of the other agent is correctly inferred. We consider that the goal is correctly inferred only when the estimate of the other’s goal remains accurate until the end of the game, so that we avoid counting the episodes in which the agent might infer the correct goal by chance at some intermediate step in the game. In all the games, the SOM agent learns to infer the other player’s goal with a mean accuracy ranging

from $\sim 60 - 80\%$. Comparing the second plot in Figure 2 with the left plot in Figure 7, one can observe that the SOM agent starts distinguishing Other from Neither coins after approximately 2M training episodes, which coincides with the time when the mean accuracy of the inferred goal converges to $\sim 75\%$. The Door Game (right) presents higher variance since the agents learn to use and infer the other’s goal at different stages during training.




Figure 8. Inference Step Distribution: Cumulative distribution of the step t_{inf} at which the goal of the other player is correctly inferred (i.e. $\tilde{z}_{other}^t = z_{other}, \forall t \geq t_{inf}$) for the Coin (left), Recipe (center) and Door (right) games. We define this step so that $\tilde{z}_{other} = z_{other}$ for all the remaining steps in the game. The distribution is computed over the subset of runs in which the goal is correctly inferred before the end of the game ($\sim 70 - 80\%$ of all runs). A total of 1000 runs with trained SOM models were used to compute this distribution.

Figure 8 shows the cumulative distribution of the step at which the goal of the other player is correctly inferred (and remains the same until the end of the game). The cumulative distribution is computed over the episodes in which the goal is correctly inferred before the end of the game. In the Coin (blue) and Recipe (red) games, 80% of the times the agent correctly infers the goal of the other, it does so in the first five steps. The distribution for the Door (green) game indicates that the agent needs more steps on average to correctly infer the goal. This explains in part why the SOM agent only slightly outperforms the SPP baseline. If the agent does not infer the other’s goal early enough in the episode, it cannot efficiently use it to maximize its return.




Figure 9. Performance Variation with Number of Inference Steps: Average reward (blue) and average fraction of episodes in which the goal of the other agent is correctly inferred (red) obtained by the SOM agent as a function of the number of inference steps used for estimating the other’s goal for the Coin (left), Recipe (center), and Door (right) games. The error bars represent 1 standard deviation.

Figure 9 shows how the performance of the agent varies with

the number of optimization updates performed on \tilde{z}_{other} at each step in the game. As expected, the agent’s reward (blue) generally increases with the number of inference steps, as does the fraction of episodes in which the goal is correctly inferred. One should note that increasing the number of inference steps from 10 to 20 only translates into less than 0.45% performance gain, while increasing it from 1 to 5 translates into a performance gain of 6.9% on the Coin game, suggesting that there is a certain threshold above which increasing the number of inference steps will not significantly improve performance.

5. Discussion

Summary. In this paper, we introduced a new approach for inferring other agents’ hidden goals from their behavior and using those estimates to choose actions. We demonstrated that the agents are able to estimate others’ hidden goals in both cooperative and competitive settings, which enables them to converge to better policies. In the proposed tasks, using an explicit model of the other player led to better performance than simply considering the other agent as part of the environment.

Strengths. Some of the main advantages of our method are its simplicity and flexibility. This method does not require any extra parameters to model other agents in the environment, can be trained with any reinforcement learning algorithm, and can be easily integrated with any network architecture. SOM can also be adapted to settings with more than two players, since the agent can use its own policy to model the behavior of any number of agents and infer their goals. Moreover, it can be easily generalized to numerous other environments and tasks.

Limitations. Our approach is based on the assumption that the agents are identical or that their transition functions are independent and identically distributed. Hence, the framework is expected to be more suitable for symmetric games, in which the agents share a fixed set of goals and have similar abilities, and we expect a degradation of performance for asymmetric games. Our experiments confirm this observation. Another limitation of SOM is that it requires a longer training time than other baselines, since we backpropagate through the network at each step. However, their online nature is essential in adapting to the behavior of other agents in the environment.

Future Work. We plan to extend this work by evaluating the models on more complex environments and model deviations from the assumption that the players have identical policies, given a certain goal and state of the world. Another important avenue for future research is to design models that can adapt to non-stationary strategies of others in the environment, as well as to tasks with hierarchical goals.

Acknowledgements

This work was partially supported by ONR grant N00014-13-1-0646. The authors wish to thank Alex Peysakhovich and Adam Lerer for helpful discussions.

References

- Abbeel, P. and Ng, A. Y. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 1. ACM, 2004.
- Axelrod, R. Agent-based modeling as a bridge between disciplines. *Handbook of computational economics*, 2: 1565–1584, 2006.
- Chandrasekaran, A., Yadav, D., Chattopadhyay, P., Prabhu, V., and Parikh, D. It takes two to tango: Towards theory of ai’s mind. *arXiv preprint arXiv:1704.00717*, 2017.
- Clevert, D.-A., Unterthiner, T., and Hochreiter, S. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.
- Das, A., Kottur, S., Moura, J. M., Lee, S., and Batra, D. Learning cooperative visual dialog agents with deep reinforcement learning. *arXiv preprint arXiv:1703.06585*, 2017.
- Davidson, A. Using artificial neural networks to model opponents in texas holdem. *Unpublished manuscript*, 1999.
- Foerster, J., Assael, Y. M., de Freitas, N., and Whiteson, S. Learning to communicate with deep multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 2137–2145, 2016a.
- Foerster, J. N., Assael, Y. M., de Freitas, N., and Whiteson, S. Learning to communicate to solve riddles with deep distributed recurrent q-networks. *arXiv preprint arXiv:1602.02672*, 2016b.
- Foerster, J. N., Chen, R. Y., Al-Shedivat, M., Whiteson, S., Abbeel, P., and Mordatch, I. Learning with opponent-learning awareness. *arXiv preprint arXiv:1709.04326*, 2017.
- Hadfield-Menell, D., Russell, S. J., Abbeel, P., and Dragan, A. Cooperative inverse reinforcement learning. In *Advances in neural information processing systems*, pp. 3909–3917, 2016.
- Hadfield-Menell, D., Milli, S., Abbeel, P., Russell, S. J., and Dragan, A. Inverse reward design. In *Advances in Neural Information Processing Systems*, pp. 6768–6777, 2017.
- He, H., Boyd-Graber, J., Kwok, K., and Daumé III, H. Opponent modeling in deep reinforcement learning. In *Proceedings of The 33rd International Conference on Machine Learning*, pp. 1804–1813, 2016.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Jang, E., Gu, S., and Poole, B. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- Kingma, D. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kleiman-Weiner, M., Ho, M. K., Austerweil, J. L., Littman, M. L., and Tenenbaum, J. B. Coordinate to cooperate or compete: abstract goals and joint intentions in social interaction. In *COGSCI*, 2016.
- Lazaridou, A., Peysakhovich, A., and Baroni, M. Multi-agent cooperation and the emergence of (natural) language. *arXiv preprint arXiv:1612.07182*, 2016.
- Leibo, J. Z., Zambaldi, V., Lanctot, M., Marecki, J., and Graepel, T. Multi-agent reinforcement learning in sequential social dilemmas. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pp. 464–473. International Foundation for Autonomous Agents and Multiagent Systems, 2017.
- Lerer, A. and Peysakhovich, A. Maintaining cooperation in complex social dilemmas using deep reinforcement learning. *arXiv preprint arXiv:1707.01068*, 2017.
- Lockett, A. J., Chen, C. L., and Miikkulainen, R. Evolving explicit opponent models in game playing. In *Proceedings of the 9th annual conference on Genetic and evolutionary computation*, pp. 2106–2113. ACM, 2007.
- Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, P., and Mordatch, I. Multi-agent actor-critic for mixed cooperative-competitive environments. *arXiv preprint arXiv:1706.02275*, 2017.
- Maddison, C. J., Mnih, A., and Teh, Y. W. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pp. 1928–1937, 2016.
- Mordatch, I. and Abbeel, P. Emergence of grounded compositional language in multi-agent populations. *arXiv preprint arXiv:1703.04908*, 2017.

- Ng, A. Y., Russell, S. J., et al. Algorithms for inverse reinforcement learning. In *Icml*, pp. 663–670, 2000.
- Omidshafiei, S., Pazis, J., Amato, C., How, J. P., and Vian, J. Deep decentralized multi-task multi-agent rl under partial observability. *arXiv preprint arXiv:1703.06182*, 2017.
- Russell, S. Learning agents for uncertain environments. In *Proceedings of the eleventh annual conference on Computational learning theory*, pp. 101–103. ACM, 1998.
- Saxe, A. M., McClelland, J. L., and Ganguli, S. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*, 2013.
- Shapley, L. S. Stochastic games. *Proceedings of the National Academy of Sciences*, 39(10):1095–1100, 1953. ISSN 0027-8424. doi: 10.1073/pnas.39.10.1095. URL <http://www.pnas.org/content/39/10/1095>.
- Sukhbaatar, S., Szlam, A., Synnaeve, G., Chintala, S., and Fergus, R. Mazebase: A sandbox for learning from games. *arXiv preprint arXiv:1511.07401*, 2015.
- Sukhbaatar, S., Fergus, R., et al. Learning multiagent communication with backpropagation. In *Advances in Neural Information Processing Systems*, pp. 2244–2252, 2016.
- Sukthankar, G., Geib, C., Bui, H. H., Pynadath, D., and Goldman, R. P. *Plan, activity, and intent recognition: Theory and practice*. Newnes, 2014.