

Learning Existing Social Conventions via Observationally Augmented Self-Play

Anonymized for review

Author main disciplines: behavioral economics, AI

Abstract

In order for artificial agents to coordinate effectively with people, they must act consistently with existing conventions (e.g. how to navigate in traffic, which language to speak, or how to coordinate with teammates). A group's conventions can be viewed as a choice of equilibrium in a coordination game. We consider the problem of an agent learning a policy for a coordination game in a simulated environment and then using this policy when it enters an existing group. When there are multiple possible conventions we show that learning a policy via multi-agent reinforcement learning (MARL) is likely to find policies which achieve high payoffs at training time but fail to coordinate with the real group into which the agent enters. We assume access to a small number of samples of behavior from the true convention and show that we can augment the MARL objective to help it find policies consistent with the real group's convention. In three environments from the literature - traffic, communication, and team coordination - we observe that augmenting MARL with a small amount of imitation learning greatly increases the probability that the strategy found by MARL fits well with the existing social convention. We show that this works even in an environment where standard training methods very rarely find the true convention of the agent's partners.

Introduction

A ubiquitous feature of social interaction is a need for individuals to coordinate (Lewis 1969; Shoham and Tennenholtz 1997). A common solution to the coordination problem is the establishment of social conventions which control daily tasks such as choosing which side of the road to drive on, who should get right of way during walking, what counts as polite, what language to speak, or how a team should apportion tasks. If we seek to construct artificial agents that can coordinate with humans, they must be able to act according to existing conventions.

In game theory, Nash equilibria are strategies for all players such that if everyone behaves according to them no individual can improve their payoff by deviating. In game theoretic models a convention is one of multiple possible equilibria in a coordination game (Lewis 1969). Stated in these terms our agent's task is to construct a policy that does well when paired with the equilibrium being played by existing agents.

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

There has been great recent progress in constructing policies that can do well in both single and multi-agent environments using deep reinforcement learning (Mnih et al. 2015; Silver et al. 2017). Deep RL methods typically require orders of magnitude more experience in an environment to learn good policies than humans (Lake et al. 2017), so agents are typically trained in simulation before being deployed onto the real task. For single-agent tasks such as robotics, simulation consists of a simulated model of the environment such as MuJoCo (Todorov, Erez, and Tassa 2012).

In zero-sum two-player environments (e.g. Go), it is the policy of the other player that is simulated during training. Typically, policies for both players are trained simultaneously or iteratively, in a process called self-play. Self-play produces successful policies because if self-play converges then it converges to an equilibrium of the game (Fudenberg and Levine 1998) and in two-player, zero-sum games all equilibria are minimax/maximin strategies (von Neumann 1928). Thus, a fully converged strategy is guaranteed to be unexploitable for the task of interest (e.g. playing Go with a human champion).

Constructing agents that can cooperate and coordinate with each other to achieve goals (e.g. work together as team to finish a task) has been a long running topic in multi-agent reinforcement learning (MARL) research (Walker and Wooldridge 1995; Stone and Sutton 2001; Kapetanakis and Kudenko 2002b; Lowe et al. 2017). However, this literature typically assumes that the cooperating agents will continue to interact with those with whom they have been co-trained (this is sometimes referred to as "centralized training with distributed execution"). In this case, if MARL converges, it finds an equilibrium and since agents will play with the same partners they trained with they will achieve these equilibrium payoffs.

Unfortunately, agents are no longer guaranteed equilibrium payoffs if there are multiple equilibria and agents must coordinate with those they were *not* trained with (in other words, when we remove the centralized training assumption). For example, training self-driving cars in a virtual environment may lead to agents that avoid crashing into each other during training, but drive on the wrong side of the road relative to the society they will enter.

In this paper we propose to give the agent access to a small amount of observations of existing social behavior, i.e. samples of (state, action) pairs from the test time environ-

ment. We focus on how such data, though it is not enough to purely clone good policies, can be used in the training process to learn the best response to the policies of the future partners. Our key assumption is that the existing environment already has some existing, stable, social conventions. Thus, our assumption is that future partners will be playing *some* equilibrium strategies. In simple environments, agents could simply enumerate all possible equilibria and choose the one most consistent with the data. However, in more complex environments this becomes intractable. We propose to guide self-play toward the correct equilibrium by training with a joint MARL and behavioral cloning objective. We call this method observationally augmented self-play (OSP).

We consider OSP in several multi-agent situations with multiple conventions: a multi-agent traffic game (Sukhbaatar, Fergus, and others 2016), a particle environment combining navigation and communication (Mordatch and Abbeel 2017; Lowe et al. 2017) and a Stag Hunt game where agents must take risks to accomplish a joint goal (Peysakhovich and Lerer 2017). In each of these games we find that self-play can converge to multiple, incompatible, conventions. We find that OSP is able to learn correct conventions in these games with a small amount of observational data. Our results in the Markov Stag Hunt show that OSP can learn conventions it observes even when those conventions are very unlikely to be learned using MARL alone.

We do not claim that OSP is the ultimate approach to constructing agents that can learn social conventions. The success of OSP depends on both the game and the type of objective employed, thus an exploration of alternative algorithms is an important direction for future work. Our key result is that the combination of (a) a small number of samples from trajectories of a multi-agent game, and (b) knowledge that test time agents are playing some equilibrium gives much stronger test time performance than either component alone.

Related Work

OSP is related to existing work on adding reward shaping in MARL (Kapetanakis and Kudenko 2002a; Devlin and Kudenko 2011). However the domain of interest differs slightly as reward shaping is typically used to cause all agents in a group to converge to a high-payoff equilibrium whereas we are interested in using shaping to guide training to select the correct test-time equilibrium.

Recent work has pointed out that strategies learned via a single instance of independent MARL can overfit to other agents’ policies during training (Lanctot et al. 2017). This work differs from ours in that it suggests the training of a single best response to a mixture of heterogeneous policies. This increases the robustness of agent policies but does not solve the problem of multiple, incompatible conventions that we study here.

The approach of combining supervised learning from trajectories with RL has been studied in the single agent case (Hester et al. 2017). In that work the trajectories are expert demonstrations and are used to guide RL to an optimal policy. In this case supervision is used to speed up learning while in our work the trajectories are used to select among many

possible optima (equilibria) which may be equivalent at training time but not at test time. However this literature explores many methods for combining imitation learning and RL and some of these techniques may be interesting to consider in the multi-agent setting.

Conventions in Markov Games

A partially observed Markov game G (Littman 1994) consists of a set of players $P = \{1, \dots, N\}$, a set of states \mathcal{S} , a set of actions for every player A_i with the global set $\mathcal{A} = \times_{i \in P} A_i$, a transition function $\tau : \mathcal{S} \times \mathcal{A} \rightarrow \Delta \mathcal{S}$, a reward function for each player that takes as input the current state and actions $R_i : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. Players have observation functions $O_i : \mathcal{S} \rightarrow \mathcal{O}_i$ and can condition their behavior on these observations. Markov policies for each player are functions $\pi_i : \mathcal{O}_i \rightarrow A_i$. Let Π_i denote the set of all policies available to a player and $\Pi = \times_{i \in P} \Pi_i$ be the set of joint policies.¹

We use the standard notation π_{-i} to denote the policy vector for all players other than i . A set of policies $\pi \in \Pi$ and a (possibly random) initial state $s_0 \in \mathcal{S}$ defines a (random) trajectory of rewards for each player. We let the value function $V_i(s, \pi_i, \pi_{-i})$ denote the discounted expectation of this trajectory. The best response starting at state s for player i to π_{-i} is $BR(s, \pi_{-i}) = \operatorname{argmax}_{\pi_i \in \Pi_i} V_i(s, \pi_i, \pi_{-i})$. We let s_0 be the (possibly random) initial state of the game.

There are many ways to extend the notion of a Nash equilibrium to the case of stochastic games. We will consider the Markov perfect best response. We denote by $BR(\pi_{-i})$ the policy (or policies) which is a best response starting at *any* state and consider equilibria to be policies for each player π such that each $\pi_i \in BR(\pi_{-i})$.²

We consider games with multiple, incompatible, conventions. Formally we say that conventions (equilibrium policy sets) π and π' are incompatible if the compound policy (π_i, π'_{-i}) is not an equilibrium.

The goal of training is to compute a policy π_i which is a best response to the existing convention π_{-i}^{test} . During training, the agent has access to the environment but receives only a limited set of observations of π^{test} in the form of a set \mathcal{D} of state-action pairs sampled from π^{test} . This is summarized in Figure 1.

We denote a generic element of \mathcal{D} by $d = (s, a_k)$ which is a (state, action) pair for agent k . Let \mathcal{D}_j denote the subset of \mathcal{D} which includes actions for agent j .

The dataset \mathcal{D} may be insufficient to identify a best response to all possible policies π^{test} consistent with \mathcal{D} . However, the set of equilibrium policy sets is typically much

¹We consider only Markov policies in this work so that we can work with individual state-action pairs, although our same approach could be applied across observed trajectories to learn non-Markov policies (i.e. policies conditioned on their full history).

²There are weaker notions, for example, requiring that policies are best responses at every state reached during play. It is known that Markov perfect equilibria are harder to find by learning (Fudenberg and Levine 1998) and it is interesting to consider whether different kinds of choices (e.g. on-policy vs. off-policy learning) can make stronger or weaker convergence guarantees. However, these questions are outside the scope of this paper.

smaller than all possible policy sets. Therefore, if we assume that all agents are minimizing regret then we must only consider equilibrium policy sets consistent with \mathcal{D} .

Given a game G and dataset \mathcal{D} , a brute force approach to learn a policy compatible with the conventions of the group the agent will enter would be to compute the equilibrium π^* of G that maximizes the likelihood of \mathcal{D} . Formally this is given by

$$\pi^* \in \operatorname{argmax}_{\pi} \sum_i \sum_{(s,a) \in \mathcal{D}_i} \log(\pi_i(s,a))$$

s.t. $\forall i, \pi_i \in BR(\pi_{-i})$.

This constrained optimization problem quickly becomes intractable and we will instead try to find an equilibrium using multi-agent learning, and use \mathcal{D} to increase the probability learning converges to it.

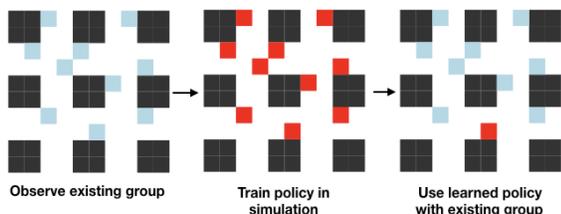


Figure 1: A graphical description of our problem. Our central question is how to use observational data to make sure that training converges to the same convention as the group which our agent will enter.

Observationally Initialized Best Response Dynamics

To get an intuition about how data can be used during training we will first study a learning rule where analytic results are more easily obtained. We will begin with the simplest multi-agent learning algorithm: *best response dynamics* in a 2 player Markov game where equilibria are in pure policies and are incompatible.

In best response dynamics each player begins with a policy initialized at π_i^0 . Players alternate updating their policy with the best response $\pi_i^t = BR(\pi_{-i}^{t-1})$. When there are multiple best responses we assume there is a (non-randomized) tie-breaking rule used to select one. Given an equilibrium A we denote by Π^A the basin of attraction of A (the set of initial states from which BR dynamics converge to A .)

A naive way to use the observed data is to force the policies be consistent with \mathcal{D} at each step of the best response dynamic by changing it at each states where it differs. However, this can introduce new equilibria to the game.³

³For a simple example consider a game where 5 agents choose an action A or B and receive reward equal to the number of other agents whose actions match theirs. In this case there are equilibria where all agents choose A or all agents choose B . If we restrict one agent to always choose A we can introduce a new equilibrium where 4 agents choose B and one agent chooses A .

In the context of reward shaping it is well known that the way to avoid the introduction of new equilibria is to use potential based reward shaping (Devlin and Kudenko 2011), or, equivalently, use our information to only change the initialization of learning (Wiewiora 2003). We will follow this advice and study *observationally initialized best response dynamics*. We begin with a policy π_0 chosen at random. However, for every player i and state-action pair (s, a) in the data \mathcal{D} we form $\bar{\pi}^0$ by setting the corresponding action of $\pi_i^0(s)$ to a . We then perform best response dynamics from this new initialization.

We will now discuss a class of games for which we can guarantee that observationally initialized best response dynamics have a larger basin of attraction for the equilibrium from which \mathcal{D} was drawn relative to standard best response dynamics. This class is a generalization of a commonly used matrix game class: games with *strategic complements* (Bulow, Geanakoplos, and Klemperer 1985). For our purposes strategic complements corresponds to assuming that one’s partners behave more consistently with some convention than one is also more incentivized to behave according to that convention.⁴ In existing work strategic complements are defined with respect to a single action rather than a Markov policy. To generalize to Markov games we introduce a notion of distance between policies:

Definition 1 (Policy Closeness) *Given a player i and target policy A we say that policy π is **weakly closer to A than policy π'** if on all states either $\pi(s) = \pi'(s)$ or $\pi(s) = A(s)$. We denote this by $\pi \succeq_A \pi'$.*

Policy closeness gives us a partial ordering on policies which we use to generalize the idea of strategic complements.

Definition 2 (Markov Strategic Complements) *A Markov game exhibits **Markov strategic complements** if for any equilibrium $A = (A_1, A_2)$ we have that $\pi_i \succeq_{A_i} \pi'_i$ implies that $BR(\pi_i) \succeq_{A_{-i}} BR(\pi'_i)$.*

Let \mathcal{D}_A be a dataset drawn from equilibrium A by sampling states and their equilibrium actions. Let $\bar{\Pi}^A(\mathcal{D}_A)$ be the basin of attraction of A given observationally initialized best response dynamics.

Theorem 1 *If a game where best-response dynamics always converge exhibits Markov strategic complements then for any \mathcal{D}_A drawn from a equilibrium A $\Pi^A \subseteq \bar{\Pi}^A(\mathcal{D}_A)$ and there exists $(s, A_i(s))$ such that if $(s, A_i(a)) \in \mathcal{D}_A$ then $\Pi^A \subset \bar{\Pi}^A(\mathcal{D}_A)$.*

We relegate the proof to the Appendix. Roughly, it has two steps: first, we show that if π_0 is in the basin of attraction of A then anything closer to A is also in the basin. Second, we show that there is an initial state that is not in the basin

⁴In economic applications the notion of strategic complements is utilized in production games and roughly corresponds either the idea of network effects (the more people use some product the higher a consumer’s utility is from using that product) or a joint monotonicity condition (if Firm X produces cars and firm Y produces car washing materials if firm X produces more cars then firm Y sees higher demand for car washing materials). See the Supplement for a more formal discussion.

of attraction of best response dynamics but is in the basin of attraction of observationally initialized best response dynamics. Because initialization can increase the basin of attraction without introducing any new equilibria the observed data can strictly improve the probability that we learn a policy consistent with the observed agents.

Observationally Augmented Self-Play

We wish to use the insights from initialization in environments where function approximation (e.g. with deep neural networks) is required. However, if the policy is computed via function approximation, it is not clear how to ‘initialize’ its value at particular states. Specifically, the policy at the small number of states in \mathcal{D} can only be expected to generalize if the approximation captures the regularities of the game, which will only be true after some interaction in the environment. Therefore, we consider a procedure where consistency with \mathcal{D} is enforced during training and smoothly decays over time.

We consider training with stochastic gradient descent using a loss function which is a linear combination of the likelihood of the data (a supervised objective) plus the policy gradient estimator of the reward in the game (we denote by \mathcal{L}^{PG} the negative of this quantity). Formally, each agent receives rewards given by

$$\mathcal{L}_i^{OSP}(\pi_i, \pi_{-i}) = \mathcal{L}^{PG}(\pi_i | \pi_{-i}) + \lambda \sum_{(s,a) \in \mathcal{D}_i} \log(\pi_i(s, a))$$

with respect to the parameters θ_i of its policy.⁵

We optimize the joint objective using MARL. At any time t when making an update to the parameters of each agent θ_i we will take gradient steps of the form

$$\nabla_t^{OSP}(\theta_i, \mathcal{D}) = \nabla_t^{PG}(\theta_i) + \lambda_t \nabla^{sup}(\theta_i, \mathcal{D}).$$

Where ∇^{PG} is our policy gradient and ∇^{sup} is the gradient of the supervised objective at θ_i .⁶

Experiments

Our main analysis is experimental and we use three environments from the literature: traffic, language games, and risky coordination. Our main results are:

Result 1 (Experiments 1,2,3) *OSP greatly increases the probability our agent learns a convention compatible with test time agents in situations where standard self-play by itself does not guarantee good test time performance and \mathcal{D} is insufficient to learn a good policy by behavioral cloning alone.*

⁵Note this is different from reward shaping as the probability that a state is reached does not affect the supervised loss of the policy.

⁶As with the best response dynamics above using a compound objective with a constant λ can, in theory, introduce new equilibria during training. To be sure this does not occur we can anneal the weight on the supervised loss over time with $\lim_{t \rightarrow \infty} \lambda_t = 0$. In practice, however, using a fixed λ in our environments appeared to create policies that were still consistent with test time equilibria thus suggesting that if new equilibria were introduced they did not have large basins of attraction for our policy-gradient based learning procedures.

Result 2 (Experiment 3) *OSP can find conventions that have a small basin of attraction for MARL alone. Thus OSP can be used in situations where self-play will rarely find a test-time compatible convention.*

For all experiments, we represent the states using simple neural networks. The first two experiments have relatively low dimensional state representations so we use two layer MLPs with 128 hidden units per layer and ELU nonlinearities. Our third experiment has a grid structure so we represent the state using the convolutional neural network architecture from (Peysakhovich and Lerer 2017).

For RL training we use A3C using a multi-agent variant of the pytorch-a3c package (Kostrikov 2018) run on 40 threads. We use 20 step returns and $\gamma = .99$. We use the Adam method for optimization (Kingma and Ba 2014). For OSP, we add the supervised term to the A3C loss with $\lambda = 1$, using minibatches from \mathcal{D} of size 20. Environment-specific parameters are detailed in the subsections below. In each experiment we compare the performance of OSP for various sizes of \mathcal{D} . We populate \mathcal{D} with actions for all agents for states sampled at uniform intervals from true test time play.

Traffic

We first consider a multi-agent traffic navigation game inspired by (Sukhbaatar, Fergus, and others 2016; Resnick et al. 2018). Each of 10 agents begins at a random edge of a grid and can move around in each of the 4 cardinal directions or stand still (see Figure 2). Each agent is given a goal to reach. When the agent reaches the goal they receive a reward of +1 and then a new goal. If agents collide with another agent or wall they receive a negative reward (−5 for colliding with another agent, −.1 for colliding with a wall). Agents do not have full observation of the environment. Rather, they have access to the position of their goal relative to themselves and a local view of any other agents nearby.

We train agents for a total of 800,000 episodes. We found it necessary to ramp the collision penalty linearly over the first 400,000 episodes to avoid agents becoming stuck in the local minima of never moving.

We train 20 replicates and see that two incompatible conventions emerge. This can be seen in Figure 2 where we plot payoffs to an agent from one replicate paired with 9 agents from another. We visualize the conventions by taking the empirical average action taken by any agent in any of the possible traffic coordinate locations (Figure 2 panel 3). We find that the two conventions that emerge are similar to the real world: either agents learn to drive on the left of the road or they learn to drive on the right.

We now train agents using OSP and insert them into environments with 9 pre-converged agents. The test time payoffs to the transplanted agent for various sizes of \mathcal{D} are shown in Figure 2 panel 4 top. The dashed line corresponds to the expected payoff of an agent trained using standard self-play (no observations). We see 20 observations (2 observations for each of the 10) agents is sufficient to guide OSP to compatible conventions. The bottom panel shows that this is not enough data to train a good strategy via behavioral cloning alone (i.e. using just the supervised objective).

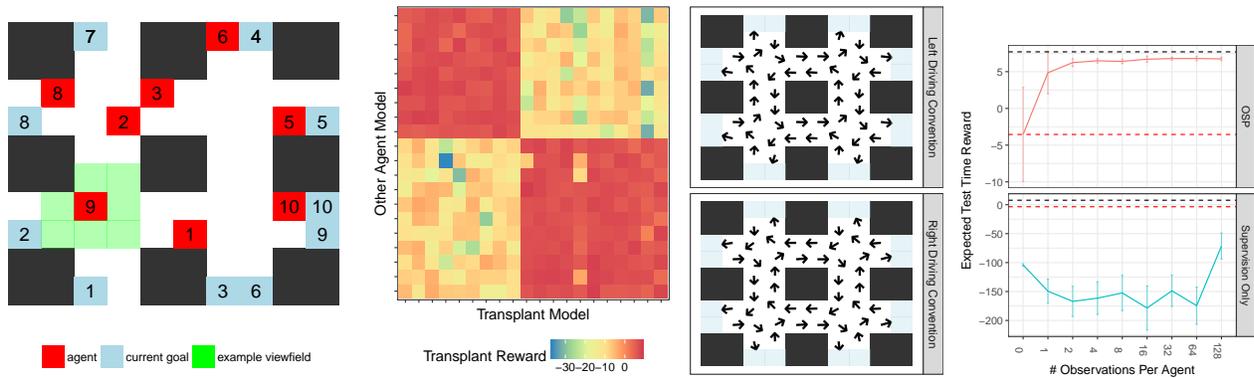


Figure 2: In the traffic game 10 agents navigate a grid world to reach goals (first panel). When agents reach their goal, they receive a reward and are assigned a new goal. Agents have a local field of view (indicated in green) and know the relative position of their next goal from their current position. Two incompatible conventions emerge across multiple training runs. MARL agents achieve high rewards with their training partners (diagonal) but lower rewards with incompatible partners (second panel). Visualizing these conventions shows that they differ in which side of the road to drive on (third panel). Given only 2 observations for each agent OSP learns the correct test time convention (fourth panel), approaching the reward of centralized training (black dashed line), while supervised training does not. The dashed line indicates the average reward of separate self-play training, and error bars represent 95% confidence intervals derived from independent replicates.

Language

An important type of convention is language. There is a resurgence of interest in the deep RL community in using communication games to construct agents which are able to communicate (Foerster et al. 2016; Lazaridou, Peysakhovich, and Baroni 2017; Lowe et al. 2017).

We now apply OSP to the cooperative communication task in the particle environment studied by (Lowe et al. 2017) and (Mordatch and Abbeel 2017). In this environment there are two agents, a speaker and a listener, and three landmarks (blue, red, or green). One of the landmarks is randomly chosen as a goal and the reward for both agents at each step is equal to the distance of the listener from the landmark. However, which landmark is the goal during a particular episode is only known to the speaker who is able to produce a communication output from a set of 20 symbols. To solve the cooperation task, agents thus need to evolve a simple ‘language’. This language only requires 3 symbols to be used, but this still allows for at least 6840 incompatible conventions (one symbol per landmark).

In this experiment we use a lower discount factor of $\gamma = .8$ and as suggested by Lowe et al. we also use a centralized critic. It was shown in prior work that if artificial agents learn language by self-play they can learn arbitrary languages which may not be compatible with new partners (Lazaridou, Peysakhovich, and Baroni 2017). Indeed, when we pair two agents who were trained separately they clearly do not speak the same language - i.e. they cannot communicate and so achieve low payoffs.

We look at the effect of adding observational data to the training of either a speaker or listener (we train a total of 135 replicates to convergence). In the case of the speaker (whose policy is a simple map from goals to communication symbols) supervision is sufficient to learn the a good test-time language.

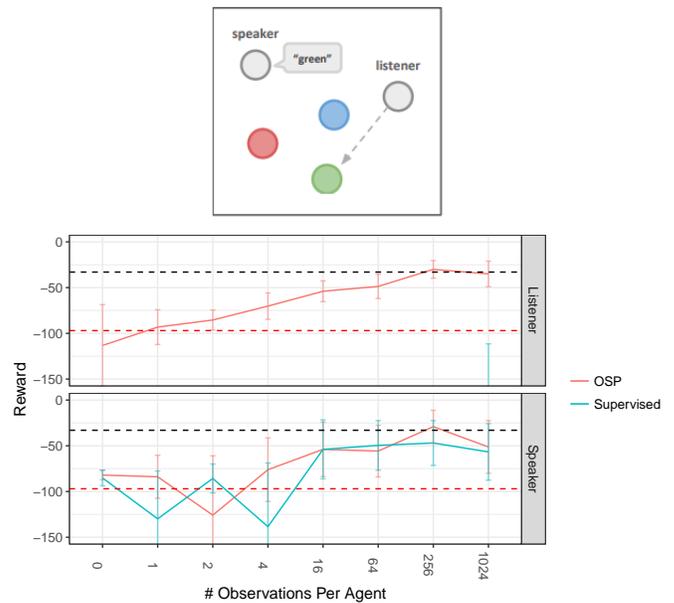


Figure 3: In the speaker/listener particle environment (first panel taken from the original paper) the speaker must communicate information to a listener who navigates to a landmark. When a pair are trained together they can reach high payoffs (black dashed line) but pairs trained separately with self-play (red dashed line) perform quite poorly. OSP leads to agents that reach high payoffs with their test time partner. Supervision alone is insufficient to construct good Listener agents (plot is below the y axis for $N < 1024$). Error bars represent confidence intervals derived from independent replicates.

However, pure behavioral cloning fails catastrophically for the listener. Again, OSP with a relatively small number of observations is able to achieve high payoffs (Figure 3).

Risky Coordination

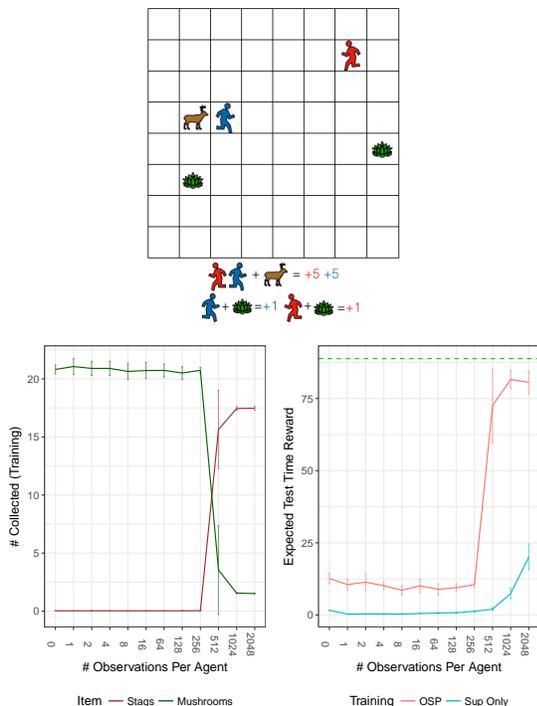


Figure 4: In the Stag Hunt two agents either coordinate to catch a Stag or choose to grab plants alone. Hunting is risky because hunting without a partner is ineffective. In this game standard independent MARL typically converges to the safe equilibrium rather than discovering the risky high-payoff Hunt equilibrium (middle panel shows training behavior at convergence). An agent trained separately cannot coordinate when matched with a test-time partner who choose to hunt. However OSP, with samples from the Hunt equilibrium, can indeed converge to the Hunting convention during training reaching a payoff almost as high as two Hunters trained together (green line). Error bars represent confidence intervals derived from independent replicates.

We now consider a risky coordination game known as the Stag Hunt. The matrix game version of the Stag Hunt has both agents choosing either to Hunt (an action that require coordination) or Forage (a safe action). Foraging yields a sure (low) payoff whereas Hunting yields a high payoff if the other agent chooses to Hunt also and a low payoff if one shows up alone. It is known that in both matrix and Markov versions of Stag Hunt games many standard self-play based algorithms yield agents that converge to the inefficient equilibrium in which both agents choose safe actions. This happens because while our partner is not hunting effectively (i.e. early in training), the payoff to hunting ourselves is quite

low. Thus, the basin of attraction of joint hunting is much smaller than the basin of attraction of both foraging.

This situation is different from the ones in traffic and language: here there are multiple conventions (hunt or forage) but they are not payoff equivalent (hunting is better) nor do they have similar sized basins of attraction (hunting is very difficult to find via standard independent MARL).

We use the Markov version of Stag Hunt introduced by Peysakhovich and Lerer where two agents live on a 8×8 grid. Two plants and a stag are placed at random locations. If an agent moves over a plant, it receives a reward of 1. Stepping on the stag gives a reward of 5 to both players if they step on it simultaneously, otherwise there is no reward. When either a plant or stag is stepped on, it restarts in a new location. Games last 100 rounds.

We start by constructing a test time hunting partner by inducing joint hunting strategies in 40 replicates. Because MARL by itself does not find hunt equilibria, we construct a hunting partner by training an agent under a modified payoff structure (payoff of 0 for plants; payoff of 0.1 for unilateral hunting).

We then test whether we can train agents in the *original* game who can coordinate with test time partners that hunt. We use OSP with varying amounts of data from the hunting agents. We see that with moderate amounts of data OSP often converges to the hunting convention at test time even though two agents trained together using independent MARL fail to find the high payoff joint stag equilibrium in *any* of the replicates. As a result, OSP outperforms even centralized self-play because the observations of the risky partner guide the agent to a better equilibrium. As with the traffic and language environments above we see that pure behavioral cloning is insufficient to construct good test time strategies (Figure 4).

Conclusion

Conventions are an important part of social behavior and many multi-agent environments support multiple conventions as equilibria. If we want to construct artificial agents that can adapt to people (rather than requiring people to adapt to them) these agents need to be able to act according to the existing social conventions. In this paper we have discussed how a simple procedure that combines small amounts of imitation learning with self-play can lead to agents that can learn social conventions.

There are many open questions remaining in the study of conventions and building agents that can learn them quickly. OSP uses a straightforward combination of RL and behavioral cloning. It would be interesting to explore whether ideas from the learning with expert demonstrations literature (Hester et al. 2017). In addition, OSP follows current deep RL paradigms splits strategy construction into a training and a test phase. An interesting extension is to consider the OSP training strategies can be fine-tuned during test time.

There is a large recent interest in hybrid systems which include both human and artificially intelligent participants (Shirado and Christakis 2017). Thus, another key extension of our work is to understand whether techniques like OSP can construct agents that can interact with humans in more complex environments.

References

- [Bulow, Geanakoplos, and Klemperer 1985] Bulow, J. I.; Geanakoplos, J. D.; and Klemperer, P. D. 1985. Multimarket oligopoly: Strategic substitutes and complements. *Journal of Political economy* 93(3):488–511.
- [Devlin and Kudenko 2011] Devlin, S., and Kudenko, D. 2011. Theoretical considerations of potential-based reward shaping for multi-agent systems. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, 225–232. International Foundation for Autonomous Agents and Multiagent Systems.
- [Foerster et al. 2016] Foerster, J.; Assael, Y.; de Freitas, N.; and Whiteson, S. 2016. Learning to communicate with deep multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems*, 2137–2145.
- [Fudenberg and Levine 1998] Fudenberg, D., and Levine, D. K. 1998. *The theory of learning in games*, volume 2. MIT press.
- [Hester et al. 2017] Hester, T.; Vecerik, M.; Pietquin, O.; Lanctot, M.; Schaul, T.; Piot, B.; Horgan, D.; Quan, J.; Sendonaris, A.; Dulac-Arnold, G.; et al. 2017. Deep q-learning from demonstrations. *arXiv preprint arXiv:1704.03732*.
- [Kapetanakis and Kudenko 2002a] Kapetanakis, S., and Kudenko, D. 2002a. Improving on the reinforcement learning of coordination in cooperative multi-agent systems. In *Proceedings of the Second Symposium on Adaptive Agents and Multi-agent Systems (AISB02)*.
- [Kapetanakis and Kudenko 2002b] Kapetanakis, S., and Kudenko, D. 2002b. Reinforcement learning of coordination in cooperative multi-agent systems. *AAAI/IAAI 2002*:326–331.
- [Kingma and Ba 2014] Kingma, D., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- [Kostrikov 2018] Kostrikov, I. 2018. Pytorch implementations of asynchronous advantage actor critic. <https://github.com/ikostrikov/pytorch-a3c>.
- [Lake et al. 2017] Lake, B. M.; Ullman, T. D.; Tenenbaum, J. B.; and Gershman, S. J. 2017. Building machines that learn and think like people. *Behavioral and Brain Sciences* 40.
- [Lanctot et al. 2017] Lanctot, M.; Zambaldi, V.; Grusly, A.; Lazaridou, A.; Perolat, J.; Silver, D.; Graepel, T.; et al. 2017. A unified game-theoretic approach to multiagent reinforcement learning. In *Advances in Neural Information Processing Systems*, 4193–4206.
- [Lazaridou, Peysakhovich, and Baroni 2017] Lazaridou, A.; Peysakhovich, A.; and Baroni, M. 2017. Multi-agent cooperation and the emergence of (natural) language. In *International Conference on Learning Representations*.
- [Lewis 1969] Lewis, D. 1969. *Convention: A philosophical study*. John Wiley & Sons.
- [Littman 1994] Littman, M. L. 1994. Markov games as a framework for multi-agent reinforcement learning. In *Machine Learning Proceedings 1994*. Elsevier. 157–163.
- [Lowe et al. 2017] Lowe, R.; Wu, Y.; Tamar, A.; Harb, J.; Abbeel, P.; and Mordatch, I. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *arXiv preprint arXiv:1706.02275*.
- [Mnih et al. 2015] Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518(7540):529–533.
- [Mordatch and Abbeel 2017] Mordatch, I., and Abbeel, P. 2017. Emergence of grounded compositional language in multi-agent populations. *arXiv preprint arXiv:1703.04908*.
- [Peysakhovich and Lerer 2017] Peysakhovich, A., and Lerer, A. 2017. Prosocial learning agents solve generalized stag hunts better than selfish ones. *arXiv preprint arXiv:1709.02865*.
- [Resnick et al. 2018] Resnick, C.; Kulikov, I.; Cho, K.; and Weston, J. 2018. Vehicle community strategies. *arXiv preprint arXiv:1804.07178*.
- [Shirado and Christakis 2017] Shirado, H., and Christakis, N. A. 2017. Locally noisy autonomous agents improve global human coordination in network experiments. *Nature* 545(7654):370–374.
- [Shoham and Tennenholtz 1997] Shoham, Y., and Tennenholtz, M. 1997. On the emergence of social conventions: modeling, analysis, and simulations. *Artificial Intelligence* 94(1-2):139–166.
- [Silver et al. 2017] Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; et al. 2017. Mastering the game of go without human knowledge. *Nature* 550(7676):354.
- [Stone and Sutton 2001] Stone, P., and Sutton, R. S. 2001. Scaling reinforcement learning toward robocup soccer. In *ICML*, volume 1, 537–544. Citeseer.
- [Sukhbaatar, Fergus, and others 2016] Sukhbaatar, S.; Fergus, R.; et al. 2016. Learning multiagent communication with backpropagation. In *Advances in Neural Information Processing Systems*, 2244–2252.
- [Todorov, Erez, and Tassa 2012] Todorov, E.; Erez, T.; and Tassa, Y. 2012. Mujoco: A physics engine for model-based control. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, 5026–5033. IEEE.
- [von Neumann 1928] von Neumann, J. 1928. Zur theorie der gesellschaftsspiele. *Mathematische annalen* 100(1):295–320.
- [Walker and Wooldridge 1995] Walker, A., and Wooldridge, M. 1995. Understanding the emergence of conventions in multi-agent systems. In *ICMAS*, volume 95, 384–389.
- [Wiewiora 2003] Wiewiora, E. 2003. Potential-based shaping and q-value initialization are equivalent. *Journal of Artificial Intelligence Research* 19:205–208.

Appendix

Relationship Between Markov Strategic Complements and Strategic Complements

The original definition of strategic complements comes from games with continuous actions used to model multiple firms in a market. In the simplest example we have multiple firms which produce units of goods (x_1, x_2, \dots, x_n) . The revenue function of each firm i is $R_i(x_i, x_{-i})$ where R_i is smooth, strictly concave, increasing and has $R_i(x_i = 0, x_{-i}) = 0$.

The goods are strategic complements if $\frac{\partial^2 R_i}{\partial x_i \partial x_j} > 0$, in other words goods are strategic complements if “more ‘aggressive’ play... by one firm... raises the marginal profitabilities [of the others].” (Bulow, Geanakoplos, and Klemperer 1985) Firms have costs of production given by $c_i(x_i)$ which has $c(0) = 0$, $c'(0) = 0$, is convex, and increasing. Thus each firm’s objective function is

$$R_i(x_i, x_{-i}) - c_i(x_i)$$

If firm $-i$ is producing x_{-i} then firm i ’s best response $x_i^*(x_{-i})$ sets

$$\frac{\partial R_i(x_i^*(x_{-i}), x_{-i})}{\partial x_i} = \frac{\partial c_i(x_i^*(x_{-i}))}{\partial x_i}.$$

Given the definition of strategic complements above this means that $\frac{\partial x_i^*}{\partial x_j} > 0$ for all other firms j .

Strategic complements implies our Markov strategic complements in a matrix game with multiple equilibria (since any firm changing their production level higher or lower causes other firms to also want to change their production). Markov strategic complements is weaker than strategic complements in matrix games since it only pins down how best responses to shift when others change to equilibrium actions rather than any action shift (though if action spaces in each state were totally ordered one could amend the definition to keep all of the properties).

Proof of Main Theorem

Proof 1

Lemma 1: *In a Markov strategic complements (MSC) game, any policy π in the basin of attraction of an equilibrium A remains there under observational initialization, i.e. $\pi_i \in \Pi^A \implies \bar{\pi}_i(\mathcal{D}_A) \in \Pi^A$.*

We define the operator $BR^{(k)}$ as k iterations of the best response operator;

$$BR^{(k)}(\pi_i) = BR^i(BR^{-i}(BR^i(\dots BR^{-i}(\pi_i))))).$$

Consider an initial policy $\pi_i \in \Pi^A$ for some equilibrium A . There exists k_c such that $BR^{(k_c)} = A_i$. Now consider an observationally initialized policy $\bar{\pi}_i(\mathcal{D}_A)$ for some dataset \mathcal{D} drawn from A_i . By definition, this implies that $\bar{\pi}_i \succ_A \pi_i$. Now, since the game is MSC,

$$BR^{-i}(\bar{\pi}_i(\mathcal{D}_A)) \succ_A BR^{-i}(\pi_i).$$

By repeated application of the MSC property, we find that for all k ,

$$BR^{(k)}(\bar{\pi}_i(\mathcal{D}_A)) \succ_A BR^{(k)}(\pi_i).$$

To conclude, we note that $BR^{(k_c)}(\bar{\pi}_i(\mathcal{D}_A)) \succ_A A_i$, which implies $BR^{(k_c)}(\bar{\pi}_i(\mathcal{D}_A)) = A_i$.

Lemma 2: *In a MSC game with a finite number of states, there exists a state s such that for any \mathcal{D}_A that contains the state-action pair $(s, A_i(s))$, there is a policy not in the basin of attraction of A but which enters the basin of attraction of A under observational initialization.*

Consider a policy $\pi_i^0 \notin \Pi_A$, and order the states lexicographically (s_1, s_2, \dots, s_M) . Now consider the sequence of policies π_i^k where $\pi_i^k(s_l) = A_i(s_l)$ for $l < k$ and $\pi_i^k(s_l) = \pi_i^0(s_l)$ for $l \geq k$. We know that $\pi_i^M = A_i \in \Pi_A$, therefore there exists some t such that $\pi_i^t \notin \Pi_A$ and $\pi_i^{t+1} \in \Pi_A$. Now, consider a dataset \mathcal{D}_A containing the state-action pair $(s_{t+1}, A_i(s_{t+1}))$. Then $\bar{\pi}_i^t(\mathcal{D}_A) \succ_A \pi_i^{t+1}$. As discussed in the last section, if $\pi_i^{t+1} \in \Pi_A$ and $\bar{\pi}_i^t(\mathcal{D}_A) \succ_A \pi_i^{t+1}$, then $\bar{\pi}_i^t(\mathcal{D}_A) \in \Pi_A$. Therefore, for any dataset containing s_{t+1} , the policy π_i^t enters the basin of attraction of A under observational initialization.