

EVOLVING INTRINSIC MOTIVATIONS FOR ALTRUISTIC BEHAVIOR

Anonymous authors

Paper under double-blind review

ABSTRACT

Multi-agent cooperation is an important feature of the natural world. Many tasks involve individual incentives that are misaligned with the common good, yet a wide range of organisms from bacteria to insects and humans are able to overcome their differences and collaborate. Therefore, the emergence of cooperative behavior amongst self-interested individuals is an important question for the fields of multi-agent reinforcement learning (MARL) and evolutionary theory. Here, we study a particular class of multi-agent problems called intertemporal social dilemmas (ISDs), where the conflict between the individual and the group is particularly sharp. By combining MARL with appropriately structured natural selection, we demonstrate that individual inductive biases for cooperation can be learned in a model-free way. To achieve this, we introduce an innovative modular architecture for deep reinforcement learning agents which supports multi-level selection. We present results in two challenging environments, and interpret these in the context of cultural and ecological evolution.

1 INTRODUCTION

Nature shows a substantial amount of cooperation at all scales, from microscopic interactions of genomes and bacteria to species-wide societies of insects and humans (Maynard Smith & Szathmari, 1997). This is in spite of natural selection pushing for short-term individual selfish interests (Darwin, 1859). In its purest form, altruism can be favored by selection when cooperating individuals preferentially interact with other cooperators, thus realising the rewards of cooperation without being exploited by defectors (Hamilton, 1964a;b; Dawkins, 1976; Santos et al., 2006; Fletcher & Doebeli, 2009). However, many other possibilities exist, including kin selection, reciprocity and group selection (Nowak, 2006; Úbeda & Duéñez-Guzmán, 2011; Trivers, 1971; Nowak & Sigmund, 2005; Wilson, 1975; Smith, 1964).

Lately the emergence of cooperation among self-interested agents has become an important topic in multi-agent deep reinforcement learning (MARL). Leibo et al. (2017) and Hughes et al. (2018) formalize the problem domain as an *intertemporal social dilemma* (ISD), which generalizes matrix game social dilemmas to Markov settings. Social dilemmas are characterized by a trade-off between collective welfare and individual utility. As predicted by evolutionary theory, self-interested reinforcement-learning agents are typically unable to achieve the collectively optimal outcome, converging instead to defecting strategies (Leibo et al., 2017; Pérolat et al., 2017). The goal is to find multi-agent training regimes in which individuals resolve social dilemmas, i.e., cooperation emerges.

Previous work has found several solutions, belonging to three broad categories: 1) opponent modelling (Foerster et al., 2017; Kleiman-Weiner et al., 2016), 2) long-term planning using perfect knowledge of the game’s rules (Lerer & Peysakhovich, 2017; Peysakhovich & Lerer, 2018) and 3) a specific intrinsic motivation function drawn from behavioral economics (Hughes et al., 2018). These hand-crafted approaches run at odds with more recent end-to-end model-free learning algorithms, which have been shown to have a greater ability to generalize (e.g. (Espeholt et al., 2018)). We propose that evolution can be applied to remove the hand-crafting of intrinsic motivation, similar to other applications of evolution in deep learning.

Evolution has been used to optimize single-agent hyperparameters (Jaderberg et al., 2017), implement black-box optimization (Wierstra et al., 2008), and to evolve neuroarchitectures (Miller et al., 1989; Stanley & Miikkulainen, 2002), regularization (Chan et al., 2002), loss functions (Jaderberg

et al., 2018; Houthoofd et al., 2018), behavioral diversity (Conti et al., 2017), and entire reward functions (Singh et al., 2009). These principles tend to be driven by single-agent search and optimization or competitive multi-agent tasks. Therefore there is no guarantee of success when applying them in the ISD setting. More closely related to our domain are evolutionary simulations of predator-prey dynamics (Yong & Miikkulainen, 2001), which used enforced subpopulations to evolve populations of neurons which are sampled to form the hidden layer of a neural network.¹

To address the specific challenges of ISDs, the system we propose distinguishes between optimization processes that unfold over two distinct time-scales: (1) the fast time-scale of learning and (2) the slow time-scale of evolution (similar to Hinton & Nowlan, 1987). In the former, individual agents repeatedly participate in an intertemporal social dilemma using a fixed intrinsic motivation. In the latter, that motivation is itself subject to natural selection in a population. We model this intrinsic motivation as an additional additive term in the reward of each agent (Chentanez et al., 2005). We implement the intrinsic reward function as a two-layer fully-connected feed-forward neural network, whose weights define the genotype for evolution. We propose that evolution can help mitigate this intertemporal dilemma by bridging between these two timescales via an intrinsic reward function.

Evolutionary theory predicts that evolving individual intrinsic reward weights across a population who interact uniformly at random does not lead to altruistic behavior (Axelrod & Hamilton, 1981). Thus, to achieve our goal, we must structure the evolutionary dynamics (Nowak, 2006). We first implement a “Greenbeard” strategy (Dawkins, 1976; Jansen & van Baalen, 2006) in which agents choose interaction partners based on an honest, real-time signal of cooperativeness. We term this process *assortative matchmaking*. Although there is ecological evidence of assortative matchmaking (Keller & G. Ross, 1998), it cannot explain cooperation in all taxa (Grafen et al., 1990; Henrich, 2004; Gardner & West, 2010). Moreover it isn’t a general method for multi-agent reinforcement learning, since honest signals of cooperativeness are not normally observable in the ISD models typically studied in deep reinforcement learning.

To address the limitations of the assortative matchmaking approach, we introduce an alternative modular training scheme loosely inspired by ideas from the theory of multi-level (group) selection (Wilson, 1975; Henrich, 2004), which we term *shared reward network* evolution. Here, agents are composed of two neural network modules: a policy network and a reward network. On the fast timescale of reinforcement learning, the policy network is trained using the modified rewards specified by the reward network. On the slow timescale of evolution, the policy network and reward network modules evolve separately from one another. In each episode every agent has a distinct policy network but the same reward network. As before, the fitness for the policy network is the individual’s reward. In contrast, the fitness for the reward network is the collective return for the entire group of co-players. In terms of multi-level selection theory, the policy networks are the lower level units of evolution and the reward networks are the higher level units. Evolving the two modules separately in this manner prevents evolved reward networks from overfitting to specific policies. This evolutionary paradigm not only resolves difficult ISDs without handcrafting but also points to a potential mechanism for the evolutionary origin of social inductive biases.

2 METHODS

We varied and explored different combinations of parameters, namely: (1) environments {Harvest, Cleanup}, (2) reward network features {prospective, retrospective}, (3) matchmaking {random, assortative}, and (4) reward network evolution {individual, shared, none}. We describe these in the following sections.

2.1 INTERTEMPORAL SOCIAL DILEMMAS

In this paper, we consider Markov games (Littman, 1994) within a MARL setting. Specifically we study intertemporal social dilemmas (Leibo et al., 2017; Hughes et al., 2018), defined as games in which individually selfish actions produce individual benefit on short timescales but have negative impacts on the group over a longer time horizon. This conflict between the two timescales characterizes the intertemporal nature of these games. The tension between individual and group-level rationality identifies them as social dilemmas (e.g. the famous Prisoner’s Dilemma).

¹See also Potter & Jong (2000) and Panait & Luke (2005) for reviews of other evolutionary approaches to cooperative multi-agent problems.




Figure 1: Screenshots from (a) the Cleanup game, (b) the Harvest game. The size of the agent-centered observation window is shown in (b). The same size observation was used in all experiments.

We consider two dilemmas, each implemented as a partially observable Markov game on a 2D grid (see Figure 1). In the *Cleanup* game, agents tried to collect apples (reward +1) that spawned in a field at a rate inversely related to the cleanliness of a geographically separate aquifer. Over time, this aquifer filled up with waste, lowering the respawn rate of apples linearly, until a critical point past which no apples could spawn. Episodes were initialized with no apples present and zero spawning, thus necessitating cleaning. The dilemma occurred because in order for apples to spawn, agents must leave the apple field and clean, which conferred no reward. However if all agents declined to clean (defect), then no rewards would be received by any. In the *Harvest* game, again agents collected rewarding apples. The apple spawn rate at a particular point on the map depended on the number of nearby apples, falling to zero once there were no apples in a certain radius. There is a dilemma between the short-term individual temptation to harvest all the apples quickly and the consequential rapid depletion of apples, leading to a lower total yield for the group in the long-term. For more details, see the Appendix.

2.2 MODELING SOCIAL PREFERENCES AS INTRINSIC MOTIVATIONS

In our model, there are three components to the reward that enter into agents’ loss functions (1) total reward, which is used for the policy loss, (2) extrinsic reward, which is used for the extrinsic value function loss and (3) intrinsic reward, which is used for the intrinsic value function loss.

The *total reward* for player i is the sum of the extrinsic reward and an intrinsic reward as follows:

$$r_i(s_i, a_i) = r_i^E(s_i, a_i) + u_i(\mathbf{f}_i). \quad (1)$$

The *extrinsic reward* $r_i^E(s, a)$ is the environment reward obtained by player i when it takes action a_i from state s_i , sometimes also written with a time index t . The *intrinsic reward* $u(\mathbf{f})$ is an aggregate social preference across features \mathbf{f} and is calculated according to the formula,

$$u_i(\mathbf{f}_i|\theta) = \mathbf{v}^T \sigma(\mathbf{W}^T \mathbf{f}_i + \mathbf{b}), \quad (2)$$

where σ is the ReLU activation function, and $\theta = \{\mathbf{W}, \mathbf{v}, \mathbf{b}\}$ are the parameters of a 2-layer neural network with 2 hidden nodes. These parameters are evolved based on fitness (see Section 2.3). The elements of $\mathbf{v} = (v_1, v_2)$ can be seen to approximately correspond to a linear combination of the coefficients related to advantagenous and disadvantagenous inequity aversion mentioned in Hughes et al. (2018), which were found via grid search in this previous work, but are here evolved.

The feature vector \mathbf{f}_i is a player-specific quantity that other agents can transform into intrinsic reward via their reward network. Each agent has access to the same set of features, with the exception that its own feature is demarcated specially. The features themselves are a function of recently received or expected future (extrinsic) reward for each agent. In Markov games the rewards received by different players may not be aligned in time. Thus, any model of social preferences should not be

overly influenced by the precise temporal alignment of different players’ rewards. Intuitively, they ought to depend on comparing temporally averaged reward estimates between players, rather than instantaneous values. Therefore, we considered two different ways of temporally aggregating the rewards.

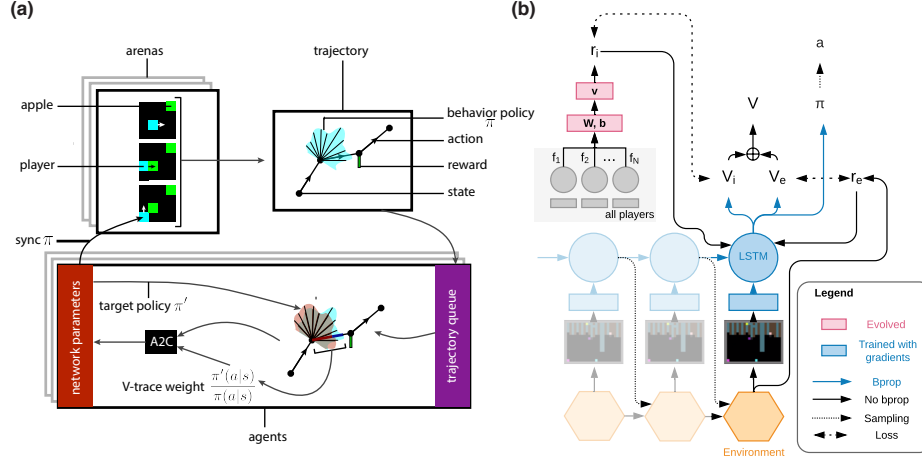


Figure 2: (a) Agent A_j adjusts policy $\pi_j(s, a|\phi)$ using off-policy importance weighted actor-critic (V-Trace) (Espeholt et al., 2018) by sampling from a queue with (possibly stale) trajectories recorded from 500 actors acting in parallel arenas. (b) The architecture includes intrinsic and extrinsic value heads, a policy head, and evolution of the reward network.

The *retrospective* method derives intrinsic reward from whether an agent judges that other agents have been actually (extrinsically) rewarded in the recent past. The *prospective* variant derives intrinsic reward from whether other agents are expecting to be (extrinsically) rewarded in the near future.² For the retrospective variant, $f_{ij} = e_j^t$, where the temporally decayed reward e_j^t for the agents $j = 1, \dots, N$ are updated at each timestep t according to

$$e_j^t = \eta e_j^{t-1} + r_j^{E,t}, \quad (3)$$

and $\eta = 0.975$. The prospective variant uses the value estimates V_j^{est} for f_{ij} and has a stop-gradient before the reward network module so that gradients don’t flow back into other agents.

2.3 ARCHITECTURE AND TRAINING

We used the same training framework as in Jaderberg et al. (2018), which performs distributed asynchronous training in multi-agent environments, including population-based training (PBT) (Jaderberg et al., 2017). We trained a population of 50 agents³ with policies $\{\pi_i\}$, from which we sampled 5 players in order to populate each of 500 arenas running in parallel. Within each arena, an episode of the environment was played with the sampled agents, before resampling new ones. Agents were sampled using one of two matchmaking processes (described in more detail below). Episode trajectories lasted 1000 steps and were written to queues for learning, from which weights were updated using V-Trace (Figure 2(a)). More details are in the Appendix.

The set of weights evolved included learning rate, entropy cost weight, and reward network weights θ^4 . The parameters of the policy network ϕ were inherited in a Lamarckian fashion as in (Jaderberg et al., 2017). Furthermore, we allowed agents to observe their last actions $a_{i,t-1}$, last intrinsic rewards ($r_{i,t-1}^E(s_i, a_i)$), and last extrinsic rewards ($u_{i,t-1}(f_i)$) as input to the LSTM in the agent’s neural network.

²Our terms prospective and retrospective map onto the terms intentional and consequentialist respectively as used by Lerer & Peysakhovich (2017); Peysakhovich & Lerer (2018).

³Similar to as in (Espeholt et al., 2018) we distinguish between an “agent” which acts in the environment according to some policy, and a “learner” which updates the parameters of a policy. In principle, a single agent’s policy may depend on parameters updated by several separate learners.

⁴We can imagine that the reward weights are simply another set of optimization hyperparameters since they enter into the loss.

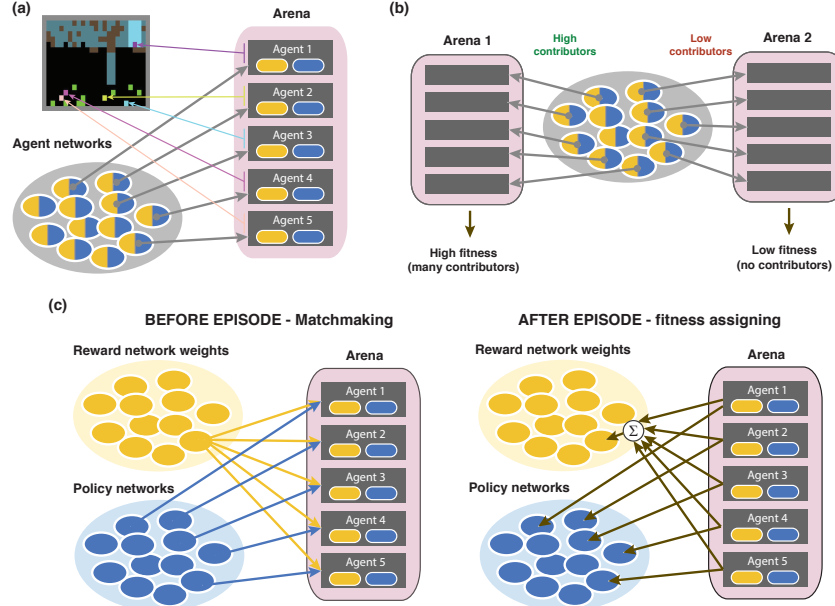


Figure 3: (a) Agents assigned and evolved with individual reward networks. (b) Assortative matchmaking, which preferentially plays cooperators with other cooperators and defectors with other defectors. (c) A single reward network is sampled from the population and assigned to all players, while 5 policy networks are sampled and assigned to the 5 players individually. After the episode, policy networks evolve according to individual player returns, while reward networks evolve according to aggregate returns over all players.

The objective function was identical to that presented in Espeholt et al. (2018) and comprised three components: (1) the value function gradient, (2) policy gradient, and (3) entropy regularization, weighted according to hyperparameters baseline cost and entropy cost (see Figure 2(b)).

Evolution was based on a fitness measure calculated as a moving average of total episode return, which was a sum of apples collected minus penalties due to tagging, smoothed as follows:

$$F_j^i = (1 - \nu)F_j^{i-1} + \nu R_j^i, \quad (4)$$

where $\nu = 0.001$ and R_j^i is the return obtained on episode i by agent j (or reward network j in the case of the shared reward network evolution (see Section 2.5 and Appendix for details)).

2.4 RANDOM VS. ASSORTATIVE MATCHMAKING

Matches were determined according to two methods: (1) random matchmaking and (2) assortative matchmaking. Random matchmaking simply selected uniformly at random from the pool of agents to populate the game, while cooperative matchmaking first ranked agents within the pool according to a metric of recent cooperativeness, and then grouped agents such that players of similar rank played with each other. This ensured that highly cooperative agents played only with other cooperative agents, while defecting agents played only with other defectors. For Cleanup, cooperativeness was calculated based on the amount of steps in the last episode the agent chose to clean. For Harvest, it was calculated based on the difference between the the agent’s return and the mean return of all players, so that having less return than average yielded a high cooperativeness ranking. Cooperative metric-based matchmaking was only done with either individual reward networks or no reward networks (Figure 3(b)). We did not use cooperative metric-based matchmaking for our multi-level selection model, since these are theoretically separate approaches.

2.5 INDIVIDUAL VS. SHARED REWARD NETWORKS

Building on previous work that evolved either the intrinsic reward (Jaderberg et al., 2018) or the entire loss function (Houthoofd et al., 2018), we separately evolved the reward network within its own

population, thereby allowing different modules of the agent to compete only with like components. This allowed for independent exploration of hyperparameters via separate credit assignment of fitness and thus considerably more of the hyperparameter landscape could be explored compared with using only a single pool. In addition, reward networks could be randomly assigned to any policy network, and so were forced to generalize to a wide range of policies. In a given episode, 5 separate policy networks were paired with the same reward network, which we term a *shared reward network*. In line with (Jaderberg et al., 2017), the fitness determining the copying of policy network weights and evolution of optimization-related hyperparameters (entropy cost and learning rate) were based on individual agent return. By contrast, the reward network parameters were evolved according to fitness based on total episode return across the group of co-players (Figure 3(c)).

This contribution is distinct from previous work which evolved intrinsic rewards (e.g. Jaderberg et al., 2018) because (1) we evolve over social features rather than a remapping of environmental events, and (2) reward network evolution is motivated by dealing with the inherent tension in ISDs, rather than merely providing a denser reward signal. In this sense it’s closer to evolving a form of communication for social cooperation, rather than learning reward-shaping in a sparse-reward environment. We allow for multiple agents to share the same components, and as we shall see, in a social setting, this winds up being critical. Shared reward networks provide a biologically principled method that mixes between group fitness on a long timescale and individual reward on a short timescale. This contrasts with hand-crafted means of aggregation, as in previous work (Chang et al., 2004; Mataric, 1994).

3 RESULTS

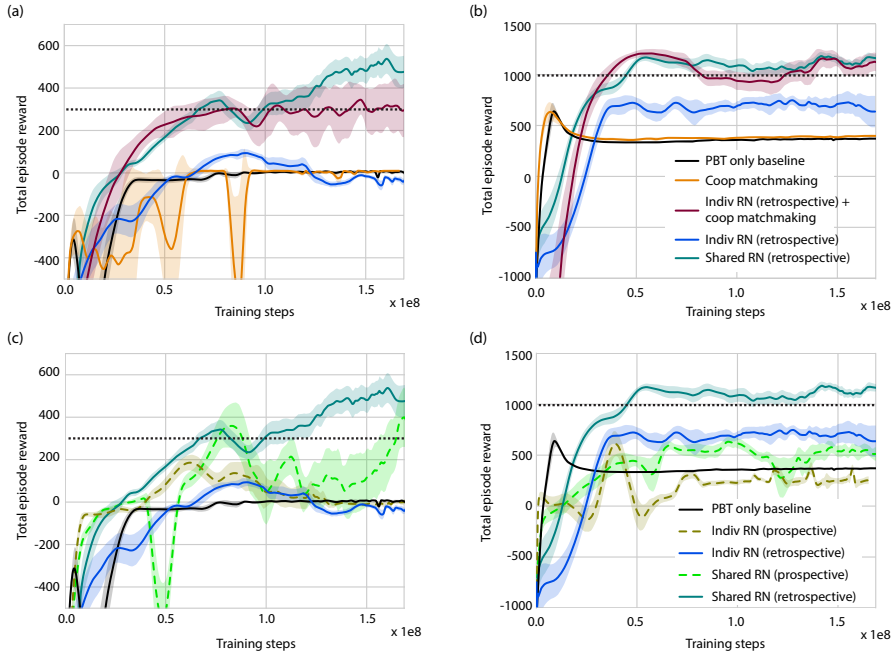


Figure 4: Total episode rewards, aggregated over players. (a), (b): Comparing retrospective (backward-looking) reward evolution with assortative matchmaking and PBT-only baseline in Cleanup (a) and Harvest (b). (c), (d): Comparing prospective (forward-looking) with retrospective (backward-looking) reward evolution in Cleanup (c) and Harvest (d). The black dotted line indicates performance from Hughes et al. (2018). The shaded region shows standard error of the mean, taken over the population of agents.

As shown in Figure 4, PBT without using an intrinsic reward network performs poorly on both games, where it asymptotes to 0 total episode reward in Cleanup and 400 for Harvest (the number of apples gained if all agents collect as quickly as they can).

Figures 4(a) and (b) compare random and assortative matchmaking with PBT and reward networks using retrospective social features. When using random matchmaking, individual reward network

agents perform no better than PBT on Cleanup, and only moderately better at Harvest. Hence there is little benefit to adding reward networks over social features if players have separate networks, evolved selfishly. The assortative matchmaking experiments used either no reward network ($u(f) = 0$) or individual reward networks. Without a reward network, performance was the same as the PBT baseline. With individual reward networks, performance was very high, indicating that both conditioning the internal rewards on social features and a preference for cooperative agents to play together were key to resolving the dilemma. On the other hand, shared reward network agents perform as well as assortative matchmaking and the handcrafted inequity aversion intrinsic reward from (Hughes et al., 2018), even using random matchmaking. This implies that agents didn’t necessarily need to have immediate access to honest signals of other agents’ cooperativeness to resolve the dilemma; it was enough to simply have the same intrinsic reward function, evolved according to collective episode return.

Figures 4(c) and (d) compare the retrospective and prospective variants of reward network evolution. The prospective variant, although better than PBT when using a shared reward network, generally results in worse performance and more instability. This is likely because the prospective variant depends on agents learning good value estimates before the reward networks become useful, whereas the retrospective variant only depends on environmentally provided reward and thus does not suffer from this issue.

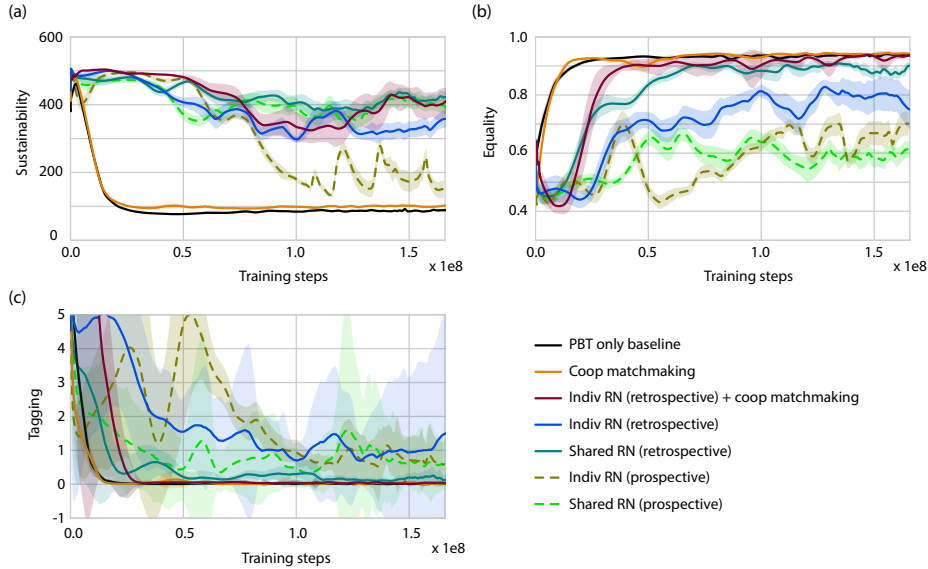


Figure 5: Social outcome metrics for Harvest. (a) Sustainability. (b) Equality. (c) Total amount of tagging. The shaded region shows the standard error of the mean.

We next plot various social outcome metrics in order to better capture the complexities of agent behavior (see Figure 5 for Harvest, see Appendix for Cleanup). Sustainability measures the average time step on which agents received positive reward, averaged over the episode and over agents. Figure 5(a) shows that having no reward network results in players collecting apples extremely quickly, compared with much more sustainable behavior with reward networks. Equality is calculated as $\mathbb{E}(1 - G(\mathbf{R}))$, where $G(\mathbf{R})$ is the Gini coefficient over individual returns. Figure 5(b) demonstrates that having the prospective version of reward networks tends to lead to lower equality, while retrospective variant has very high equality. Tagging measures the average number of times a player fined another player throughout the episode. Figure 5(c) shows that there is a higher propensity for tagging when using either a prospective reward network or an individual reward network, compared to the retrospective shared reward network. This explains the performance shown in Figure 4.

Finally, we can directly examine the weights of the final retrospective shared reward networks which were best at resolving the ISDs. Interestingly, the final weights evolved in the second layer suggest that resolving each game might require a different set of social preferences. In Cleanup, one of the final layer weights v_2 evolved to be close to 0, whereas in Harvest, v_1 and v_2 evolved to be of large magnitude but opposite sign. We can see a similar pattern with the biases \mathbf{b} . We interpret this to mean that Cleanup required a less complex reward network: it was enough to simply find other




Figure 6: Distribution of layer 2 weights and biases of evolved retrospective shared reward network at 1.5×10^8 training steps for (a) Cleanup, and (b) Harvest.

agents’ being rewarded as intrinsically rewarding. In Harvest, however, a more complex reward function was perhaps needed in order to ensure that other agents were not over-exploiting the apples. We found that the first layer weights \mathbf{W} tended to take on arbitrary (but positive) values. This is because of random matchmaking: co-players were randomly selected and thus there was little evolutionary pressure to specialize these weights.

4 DISCUSSION

Real environments don’t provide scalar reward signals to learn from. Instead, organisms have developed various internal drives based on either primary or secondary goals (Baldassarre & Mirolli, 2013). Here we examined intrinsic rewards based on features derived from other agents in the environment. In accord with evolutionary theory (Axelrod & Hamilton, 1981; Nowak, 2006), we found that naïvely implementing natural selection via genetic algorithms did not lead to the emergence of cooperation. Furthermore, assortative matchmaking was sufficient to generate cooperative behavior in cases where honest signals were available. Finally, we proposed a new multi-level evolutionary paradigm based on shared reward networks that achieves cooperation in more general situations.

Why does evolving intrinsic social preferences promote cooperation? Firstly, evolution ameliorates the intertemporal choice problem by distilling the long timescale of collective fitness into the short timescale of individual reinforcement learning, thereby improving credit assignment between selfish acts and their temporally displaced negative group outcomes (Hughes et al., 2018). Secondly, it mitigates the social dilemma itself by allowing evolution to expose social signals that correlate with, for example, an agent’s current level of selfishness. Such information powers a range of mechanisms for achieving mutual cooperation like competitive altruism (Hardy & Van Vugt, 2006), other-regarding preferences (Cooper & Kagel, 2016), and inequity aversion (Fehr & Schmidt, 1999). In accord, laboratory experiments show that humans cooperate more readily when they can communicate (Ostrom et al., 1992; Janssen et al., 2010).

The shared reward network evolution model was inspired by multi-level selection; yet it does not correspond to the prototypical case of that theory since its lower level units of evolution (the policy networks) are constantly swapping which higher level unit (reward network) they are paired with. Nevertheless, there are a variety of ways in which we see this form of modularity arise in nature. For example, free-living microorganisms occasionally form multi-cellular structures to solve a higher order adaptive problem, like slime mold forming a spore-producing stalk for dispersal (West et al., 2006), and many prokaryotes can incorporate plasmids (modules) found in their environment or received from other individuals as functional parts of their genome, thereby achieving cooperation in social dilemmas (Griffin et al., 2004; McGinty et al., 2011). Alternatively, in humans a reward network may represent a shared “cultural norm”, with its fitness based on cultural information accumulated from the groups in which it holds sway. In this way, the spread of norms can occur independently of the success of individual agents (Boyd & Richerson, 2009).

For future work, we suggest investigating alternative evolutionary mechanisms for the emergence of cooperation, such as kin selection (Griffin & West, 2002) and reciprocity (Trivers, 1971). It would be interesting to see whether these lead to different weights in a reward network, potentially hinting at the evolutionary origins of different social biases. Along these lines, one might consider studying an emergent version of the assortative matchmaking model along the lines suggested by Henrich (2004), adding further generality and power to our setup. Finally, it would be fascinating to determine how an evolutionary approach can be combined with multi-agent communication to produce that most paradoxical of cooperative behaviors: cheap talk.

REFERENCES

- Robert Axelrod and William D. Hamilton. The evolution of cooperation. *Science*, 211(4489):1390–1396, 1981. ISSN 00368075, 10959203. URL <http://www.jstor.org/stable/1685895>.
- Gianluca Baldassarre and Marco Mirolli. *Intrinsically Motivated Learning in Natural and Artificial Systems*. 01 2013.
- Robert Boyd and Peter J. Richerson. Culture and the evolution of human cooperation. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 364(1533):3281–3288, 2009. ISSN 0962-8436. doi: 10.1098/rstb.2009.0134. URL <http://rstb.royalsocietypublishing.org/content/364/1533/3281>.
- ZSH Chan, HW Ngan, AB Rad, and TK Ho. Alleviating ‘overfitting’ via genetically-regularised neural network. *Electronics Letters*, 38(15):1, 2002.
- Yu-Han Chang, Tracey Ho, and Leslie P Kaelbling. All learning is local: Multi-agent learning in global reward games. In *Advances in neural information processing systems*, pp. 807–814, 2004.
- Nuttapong Chentanez, Andrew G. Barto, and Satinder P. Singh. Intrinsically motivated reinforcement learning. In L. K. Saul, Y. Weiss, and L. Bottou (eds.), *Advances in Neural Information Processing Systems 17*, pp. 1281–1288. MIT Press, 2005. URL <http://papers.nips.cc/paper/2552-intrinsically-motivated-reinforcement-learning.pdf>.
- Edoardo Conti, Vashisht Madhavan, Felipe Petroski Such, Joel Lehman, Kenneth O Stanley, and Jeff Clune. Improving exploration in evolution strategies for deep reinforcement learning via a population of novelty-seeking agents. *arXiv preprint arXiv:1712.06560*, 2017.
- David J Cooper and John H Kagel. Other-regarding preferences. *The handbook of experimental economics*, 2:217, 2016.
- Charles Darwin. *On the Origin of Species by Means of Natural Selection*. Murray, London, 1859. or the Preservation of Favored Races in the Struggle for Life.
- Richard Dawkins. The selfish gene oxford university press. *New York*, 1976.
- Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Volodymir Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, et al. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. *arXiv preprint arXiv:1802.01561*, 2018.
- Ernst Fehr and Klaus M. Schmidt. A theory of fairness, competition, and cooperation*. *The Quarterly Journal of Economics*, 114(3):817–868, 1999. doi: 10.1162/003355399556151. URL <http://dx.doi.org/10.1162/003355399556151>.
- Jeffrey A Fletcher and Michael Doebeli. A simple and general explanation for the evolution of altruism. *Proceedings of the Royal Society of London B: Biological Sciences*, 276(1654):13–19, 2009.
- Jakob N. Foerster, Richard Y. Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. Learning with opponent-learning awareness. *CoRR*, abs/1709.04326, 2017. URL <http://arxiv.org/abs/1709.04326>.
- Andy Gardner and Stuart A West. Greenbeards. *Evolution: International Journal of Organic Evolution*, 64(1):25–38, 2010.
- Alan Grafen et al. Do animals really recognize kin?. *Animal Behaviour*, 39(1):42–54, 1990.
- Ashleigh Griffin and Stuart West. Kin selection: Fact and fiction. 17:15–21, 01 2002.
- Ashleigh S Griffin, Stuart A West, and Angus Buckling. Cooperation and competition in pathogenic bacteria. *Nature*, 430(7003):1024, 2004.
- Özgür Gürerk, Bernd Irlenbusch, and Bettina Rockenbach. The competitive advantage of sanctioning institutions. *Science*, 312(5770):108–111, 2006.

- W.D. Hamilton. The genetical evolution of social behaviour. i. *Journal of Theoretical Biology*, 7(1): 1–16, July 1964a. ISSN 0022-5193.
- William D Hamilton. The genetical evolution of social behaviour. ii. *Journal of theoretical biology*, 7(1):17–52, 1964b.
- Charlie L Hardy and Mark Van Vugt. Nice guys finish first: The competitive altruism hypothesis. *Personality and Social Psychology Bulletin*, 32(10):1402–1413, 2006.
- Joseph Henrich. Cultural group selection, coevolutionary processes and large-scale cooperation. *Journal of Economic Behavior & Organization*, 53(1):85–88, 2004.
- Geoffrey E Hinton and Steven J Nowlan. How learning can guide evolution. *Complex systems*, 1(3):495–502, 1987.
- Rein Houthooft, Richard Y. Chen, Phillip Isola, Bradley C. Stadie, Filip Wolski, Jonathan Ho, and Pieter Abbeel. Evolved policy gradients. *CoRR*, abs/1802.04821, 2018. URL <http://arxiv.org/abs/1802.04821>.
- Edward Hughes, Joel Z Leibo, Matthew G Phillips, Karl Tuyls, Edgar A Duéñez-Guzmán, Antonio García Castañeda, Iain Dunning, Tina Zhu, Kevin R McKee, Raphael Koster, et al. Inequity aversion improves cooperation in intertemporal social dilemmas. In *Advances in neural information processing systems (NIPS)*, Montreal, Canada, 2018.
- Max Jaderberg, Valentin Dalibard, Simon Osindero, Wojciech M Czarnecki, Jeff Donahue, Ali Razavi, Oriol Vinyals, Tim Green, Iain Dunning, Karen Simonyan, et al. Population based training of neural networks. *arXiv preprint arXiv:1711.09846*, 2017.
- Max Jaderberg, Wojciech M Czarnecki, Iain Dunning, Luke Marris, Guy Lever, Antonio Garcia Castaneda, Charles Beattie, Neil C Rabinowitz, Ari S Morcos, Avraham Ruderman, et al. Human-level performance in first-person multiplayer games with population-based deep reinforcement learning. *arXiv preprint arXiv:1807.01281*, 2018.
- Vincent Jansen and Minus van Baalen. Altruism through beard chromodynamics. 440:663–6, 04 2006.
- Marco A Janssen, Robert Holahan, Allen Lee, and Elinor Ostrom. Lab experiments for the study of social-ecological systems. *Science*, 328(5978):613–617, 2010.
- Laurent Keller and Kenneth G. Ross. Selfish genes: A green beard in the red fire ant. 394:573–575, 08 1998.
- Max Kleiman-Weiner, Mark K. Ho, Joseph L. Austerweil, Michael L. Littman, and Joshua B. Tenenbaum. Coordinate to cooperate or compete: Abstract goals and joint intentions in social interaction. In *CogSci*, 2016.
- Joel Z Leibo, Vinicius Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. 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.
- Adam Lerer and Alexander Peysakhovich. Maintaining cooperation in complex social dilemmas using deep reinforcement learning. *CoRR*, abs/1707.01068, 2017. URL <http://arxiv.org/abs/1707.01068>.
- Michael L Littman. Markov games as a framework for multi-agent reinforcement learning. In *Machine Learning Proceedings 1994*, pp. 157–163. Elsevier, 1994.
- Maja J Mataric. Learning to behave socially. In *Third international conference on simulation of adaptive behavior*, volume 617, pp. 453–462, 1994.
- John Maynard Smith and Eors Szathmary. *The major transitions in evolution*. Oxford University Press, 1997.

- Sorcha E McGinty, Daniel J Rankin, and Sam P Brown. Horizontal gene transfer and the evolution of bacterial cooperation. *Evolution: International Journal of Organic Evolution*, 65(1):21–32, 2011.
- Geoffrey F Miller, Peter M Todd, and Shailesh U Hegde. Designing neural networks using genetic algorithms. In *ICGA*, volume 89, pp. 379–384, 1989.
- Volodymyr Mnih, Adrià Puigdomènech Badia, Mehdi Mirza, Alex Graves, Timothy P. Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. *CoRR*, abs/1602.01783, 2016. URL <http://arxiv.org/abs/1602.01783>.
- Martin A. Nowak. Five rules for the evolution of cooperation. *Science*, 314(5805):1560–1563, 2006. ISSN 0036-8075. doi: 10.1126/science.1133755. URL <http://science.sciencemag.org/content/314/5805/1560>.
- Martin A Nowak and Karl Sigmund. Evolution of indirect reciprocity. *Nature*, 437(7063):1291, 2005.
- Pamela Oliver. Rewards and punishments as selective incentives for collective action: theoretical investigations. *American journal of sociology*, 85(6):1356–1375, 1980.
- Elinor Ostrom, James Walker, and Roy Gardner. Covenants with and without a sword: Self-governance is possible. *American political science Review*, 86(2):404–417, 1992.
- Liviu Panait and Sean Luke. Cooperative multi-agent learning: The state of the art. *Autonomous agents and multi-agent systems*, 11(3):387–434, 2005.
- Julien Pérolat, Joel Z. Leibo, Vinícius Flores Zambaldi, Charles Beattie, Karl Tuyls, and Thore Graepel. A multi-agent reinforcement learning model of common-pool resource appropriation. *CoRR*, abs/1707.06600, 2017. URL <http://arxiv.org/abs/1707.06600>.
- Alexander Peysakhovich and Adam Lerer. Consequentialist conditional cooperation in social dilemmas with imperfect information. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=BkabRiQpb>.
- Mitchell A Potter and Kenneth A De Jong. Cooperative coevolution: An architecture for evolving coadapted subcomponents. *Evolutionary computation*, 8(1):1–29, 2000.
- Francisco C Santos, Jorge M Pacheco, and Tom Lenaerts. Cooperation prevails when individuals adjust their social ties. *PLoS computational biology*, 2(10):e140, 2006.
- Satinder Singh, Richard L Lewis, and Andrew G Barto. Where do rewards come from. In *Proceedings of the annual conference of the cognitive science society*, pp. 2601–2606, 2009.
- J Maynard Smith. Group selection and kin selection. *Nature*, 201(4924):1145, 1964.
- Kenneth O Stanley and Risto Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2):99–127, 2002.
- Robert L Trivers. The evolution of reciprocal altruism. *The Quarterly review of biology*, 46(1): 35–57, 1971.
- Francisco Úbeda and Edgar A Duéñez-Guzmán. Power and corruption. *Evolution: International Journal of Organic Evolution*, 65(4):1127–1139, 2011.
- Stuart A West, Ashleigh S Griffin, Andy Gardner, and Stephen P Diggle. Social evolution theory for microorganisms. *Nature Reviews Microbiology*, 4(8):597, 2006.
- Daan Wierstra, Tom Schaul, Jan Peters, and Juergen Schmidhuber. Natural evolution strategies. In *Evolutionary Computation, 2008. CEC 2008.(IEEE World Congress on Computational Intelligence). IEEE Congress on*, pp. 3381–3387. IEEE, 2008.
- David Sloan Wilson. A theory of group selection. *Proceedings of the national academy of sciences*, 72(1):143–146, 1975.
- Chern Han Yong and Risto Miikkulainen. Cooperative coevolution of multi-agent systems. *University of Texas at Austin, Austin, TX*, 2001.

A ENVIRONMENT DETAILS AND ACTION SPACE

All episodes last 1000 steps, and the total size of the playable area is 25×18 for Cleanup and 36×16 for Harvest. Games are partially observable in that agents can only observe via a 15×15 RGB window, centered on their current location. The action space consists of moving left, right, up, and down, rotating left and right, and the ability to tag each other. This action has a reward cost of 1 to use, and causes the player tagged to lose 50 reward points, thus allowing for the possibility of punishing free-riders (Oliver, 1980; Gülerk et al., 2006). The Cleanup game has an additional action for cleaning waste.

B DETAILS OF TRAINING

Training was done via joint optimization of network parameters via SGD and hyperparameters/reward network parameters via evolution in the standard PBT setup. Gradient updates were applied for every trajectory up to a maximum length of 100 steps, using a batch size of 32. Optimization was via RMSProp with $\epsilon=10^{-5}$, momentum=0, decay rate=0.99, and an RL discount factor of 0.99. The baseline cost weight (see Mnih et al. (2016)) was fixed at 0.25, and the entropy cost was sampled from $\text{LogUniform}(2 \times 10^{-4}, 0.01)$ and evolved throughout training using PBT. The learning rates were all initially set to 4×10^{-4} and then allowed to evolve.

PBT uses evolution (specifically genetic algorithms) to search over a space of hyperparameters rather than manually tuning or performing a random search, resulting in an adaptive schedule of hyperparameters and joint optimization with network parameters learned through gradient descent Jaderberg et al. (2017).

There was a mutation rate of 0.1 when evolving hyperparameters, using either multiplicative perturbations of $\pm 20\%$ for entropy cost and learning rate, and additive perturbation of ± 0.1 for reward network parameters. We implemented a burn-in period for evolution of 4×10^6 agent steps, to allow network parameters and hyperparameters to be used in enough episodes for an accurate assessment of fitness before evolution.

C SUPPLEMENTAL RESULTS

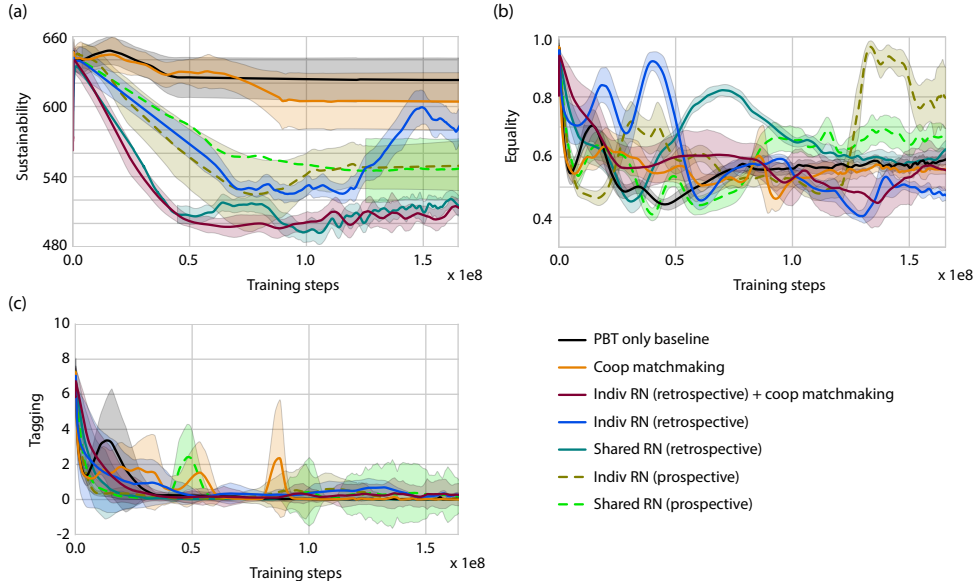


Figure 7: Social outcome metrics for Cleanup. (a) Sustainability. (b) Equality. (c) Total amount of tagging. Shaded region = SOM.