

Scaling Opponent Shaping to High Dimensional Games

Akhir Khan*
University College London
akbir.khan.13@ucl.ac.uk

Timon Willi*
University of Oxford
timon.willi@eng.ox.ac.uk

Newton Kwan*
University College London

Andrea Tacchetti
Deepmind

Chris Lu
University of Oxford

Edward Grefenstette
University College London

Tim Rocktäschel
University College London

Jakob Foerster
University of Oxford

ABSTRACT

In multi-agent settings with mixed incentives, methods developed for zero-sum games have been shown to lead to detrimental outcomes. To address this issue, *opponent shaping* (OS) methods explicitly learn to influence the learning dynamics of co-players and empirically lead to improved individual *and* collective outcomes. However, OS methods have only been evaluated in low-dimensional environments due to the challenges associated with estimating higher-order derivatives or scaling model-free meta-learning. Alternative methods that scale to more complex settings either converge to undesirable solutions or rely on unrealistic assumptions about the environment or co-players. In this paper, we successfully scale an OS-based approach to general-sum games with temporally-extended actions and long-time horizons for the first time. After analysing the representations of the meta-state and history used by previous algorithms, we propose a simplified version called SHAPER. We show empirically that SHAPER leads to improved individual and collective outcomes in a range of challenging settings from literature. We further formalize a technique previously implicit in the literature, and analyse its contribution to opponent shaping. We show empirically that this technique is helpful for the functioning of prior methods in certain environments. Lastly, we show that previous environments, such as the CoinGame, are inadequate for analysing temporally-extended general-sum interactions¹.

KEYWORDS

Multi-Agent Reinforcement Learning, Opponent Shaping, General-Sum Games

ACM Reference Format:

Akhir Khan, Timon Willi[1], Newton Kwan[1], Andrea Tacchetti, Chris Lu, Edward Grefenstette, Tim Rocktäschel, and Jakob Foerster. 2024. Scaling Opponent Shaping to High Dimensional Games. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 10 pages.

*Equal Contribution.

¹Blogpost available at sites.google.com/view/scale-os/



This work is licensed under a Creative Commons Attribution International 4.0 License.

1 INTRODUCTION

From personal assistants and chat-bots to self-driving cars and recommendation systems, the world of software is becoming increasingly *multi-agent* as these systems are continuously learning and interacting with each other in fully cooperative, fully competitive, and general-sum settings.

In this paper we investigate interacting, learning agents in general-sum settings. In such settings, commonly-used RL methods developed for zero-sum games can lead to disastrous outcomes [6]. For example, in real-world scenarios like pollution and international arms races [7, 34], such agents would fail to realise that they’re better off cooperating, even if it means they’re potentially worse off than their co-players. Poor performance could also lead to being extorted in social-dilemma-like scenarios [25, 28].

While multi-agent learning research has shown great success in strictly competitive [5, 15, 33, 36] and fully cooperative settings [11, 29], this success does not transfer to general-sum settings [10, 23]: In competitive games, agents can learn Nash equilibrium strategies by iteratively best-responding to suitable mixtures of past opponents. Similarly, best-responding to rational co-players leads to the desirable equilibria in cooperative games (assuming joint training). In contrast, many Nash equilibria coincide with globally worst welfare outcomes in general-sum settings, rendering the above learning paradigms ineffective. For example, in the iterated prisoner’s dilemma [IPD, 2, 13], naive best-response dynamics converge to unconditional mutual defection [10] rather than Nash equilibria with higher social welfare.

It is important that general-sum learning methods scale to *high-dimensional* settings, such as those with longer-time horizons, larger state spaces and temporally-extended actions, as these environments are more akin to the real world. In matrix games, cooperation and defection are clearly defined atomic actions, whilst in more complex environments such as autonomous driving, cooperation and defection are defined over sequences of actions (e.g. a path towards a cooperative/defective location). Scalable methods [16, 19, 27, 40] that manage to avoid *unconditional defection* in these settings rely heavily on *reward shaping*, which blurs the line between the problem setting (“social dilemma”) and the method.

As an alternative approach, *opponent shaping* (OS) methods recognise that the actions of any one agent influence their co-player’s policy and seek to use this mechanism to their advantage [10, 17, 22, 38]. However, many past OS methods require privileged

information to shape their opponents and are myopic since anticipating many steps is intractable. Model-Free Opponent Shaping [M-FOS, 25] and The Good Shepherd [GS, 4] solve the issues above by framing opponent shaping as a meta-learning problem, which our method inherits and builds upon. However, M-FOS presents only preliminary results on the higher-dimensional CoinGame [21] benchmark, and GS none at all.

To scale OS agents to higher-dimensional benchmarks, we systematically evaluate the architectural components of the M-FOS and GS algorithms. We identify two forms of memory—*history* and *context*. *History* captures intra-episode information and *context* inter-episode information. We find empirically that both are necessary to achieve shaping. M-FOS captures both types of memory (though not completely), whereas GS does not. However, we identify a bottleneck in the M-FOS method, as M-FOS requires two separate policies to capture *context* and *history*, where only one is necessary. Using this finding, we propose a new method, called SHAPER, removing the unnecessary bottleneck from M-FOS.

Beyond these memory components, we uncover another element used implicitly in prior work but never formally introduced or analysed: averaging across the batch of trajectories at each meta step. This ensures that the hidden states of the opponent shaping algorithm carry information from the entire batch, rather than just a single batch dimension. We formalise this technique and empirically investigate its importance. Our analysis shows that while this technique improves previous methods like M-FOS in certain environments, it is not essential for our proposed method, SHAPER, in typical environment settings. This highlights the value of our formalisation and empirical analysis in understanding and improving upon existing methods.

SHAPER outperforms previous OS methods in general-sum games with long-time horizons and temporally-extended actions. We showcase our performance on the “IPD in the Matrix” and “IMP in the Matrix” games, introduced by the *melting pot* suite [20]. These have a 30x state-space than environments previously used to evaluate OS and contain more complex interaction dynamics. Additionally, we consider shaping on sequential matrix games with 100x longer horizons than their previously used counterparts. We demonstrate empirically that our simplification of M-FOS helps scalability, that only *evolutionary-based* meta-learning is effective in these long-horizon games, and that previous evaluation environments, such as the CoinGame [21], are inadequate for analysing OS in temporally-extended, general-sum interactions.

2 BACKGROUND

What is a Game? We formalise our environments as Partially Observable Stochastic Games [32, POSG]. A POSG is given by the tuple $\mathcal{M} = \langle N, \mathcal{A}, \mathcal{O}, S, \mathcal{T}, \mathcal{I}, \mathcal{R}, \gamma \rangle$, where \mathcal{A} , \mathcal{O} , and S denote the action, observation, and state space, respectively. These parameters can be distinct at every time step and also incorporated into the transition function $\mathcal{T} : S \times \mathbf{A} \rightarrow \Delta S$, where $\mathbf{A} \equiv \mathcal{A}^n$ is the joint action of all agents. Each agent draws individual observations according to the observation function $\mathcal{I} : S \times N \rightarrow \mathcal{O}$ and obtains a reward according to their reward function $\mathcal{R} : S \times \mathbf{A} \times N \rightarrow \mathbb{R}$ where $N = \{1, \dots, n\}$. POSGs represent general-sum games. The

Algorithm 1 SHAPER Update: Given a POSG \mathcal{M} , policies $\pi_{\phi_i}, \pi_{\phi_{-i}}$ and their respective initial hidden states h_i, h_{-i} and a distribution of initial co-players ρ_{ϕ} , this algorithm updates a meta-agent policy ϕ_i over T trials consisting of E episodes.

Require: $\mathcal{M}, \phi_i, \rho_{\phi}, E, T$

```

1: for  $t = 0$  to  $T$  do
2:   Initialise trial reward  $\bar{J} = 0$ 
3:   Initialise meta-agent hidden state  $h_i = \mathbf{0}$ 
4:   Sample co-players  $\phi_{-i} \sim \rho_{\phi}$ 
5:   for  $e = 0$  to  $E$  do
6:     Initialise co-players'  $h_{-i} = \mathbf{0}$ 
7:      $J_i, J_{-i}, h'_i, h'_{-i} = \mathcal{M}(\phi_i, \phi_{-i}, h_i, h_{-i})$ 
8:     Update  $\phi_{-i}$  according to co-players' update rule.
9:      $h_i \leftarrow h'_i$ 
10:     $\bar{J} \leftarrow \bar{J} + J_i$ 
11:   end for
12:   Update  $\phi_i$  with respect to  $\bar{J}$ 
13: end for

```

single-player case, $N = \{1\}$, of POSGs are Partially Observable Markov Decision Processes (POMDPs).

What is Shaping? Shaping is acting to manipulate the co-player’s learning dynamics (and subsequent behaviour) [10], where co-players are any other participants in the game. Newer shaping methods frame shaping as a meta-learning problem [4, 17, 25]. We next present the meta-learning problem setting presented by M-FOS since our work is a simplified case of the M-FOS framework.

What is M-FOS? Conceptually, M-FOS separates the task of shaping (the meta-game) from the task of playing the game. Specifically, the meta-game is formulated as a POMDP $\langle \bar{\mathcal{A}}, \bar{\mathcal{O}}, \bar{S}, \bar{\mathcal{T}}, \bar{\mathcal{I}}, \bar{\mathcal{R}}, \bar{\gamma} \rangle$ over an underlying general-sum game, represented by a POSG \mathcal{M} , where the overline indicates the single-agent version of the elements defined for POSGs. In the “shaping” POMDP, the meta-state \bar{S} contains the policies of all players in the underlying POSG: $\bar{s}_e = (\phi_i^{e-1}, \phi_{-i}^{e-1}) \in \bar{S}$, where e indexes the episodes and $(i, -i)$ index all agents. The meta-observation is all observations of the previous episode in the underlying game, i.e., $\bar{o}_e = (o_0^{e-1}, o_1^{e-1}, \dots, o_K^{e-1})$, where K is the length of an episode. The meta-action space $\bar{\mathcal{A}}$ consists of the policy parameterisation of the inner agent i (in practice a vector conditioning the policy), i.e., $\bar{a}_e = \phi_i^e$.

M-FOS training works as follows. The meta-agent trains over a sequence of T trials (denoted “meta-episodes” in the original paper). Each trial contains E episodes. At the end of each episode e within a trial t , conditioned on both agent’s policies, the co-players update their parameters with respect to the expected episodic return $J_{-i}^e = \mathbb{E} \left[\sum_{k=0}^K \gamma^k r_{-i}^k(\phi_i^e, \phi_{-i}^e) \right]$, where K is the length of an episode. For example, if the co-players were Naive Learners, i.e., agents not accounting for the learning dynamics of the co-player, with learning rate α , the update is: $\phi_j^{e+1} = \phi_j^e + \alpha \nabla_{\phi_j} J_j^e$, for $j \in -i$.

In contrast to the co-player’s update, the meta-agent learns an update function for the parameters of their inner agent, i.e., $\bar{a}_e = \phi_i^e \sim \pi_{\theta}(\cdot | \bar{o}_e)$, where θ is the parameters of the meta-agent. The meta-agent optimises the meta-return $\bar{J} = \sum_{e=0}^E J_i^e$ (summed

over all episodes) at the end of a trial t using any model-free optimisation technique, e.g., PPO [31] or Evolution Strategies [ES, 30]. The meta-agent and the inner agent are usually represented as recurrent neural networks, such as LSTMs [14]. The meta-game setup allows the meta-agent to observe the results of the co-player’s learning dynamics, enabling it *to learn to shape*. Though not formalized or discussed in detail in the paper, the original M-FOS averages across a batch of trajectories at each update step to ensure access to all information necessary for shaping. In Section 3, we introduce a formal definition of averaging across the batch and investigate its role for shaping. While published concurrently, the Good Shepherd [4] is a simplified version of M-FOS, in which the meta-agent and underlying agent are collapsed into a single agent without memory *that only updates after each trial*. The agent is represented by a feedforward neural network and has no memory. However, as GS was evaluated on infinitely iterated matrix games, where the state usually represents a one-step history, we consider GS to have one-step history.

Where have current shaping algorithms been evaluated?

Both M-FOS and GS evaluate their shaping on infinitely-iterated matrix games. While this is a fruitful playground to discover complex strategies, such as tit-for-tat, infinitely-iterated matrix games do not contain temporally-extended actions or high dimensional state spaces. For example, in matrix games, cooperation simply consists of playing the “cooperation” action. However, in the real world, cooperation requires a repeated commitment to a cooperative strategy (where it is often unclear whether a given atomic action is cooperative). The CoinGame [21] supposedly addresses this shortcoming by incorporating IPD-like game dynamics into a gridworld. M-FOS presents very preliminary results on the CoinGame with no detailed analysis of the emerging strategies. However, as we show in Section 4 the CoinGame suffers from pathologies that enable shaping with simple strategies.

3 SHAPER: A SCALABLE OS METHOD

To introduce our method, we first analyse the role of memory in meta-learning-based OS. Memory is important because it enables a meta-agent to adapt its meta-policy within a trial since it only updates *parameters* after a trial.

If the meta-agent cannot adapt their policy within a trial, a co-player could simply learn the best-response to the meta-agent’s policy. For example, in Rock-Paper-Scissors, the meta-agent would be forced to play the fully mixed strategy, as any deviation from it will be taken advantage of by the co-player. Instead a meta-agent with memory can adapt to the co-player’s best response within a trial and potentially achieve a better meta-return, which we show in Section 5. Thus, memory is important if a meta-agent is to perform well in all general-sum games.

We define two forms of memory: *context* and *history*. Let us define *history* as a trajectory of a (partial) episode e , $\tau_e = (o_e^0, a_e^0, r_e^0, \dots, r_e^K)$, and *context* as a trajectory of a (partial) trial t , $\bar{\tau}_t = (\tau_0, \dots, \tau_E)$. *History* captures the dynamics within an episode and is crucial for implementing policies such as TFT that reward/punish the co-player based on past actions. In contrast, *context* captures the learning dynamics of the co-player as it contains the co-player’s policy changes over many parameter updates. *Context* is important for shaping

when the co-players *update dynamics* are non-stationary across a trial or need to be inferred from the changing policy itself across different episodes. It allows the shaper to adapt its inner policy to, e.g., a change of the co-player’s learning rate or implicitly infer their objective function.

Using the above definitions, we express the policies as the following: M-FOS : $a \sim \pi_\phi(\cdot \mid \tau_e, \pi_\theta(\bar{o}_e))$ and GS : $a \sim \pi_\phi(\cdot \mid o_e^t)$. M-FOS captures one-step *context* via the memory of the meta-agent, and *history* via the memory of the inner agent but requires two agents to do so. In contrast, GS captures one-step history (if given by environment) but does not require a separate inner agent.

We propose SHAPER, an algorithm requiring only one agent to capture *context and history*. This is accomplished via an RNN that retains its hidden state over episodes and only resets *after each trial*

$$\text{SHAPER: } a \sim \pi_\phi(\cdot \mid \tau_e, \bar{\tau}_t) \quad (1)$$

Compared to GS, SHAPER has access to history and context by *adding memory to the architecture and retaining the hidden state over the episodes*. Compared to M-FOS, SHAPER only requires sampling from one action space. To contrast SHAPER to M-FOS in more detail, we refer to Appendix.

SHAPER is trained as follows. Given a POSG \mathcal{M} , at the start of a trial, co-players $\phi_{-i} \sim \rho_\phi$ are sampled, where ρ_ϕ is the respective sampling distribution. SHAPER’s parameters ϕ_i and hidden state h_i are randomly initialised. During an episode of length K , agents take their actions, $a_i^k \sim \pi_{\phi_i}(\cdot \mid o_i^k, h_i^k)$. At each time step in the episode, the hidden state of the meta-agent is updated: $h_i^{k+1} = f(o_i^k, h_i^k)$. On receiving actions, the POSG returns rewards r_i^k , new observations o_i^{k+1} and a done flag d , indicating if an episode has ended.

When an inner episode terminates, the updated co-player ϕ_{-i}^{e+1} and the meta-agent’s hidden state h_i^K are passed to the next episode. This process is repeated over E episodes in a trial. When a trial terminates, the meta-agent’s policy is updated, maximising total trial reward, $\bar{J} = \sum_e^E J_i^e$. This leads to the following objective,

$$\max_{\phi_i} \mathbb{E}_{\rho(\phi), \rho(\mathcal{M})} [\bar{J}]. \quad (2)$$

In practice, the co-players optimise their parameters using some form of gradient descent, which typically involves batching the episodic trajectories. Assume $\phi_{-i}^e = G(\phi_{-i}^{e-1}, \tau_{e-1})$ is some co-player’s update function $G : \mathbb{R}^P \times \mathbb{R}^{B \times T} \rightarrow \mathbb{R}^P$, where P is the number of parameters of ϕ_{-i} , B is the batch size, i.e., number of environments for parallel training, and T is the length of the trajectory. Shaper then interacts with a co-player over a batch of environments, i.e., $\mathbf{a}_i^k \sim \pi_{\phi_i}(\cdot \mid \mathbf{o}_i^k, \mathbf{h}_i^k)$, where $\mathbf{a}_i^k \in \mathbb{R}^{B \times A}$, $\mathbf{o}_i^k \in \mathbb{R}^{B \times O}$, and $\mathbf{h}_i^k \in \mathbb{R}^{B \times H}$, and A , O , and H are the action-, observation-, and hidden-state-size respectively.

Shaper needs to account for the batched updates of the co-player because opponent shaping requires all the information determining the learning update of the co-player. For example, imagine Shaper plays with a co-player across a batch of environments with different reward functions. While the co-player updates its parameters based on a diverse set of trajectories from many reward functions, each of Shaper’s hidden states only observes the trajectory of its respective reward function. Intuitively, if the reward functions are very diverse, the update derived from the whole batch would significantly differ

Table 1: Converged reward per episode (meta-agent, co-player) for agents trained with Naive Learners. We report reward per episode for better interpretability. Mean and standard deviation reported across 100 seeds.

	CoinGame		IPD in the Matrix		IMP in the Matrix	
SHAPER	4.63 ± 0.66 ,	-3.35 ± 0.67	22.44 ± 1.12 ,	21.49 ± 0.67	0.14 ± 0.06 ,	-0.14 ± 0.06
M-FOS (ES)	3.13 ± 0.40 ,	2.27 ± 0.38	15.49 ± 1.28 ,	23.88 ± 0.93	0.11 ± 0.07 ,	-0.11 ± 0.07
M-FOS (RL)	0.94 ± 0.68 ,	-0.23 ± 0.52	7.42 ± 0.21 ,	7.28 ± 0.15	0.04 ± 0.00 ,	-0.04 ± 0.00
GS	5.44 ± 0.61 ,	-4.17 ± 0.48	16.16 ± 1.33 ,	24.35 ± 0.83	0.00 ± 0.00 ,	0.00 ± 0.00
PT-NL	0.70 ± 0.58 ,	-0.3 ± 0.47	6.33 ± 0.33 ,	6.96 ± 0.31	-0.17 ± 0.10 ,	0.17 ± 0.10
CT-NL	0.47 ± 0.83 ,	0.26 ± 0.30	5.56 ± 0.02 ,	5.56 ± 0.02	-0.10 ± 0.06 ,	0.10 ± 0.06

from the update estimated from a single environment trajectory, as observed by the hidden state. We thus define the *batched context* as a trajectory of a batched (partial) trial $\bar{\tau}_t = (\tau_0, \dots, \tau_E)$.

$$\text{Shaper: } a \sim \pi_\phi(\cdot \mid \tau_e, \bar{\tau}_t)$$

This insight leads to the consequence that Shaper needs to consolidate information across its batch of hidden states, at least at every co-player update. To address this issue, Shaper averages over its hidden states across the batch at each step, combined with a skip connection to ensure “situational” awareness of the hidden state’s respective environment (see Appendix).

$$\mathbf{h}_i^{k+1} = \lambda \left(\frac{1}{B} \sum_{l=0}^B h_{i,l}^k \right) + (1 - \lambda) \mathbf{h}_i^k \quad (3)$$

$$\mathbf{h}_i^{k+1} = f(o_i^k, \mathbf{h}_i^{k+1}) \quad (4)$$

This approach ensures that SHAPER effectively shapes its co-players while accounting for the diverse set of trajectories that inform their gradient updates, captured by the following meta-return function with the expectation over the batched gradient update of the co-player:

$$\begin{aligned} \bar{J}^{\text{ES}} = \mathbb{E}_{e \sim \mathcal{N}(\mathbf{0}, I_d)} & \left[\mathbb{E}_{\bar{\tau} \sim \pi_{\phi_i + \epsilon \sigma}; \phi_{-i}^e \sim G(\phi_{-i}^{e-1}, \tau_{e-1})} \right. \\ & \left. \times \left[\mathbb{E}_{\tau_e \sim \pi_{\phi_{-i}^e}} \left[\sum_{k=0}^K \gamma^k r^k(\phi_{-i}^e, \phi_i + \epsilon \sigma) \right] \right] \right] \quad (5) \end{aligned}$$

4 EXPERIMENTS

Here we present the test environments and our evaluation protocol for SHAPER. We also explain our ablation experiments helping us evaluate the role of memory in OS.

The **Prisoner’s Dilemma** is a well-known and widely studied general-sum game illustrating that two self-interested agents do not cooperate even if it is globally optimal. The players either cooperate (C) or defect (D) and receive a payoff (see Appendix). In the *iterated prisoner’s dilemma* (IPD), the agents repeatedly play the prisoner’s dilemma and observe the previous action of both players. Past research used the infinite IPD in their experiments [4, 10, 23, 25, 38]. In the infinite version, the exact value function and gradients thereof are calculated directly from the policy weights [10]; In our work, we consider the finitely iterated PD (f-IPD), where we cannot calculate the exact value function and have to rely on sample-based approaches such as RL and ES.

Iterated Matching Pennies (IMP) is an iterated matrix game like the IPD. The players choose heads (H) or tails (T) and receive a payoff according to both players’ choices. In contrast to the IPD, a general-sum game, IMP is a zero-sum game. In the IMP one player gets +1 for playing the same action as the other player, while the other player is rewarded for playing a *different* action. Thus, the only equilibrium strategy for each one-memory agent is to play a random policy, resulting in an expected joint payoff of (0,0). Only with intra-episode memory can a meta-agent observe a co-player’s current policy and thus shape it.

CoinGame is a multi-agent gridworld environment that simulates social dilemmas (like the IPD) with high-dimensional states and multi-step actions [21]. Two players, blue and orange, move around a wrap-around grid and pick up blue and orange coloured coins. When a player picks up a coin of any colour, this player receives a reward of +1. When a player picks up a coin of the co-player’s colour, the co-player receives a reward of −2. Whenever a coin gets picked up, a new coin of the same colour appears in a random location on the grid at the next time step. If both agents reach a coin simultaneously, then both agents pick up that coin (the coin is duplicated). When both players pick up coins without regard to colour, the expected reward is 0. In contrast to matrix games, the CoinGame requires learning from high-dimensional states with multi-step actions.

*** in the Matrix** extends matrix games to gridworld environments [35], where * is any normal-form game. For visual descriptions, see Appendix. Agents collect two types of resources into their inventory: *Cooperate* and *Defect* coins. Once an agent has collected any coin, the agent’s colour changes, representing that the agent is “ready” for interaction. Agents can fire an ‘interact’ beam to an area in front of them. If an agent’s interact beam catches a “ready” agent, both receive rewards equivalent to playing a matrix game *, where their inventory represents their policy. For example, when agent 1’s inventory is 1 *Cooperate* coin and 3 *Defect* coins, agent 1’s probability to cooperate is 25%. For all details, see the Appendix.

* in the Matrix introduces a series of novel complexities for shaping over the CoinGame and finite matrix games. The environment is substantially more demanding than the previous games—it is partially observable, has complex interactions, and much longer time horizons. For shaping, partial observability makes temporally-extended actions harder to estimate. Shapers are also incentivised to speed up co-players learning, as the environment only allows interactions after both agents have picked up a coin. We explore two




Figure 1: Evaluation results over a single trial (with co-player) compromising over 100 seeds for the CoinGame. (a) Reward, (b) SHAPER’s frequency of picking up its own colour coin, (c) state visitation, and (d) the number of coins picked up per episode. SHAPER successfully elicits exploitation with a co-player with a high state visitation for DC and strong competency.




Figure 2: Render of the IPDitM games, a multi-step, gridworld-based general-sum game. Agents with restricted visibility and orientation traverse a grid picking up either Defect or Cooperate coins. (left) shows an initial state of the game before either agent has a coin. Once agents pick up a coin, their appearance changes, and they can interact. (right) shows the orange agent having collected a coin and the blue agent firing their interact beam.

specific implementations of the game: “IPD in the Matrix” (IPDitM) and “IMP in the Matrix” (IMPitM).

For our baseline comparisons, we compare SHAPER against multiple baselines: Naive Learners (NLs), variants of M-FOS, and GS. A NL does not explicitly account for the learning of the co-player across different episodes. In all of our experiments, the co-player is a NL. We train meta-agents until convergence in their respective environments. Then, we evaluate the performance of fixed meta-agents against new co-player (NL) initialisation ϕ_{-i} . Additional implementation details and hyperparameters for each game are provided in the Appendix.²

In finite matrix games, our NL is parameterised as a tabular policy trained using PPO. In the gridworld environments, the NL is parameterised by a recurrent neural network and trained using PPO. Furthermore, in gridworlds, we compare to both M-FOS optimised with PPO and by ES. For GS, we only use ES, consistent with the original paper. We compare the performance of SHAPER to two different types of NL pairs: The first type, co-training NL (CT-NL), two NLs are initialised randomly and trained together

²The codebase is open-source [37].

using independent learning. This shows that avoiding unconditional defection is a challenge in the first place. The second type, pre-trained NL (PT-NL) instead takes an agent from a fully trained CT-NL pair and uses it as a naive shaper baseline, i.e., trains a NL as a best response to the fixed final policy. This ensures that the performance of SHAPER is not simply due to breaking the learning dynamics of the co-player, e.g., because the fully trained NL deprives a randomly initialised agent of all rewards. Specific details are provided in Appendix. Next, for our **ablations**, we consider three challenges:

Context Challenge: During a trial, after k episodes, the co-player stops updating their parameters. When they stop updating, the shaper’s optimal behaviour is to exploit the co-player’s fixed policy (effectively stop shaping). We evaluate in the IPD and choose $k = 2$. This challenge tests if shapers: 1) identify the sudden change in a co-player’s learning dynamics, and 2) deploy a more suitable exploitative policy. We hypothesise that shapers without context cannot identify the change. We evaluate SHAPER and compare against GS to understand the importance of context for shaping.

History Challenge: We reset the hidden state of SHAPER between episodes, removing its ability to use *context* to shape (SHAPER w/o context). We evaluate in IMP, and agents must infer the co-player’s current policy using only history. Finally, we evaluate SHAPER within IMP environment over short and long episode lengths (2 and 100, respectively) to limit the relative strength of *history*.

Average Challenge: We also analyse the role of averaging across the batch in matrix games by comparing the performance difference of both MFOS and Shaper with and without averaging.

5 RESULTS

Shaping in Finite Matrix Games: We evaluate SHAPER, M-FOS and GS on finite matrix games, i.e., long-time-horizon variants of the infinite matrix games used in prior work. We recreate previously reported extortion behaviour in a more challenging setting [25].

Insight 1: SHAPER shapes the best in long-horizon iterated matrix games. We inspect the converged reward for each shaping algorithm against a PPO agent in the IPD (see the Appendix). Here, SHAPER shapes its co-player more effectively than the baselines, achieving an average return of -0.13 per episode. All shaping baselines reach extortion-like policies.

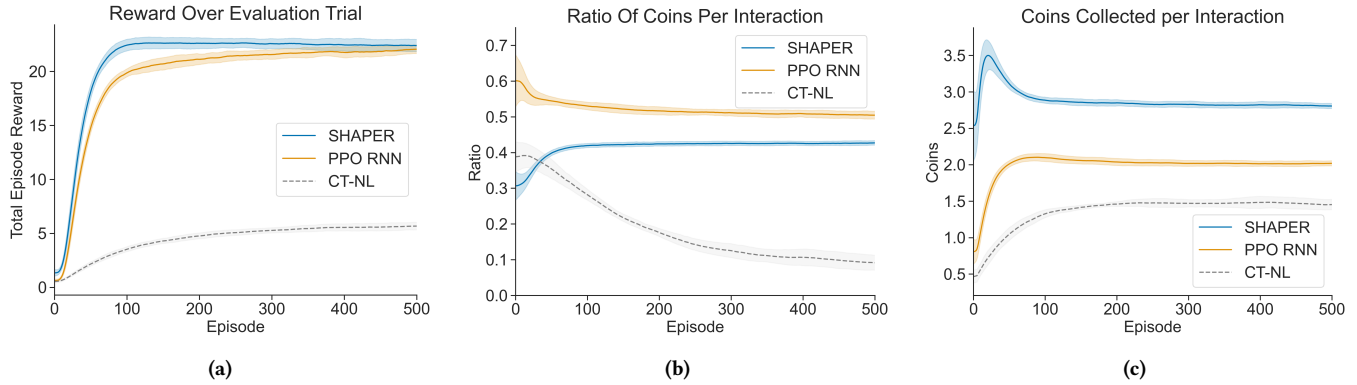


Figure 3: Evaluation results over a single trial (with co-player) comprising over 100 seeds for the IPDitM. (a) Mean reward per timestep, (b) mean ratio of picking up cooperate coins per soft-reset, (c) total number of coins picked up per soft-reset. The independent learner is shown to contrast what learning without a meta-agent would look like.

Table 2: Converged reward per step (meta-agent, co-player) for agents against Naive Learners in finite matrix games. SHAPER shapes co-players to exploitative equilibria. Mean and standard deviation over 20 randomised co-players.

	IPD	IMP
SHAPER	$-0.1 \pm 0.02, -2.8 \pm 0.05$	$0.9 \pm 0.02, -0.9 \pm 0.02$
M-FOS	$-0.6 \pm 0.14, -2.3 \pm 0.14$	$0.8 \pm 0.09, -0.8 \pm 0.09$
GS	$-1.0 \pm 0.03, -1.3 \pm 0.10$	$0.0 \pm 0.01, 0.0 \pm 0.01$
CT-NL	$-2.0 \pm 0.00, -2.0 \pm 0.00$	$0.0 \pm 0.00, 0.0 \pm 0.00$

Insight 2: Memory is important for shaping in the IMP.

In the IMP, SHAPER exploits its opponent to achieve a score of (0.9, -0.9) (see Table 2). As expected, GS cannot shape the opponent, achieving a score close to the Nash equilibrium, (0.0, 0.0). With only a single-step history, it is impossible to shape the opponent because the opponent can switch to a random strategy between episodes to achieve a score of at least 0. Thus memory is required to find shaping strategies. We find that M-FOS, an agent with memory, shapes too. Next, we present our **CoinGame** results.

Insight 3: Knowing how to navigate the gridworld and pick up coins is already enough to suppress co-player’s learning. Towards the end of meta-training, newly initialised co-players have to play against already competent meta-agents who have seen the game many times. We found that in CoinGame, it was sufficient for the meta-agents to pick up all coins before the co-player could reach them to hinder training. Therefore, we suggest checking that the co-player learns against pre-trained Naive Learners. This mitigates behaviours that prohibit the co-players from learning at all. We found that changing from a global to an egocentric observation space in the CoinGame helped co-players learn against pre-trained agents. Examples of sanity tests are found in the Appendix.

Reiterating *Insight 1*, we find meta-agents find extortion-like policies in the CoinGame. To better understand behaviour in CoinGame, we extend the five states from the IPD (S, CC, CD, DC, DD) to include the start states (SS, SC, SD, CS, DS). At the start of an episode, the state is SS until a player picks up a coin. To understand how

SHAPER shapes, we inspect the probability of the meta-agent picking up a coin of its own colour at the start, i.e., SC \rightarrow CC. For example, suppose the meta-agent were to cooperate unconditionally in the CoinGame. In that case, it only picks up coins of their own colour no matter the state and would relate to a high probability of cooperating over all states.

Figure 1b demonstrates how SHAPER shapes its co-players effectively already at the start. The difference between cooperating in SC and SS (25%, 15% resp.) highlights how SHAPER uses context to evaluate the exploitability of its co-player. In SS, when both agents have not picked up coins, SHAPER probes for exploitability by not cooperating. In SC, where the co-player has already shown they are cooperative, SHAPER also cooperates. Moreover, Figure 1c shows that CS is visited more often than DS in early episodes (18%, 15% resp.), indicating that SHAPER is shaping the co-player to form a preference for picking up their own colour. This preference is then exploited by SHAPER as indicated by the increasing visitation of DC. The meta-agent’s probability of cooperating in DC converges to 25%, i.e., occasionally rewarding the co-player, as never cooperating would probably make the co-player learn pure defection.

Insight 4: CoinGame is not suitable as a multi-step action environment. We found GS produces comparable results to SHAPER. At first, this is surprising since GS is a feedforward network and does not have access to the history (or, at most, one step). Therefore it should not be able to retaliate against a defecting agent since it has no memory of their past actions. However, a close investigation of the problem setting shows that due to particular environment dynamics, the *current state* is often indicative of *past actions*. For example, seeing two agents and a coin on the same square is a strong signal that one of the agents defected since this situation only could have arisen when either all objects spawn on the same square (occurs with a probability of 0.12% and only at the beginning of an episode) or when both agents went for the same coin and the coin respawned on top of them (see the Appendix). This illustrates that CoinGame allows for simple shaping strategies that do not require *context* or *history*, limiting its utility as a benchmark to measure temporally-extended actions.

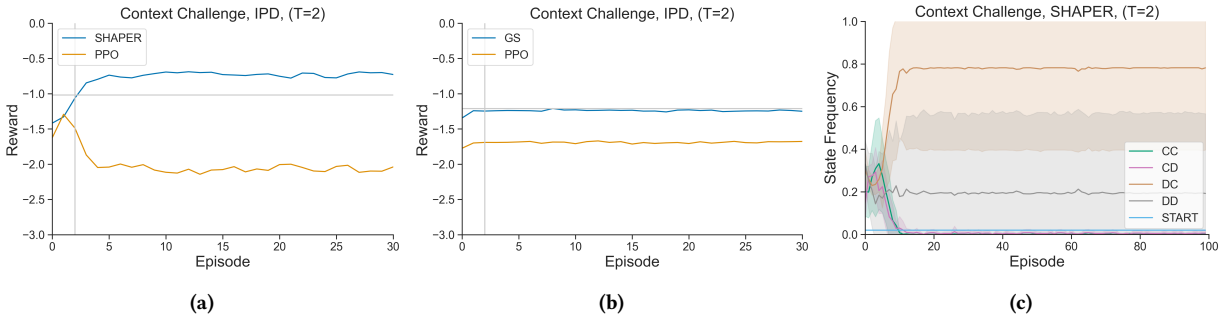


Figure 4: Hardstop Challenge: Average reward per timestep over an evaluation trial for SHAPER (a) and GS (b) against a Naive Learner in the IPD. Here GS fails to generalise to a co-player that stops learning after an unknown number of timesteps (unseen during training). (c) State Visitation through the evaluation shows SHAPER responds to co-players frozen policy by moving into either DD (the best response to a defective agent) or DC (the best response to a fully cooperative agent).

Table 3: Ablations highlighting the importance of context and history for Shaping. We report converged reward per step (meta-agent, co-player) for agents against Naive Learners.

Context Challenge: IPD	
SHAPER	-0.8, -2.0
SHAPER w/o Context	-1.25, -1.75
History Challenge: IMP (Length=2)	
SHAPER	0.5, -0.5
SHAPER w/o History	0.0, 0.0
History Challenge: IMP (Length=100)	
SHAPER	0.5, -0.5
SHAPER w/o History	0.5, -0.5

We continue with our results for the * in the **Matrix** environments. Motivated by *Insight 3*, we show that co-players learn against pre-trained agents by the number of coins collected in the Appendix.

Insight 5: SHAPER outperforms other shaping methods in the IPDitM by a considerable margin SHAPER outperforms other shaping methods in the IPDitM by a considerable margin (see Table 1), e.g., Shaper gets ~ 22.44 points against NL, where M-FOS gets ~ 15.49 . Furthermore, SHAPER finds a *collectively better equilibrium for both players* over any other shaping method, e.g., in comparison with M-FOS, Shaper achieves ($\sim 22.44, \sim 21.49$) and M-FOS gets ($\sim 15.49, \sim 23.88$).

Insight 6: Shaping in IPDitM leads to collectively and individually better outcomes. Table 1 (second column) shows that shaping (SHAPER, M-FOS, and GS) leads to collectively and individually better outcomes in IPDitM compared to PT-NL or CT-NL.

Insight 7: SHAPER shapes by picking up almost all coins at the beginning of a trial. The meta-agent picks up almost all coins in the grid in the first 20 episodes (≈ 3.5 , see Figure 3c), especially *Defect* coins. This leaves only *Cooperate* coins for co-players. Interacting with a more cooperative ratio, the co-player receives some reward, reinforcing the co-player to play a cooperative ratio in the future. Figure 3b shows the meta-agent and co-player converge to

collecting a large ratio of *Cooperate* coins ($\approx (0.4, 0.6)$), in contrast to independent learners (≈ 0.1) (grey dashed line). Interestingly, a (meta-agent, co-player) pair collects more coins ($\approx (3.0, 2.0)$) than a pair of independent agents ($\approx (1.5, 1.5)$) - this is because the independent learners maximise their return under mutual defection only by increasing interactions within an episode.

In the IMPitM, GS does not learn to shape, as expected from *Insight 2*, whereas M-FOS and SHAPER does. SHAPER and M-FOS achieve similar performances. (see Table 1).

Insight 8: SHAPER empirically tends to find better shaping policies than M-FOS in IPDitM. SHAPER outperforms M-FOS in Table 1, providing evidence that SHAPER scales to more complex policies. SHAPER demonstrates shaping, as indicated by the final rewards, which are significantly higher for both agents than M-FOS IPDitM. We postulate that as M-FOS architecture is as expressive as SHAPER, its complexities and biases hinder ES’ ability to find optimal solutions (for training training curves, see the Appendix).

In Table 4, we show that SHAPER finds policies leading to improved global welfare in cross-play with M-FOS and GS. In cross-play, the shaping algorithms are trained against Naive Learners and evaluated against each other. This experiment motivates that SHAPER’s inductive biases leads to finding more robust policies even when evaluated out of distribution. Note that Shaper vs. Shaper achieves similar scores as M-FOS vs M-FOS. However, Shaper achieves better scores against M-FOS ($7.32 \pm 0.34, 5.08 \pm 0.36$) and GS ($28.61 \pm 1.82, 20.23 \pm 1.27$). Also, note how GS achieves its highest payoff when playing against Shaper.

In our **ablations**, we find that context is beneficial for shaping in the IMP. In the “Context Challenge”, SHAPER (-0.8) outperforms SHAPER w/o Context (-1.25) (see Table 3). For shaping to occur in this challenge, we expect methods to change their strategy at $e = 2$ episodes. SHAPER demonstrates dynamic shaping by switching, yet SHAPER w/o Context’s policy does not adapt and does not exploit the stop (see Fig. 4). This result provides evidence that context is needed to shape.

In the “History Challenge”, when playing the IMP with a small number of inner-episodes ($e = 2$), we expect meta-agents without context to be unable to identify co-players’ current learning and thus cannot shape. We find that SHAPER shapes agents, whilst

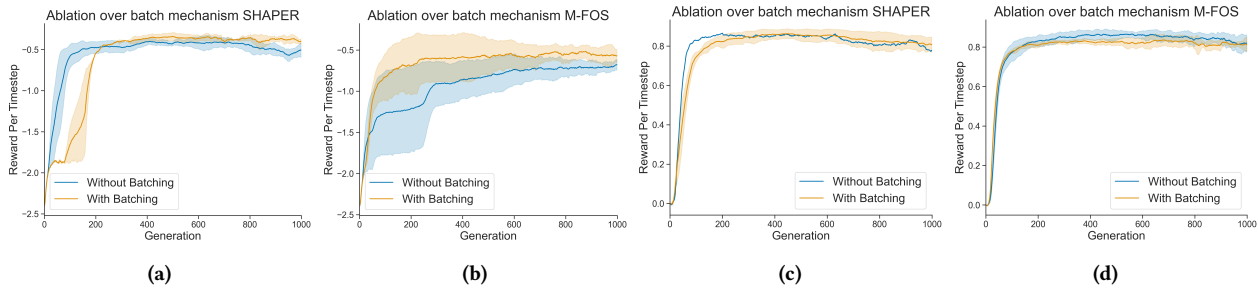


Figure 5: Reward per timestep throughout training for The “Average” challenge. Results are presented over matrix games for 5 seeds. In a) and b) we evaluate OS methods on IPD and in c) and d) we evaluate on IMP. We note that batching only helps M-FOS in the IPD. This indicates batching is only useful in sufficiently diverse environments, relative to the OS method.

Table 4: Episodic reward for in a single evaluation trial against different OS shaping methods in IPDitM. Neither agent takes gradient updates, but those with memory SHAPER and M-FOS use memory to change their policy during the trial. Results are reported for the row player in each match. We report mean and std over 5 seeds.

	SHAPER	GS	M-FOS
SHAPER	16.48 ± 0.88	28.61 ± 1.82	7.32 ± 0.34
GS	20.23 ± 1.27	0 ± 0	1.91 ± 0.27
M-FOS	5.08 ± 0.36	1.35 ± 0.28	16.25 ± 0.95

SHAPER w/o Context does not shape agents as indicated by better rewards, 0.5 vs 0.0 (see Appendix). Interestingly, we also found that with a longer inner-episode length ($E = 100$), SHAPER w/o Context uses *history* to shape its co-player (see Appendix). This shows that history can encode co-players’ learning dynamics in some environments.

In the “Average Challenge”, we find that averaging across the batch only helps M-FOS in the IPD, as it improves convergence speed. In all other scenarios, averaging across the batch did not significantly improve performance (see Figure 5b). Shaping agents must approximate, via observations, a co-players update rule. If this update is batched (such as with stochastic gradient descent), the batching mechanism should in theory provide a better estimate. If the batching mechanism is not required, this suggests experience in the update is not diverse. Comparing games, the diversity of co-player behaviours within the IMP is much less than IPD. Within the IPD, SHAPER sees no improvement with the batch mechanism compared to M-FOS (see Figure 5a - 5b). Here we postulate that given M-FOS has a limited context (1-step), batching provides M-FOS with greater context such that it can infer co-player learner. Shaper does not require averaging as it captures more context via its hidden than M-FOS does. This suggests that moving forward, OS methods should consider Context, History and Batching, as mechanisms for observing the experience / learning of co-players.

6 RELATED WORK

Opponent Shaping methods explicitly account for their opponent’s learning. Just like SHAPER, these approaches recognise that the actions of any one agent influence their co-players policy and

seek to use this mechanism to their advantage [10, 12, 17, 22, 38, 41]. However, in contrast to SHAPER, these approaches require privileged information to shape their opponents. These models are also myopic since anticipating many steps is intractable due to the difficulty of estimating higher-order gradients. Balaguer et al. [4] and Lu et al. [25] solve the issues above by framing opponent shaping as a meta reinforcement learning problem, which allows them to account for long-term shaping, where there is no need for higher-order gradients.

Algorithms for Social Dilemmas often achieve desirable outcomes in high-dimensional social dilemmas yet assume access to hand-crafted notions of adherence [40], social influence [3, 16], gifting [27] or social conventions [19]. While these approaches can achieve desirable outcomes, they change the agent’s objectives and alter the dynamics of the underlying game.

Multi-Agent Meta-Learning methods have also shown success in general-sum games with other learners [1, 18, 39]. Similar to SHAPER, they take inspiration from meta-RL - their approach is to learn the optimal initial parameterisation for the meta-agent akin to Model-Agnostic Meta Learning [9]. In contrast, SHAPER uses an approach similar to RL^2 [8], which trains an RNN-based agent to implement efficient learning for its next task. Finally, SHAPER is optimised using ES, which empirically performs better with long-time horizons than policy-gradient methods [24–26].

7 CONCLUSION

When agents interact, the actions of each agent influence the rewards and observations of others and, through their learning, ultimately affect their behaviour. Leveraging this connection is called opponent shaping, and has received considerable attention recently.

This paper introduces SHAPER, a shaping method suitable for high-dimensional games. We are the first to scale shaping successfully to long-time horizon general-sum games with temporally-extended actions, and we provide extensive performance analysis in these settings. We formalise the concept of history and context for shaping and analyse their respective roles empirically. Next, we formalise the previously implicit concept of averaging across the batch and show that it’s helpful for previous methods to learn. Future work might investigate scenarios where averaging across a batch is also necessary for SHAPER. Finally, we identify a fundamental problem in the widely-used CoinGame.

8 ETHICS STATEMENT*

Shaping can be used for good and bad. Empirically, shaping has lead learning agents to find more prosocial solutions in mixed-incentive settings. However, one can imagine scenarios where shaping is used with a negative impact on society. Assuming that learning agents will be deployed in the real world, e.g., online learning self-driving cars, it is important we understand how such agents interact. Early opponent shaping research has already shown that two naive agents mutually defect in the iterated prisoner’s dilemma and that opponent shaping leads to the more prosocial tit-for-tat strategy. It is important that we develop these methods further, investigate if they keep leading to more prosocial outcomes even in more difficult environments and if not, what improvements can we make such that they do. In our work, we show that in grid-worlds with temporally-extended actions and long-time horizons, opponent shaping tends to find more prosocial solutions than Naive Learners. Investigating shaping is important to prevent misuse of the paradigm. We are at the beginning of fundamental research in shaping and a better understanding of the necessary components to achieve shaping will help us to better control shaping agents. Opponent Shaping is still in an early phase of development, and practical implications are limited, so immediate negative societal influence is unlikely.

REFERENCES

- [1] Maruan Al-Shedivat, Trapit Bansal, Yura Burda, Ilya Sutskever, Igor Mordatch, and Pieter Abbeel. 2018. Continuous Adaptation via Meta-Learning in Nonstationary and Competitive Environments. In *International Conference on Learning Representations*.
- [2] Robert Axelrod and William D Hamilton. 1981. The evolution of cooperation. *science* 211, 4489 (1981), 1390–1396.
- [3] Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, Athul Paul Jacob, Mojtaba Komeili, Karthik Konath, Minae Kwon, Adam Lerer, Mike Lewis, Alexander H. Miller, Sasha Mitts, Adithya Renduchintala, Stephen Roller, Dirk Rowe, Weiyang Shi, Joe Spisak, Alexander Wei, David Wu, Hugh Zhang, and Markus Zijlstra. 2022. Human-level play in the game of Diplomacy by combining language models with strategic reasoning. *Science* 378, 6624 (2022), 1067–1074. <https://doi.org/10.1126/science.ade9097>
- [4] Jan Balaguer, Raphael Koster, Christopher Summerfield, and Andrea Tacchetti. 2022. The Good Shepherd: An Oracle Agent for Mechanism Design. arXiv preprint arXiv:2202.10135.
- [5] Noam Brown and Tuomas Sandholm. 2017. Libratus: the superhuman AI for no-limit poker. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*.
- [6] Allan Dafoe, Edward Hughes, Yoram Bachrach, Tatum Collins, Kevin R McKee, Joel Z Leibo, Kate Larson, and Thore Graepel. 2021. Open Problems in Cooperative AI. In *Cooperative AI workshop*.
- [7] Robyn M Dawes. 1980. Social dilemmas. *Annual review of psychology* 31, 1 (1980), 169–193.
- [8] Yan Duan, John Schulman, Xi Chen, Peter L. Bartlett, Ilya Sutskever, and Pieter Abbeel. 2016. RL²: Fast Reinforcement Learning via Slow Reinforcement Learning. arXiv preprint arXiv:1611.02779.
- [9] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 70)*. 1126–1135.
- [10] Jakob Foerster, Richard Y Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. 2018. Learning with Opponent-Learning Awareness. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. 122–130.
- [11] Jakob Foerster, Francis Song, Edward Hughes, Neil Burch, Iain Dunning, Shimon Whiteson, Matthew Botvinick, and Michael Bowling. 2019. Bayesian action decoder for deep multi-agent reinforcement learning. In *International Conference on Machine Learning*. PMLR, 1942–1951.
- [12] Kitty Fung, Qizhen Zhang, Chris Lu, Timon Willi, and Jakob Nicolaus Foerster. 2023. Analyzing the Sample Complexity of Model-Free Opponent Shaping. In *ICML Workshop on New Frontiers in Learning, Control, and Dynamical Systems*. <https://openreview.net/forum?id=Dm2fbPpU6v>
- [13] Marc Harper, Vincent Knight, Martin Jones, Georgios Koutsovoulos, Nikoleta E. Glynatsi, and Owen Campbell. 2017. Reinforcement learning produces dominant strategies for the Iterated Prisoner’s Dilemma. *PLOS ONE* 12, 12 (2017), e0188046.
- [14] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Comput.* 9, 8 (1997), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- [15] Max Jaderberg, Wojciech M. Czarnecki, Iain Dunning, Luke Marris, Guy Lever, Antonio Garcia Castañeda, Charles Beattie, Neil C. Rabinowitz, Ari S. Morcos, Avraham Ruderman, Nicolas Sonnerat, Tim Green, Louise Deason, Joel Z. Leibo, David Silver, Demis Hassabis, Koray Kavukcuoglu, and Thore Graepel. 2019. Human-level performance in 3D multiplayer games with population-based reinforcement learning. *Science* 364, 6443 (2019), 859–865. <https://doi.org/10.1126/science.aau6249>
- [16] Natasha Jaques, Angeliki Lazaridou, Edward Hughes, Caglar Gulcehre, Pedro Ortega, DJ Strouse, Joel Z Leibo, and Nando De Freitas. 2019. Social influence as intrinsic motivation for multi-agent deep reinforcement learning. In *International conference on machine learning*. PMLR, 3040–3049.
- [17] Dong-Ki Kim, Miao Liu, Matthew Riemer, Chuangchuan Sun, Marwa Abdulhai, Golnaz Habibi, Sebastian Lopez-Cot, Gerald Tesaro, and Jonathan P. How. 2021. A Policy Gradient Algorithm for Learning to Learn in Multiagent Reinforcement Learning. In *International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*. 5541–5550.
- [18] Dong Ki Kim, Miao Liu, Matthew D Riemer, Chuangchuan Sun, Marwa Abdulhai, Golnaz Habibi, Sebastian Lopez-Cot, Gerald Tesaro, and Jonathan How. 2021. A policy gradient algorithm for learning to learn in multiagent reinforcement learning. In *International Conference on Machine Learning*. PMLR, 5541–5550.
- [19] Raphael Köster, Kevin R McKee, Richard Everett, Laura Weidinger, William S Isaac, Edward Hughes, Edgar A Duéñez-Guzmán, Thore Graepel, Matthew Botvinick, and Joel Z Leibo. 2020. Model-free conventions in multi-agent reinforcement learning with heterogeneous preferences. arXiv preprint arXiv:2010.09054 (2020).
- [20] Joel Z. Leibo, Edgar A. Duéñez-Guzmán, Alexander Vezhnevets, John P. Agapiou, Peter Sunehag, Raphael Koster, Jayd Matyas, Charlie Beattie, Igor Mordatch, and Thore Graepel. 2021. Scalable Evaluation of Multi-Agent Reinforcement Learning with Melting Pot. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*. PMLR, 6187–6199.
- [21] Adam Lerer and Alexander Peysakhovich. 2017. Maintaining cooperation in complex social dilemmas using deep reinforcement learning. *CoRR abs/1707.01068* (2017).
- [22] Alistair Letcher, David Balduzzi, Sébastien Racanière, James Martens, Jakob N. Foerster, Karl Tuyls, and Thore Graepel. 2019. Differentiable Game Mechanics. *J. Mach. Learn. Res.* 20 (2019), 84:1–84:40.
- [23] Alistair Letcher, Jakob N. Foerster, David Balduzzi, Tim Rocktäschel, and Shimon Whiteson. 2019. Stable Opponent Shaping in Differentiable Games. In *7th International Conference on Learning Representations*.
- [24] Chris Lu, Jakob Grudzien Kuba, Alistair Letcher, Luke Metz, Christian Schröder de Witt, and Jakob N. Foerster. 2022. Discovered Policy Optimisation. *CoRR abs/2210.05639* (2022). <https://doi.org/10.48550/arXiv.2210.05639>
- [25] Christopher Lu, Timon Willi, Christian A Schroeder De Witt, and Jakob Foerster. 2022. Model-Free Opponent Shaping. In *International Conference on Machine Learning*. PMLR, 14398–14411.
- [26] Chris Lu, Timon Willi, Alistair Letcher, and Jakob Nicolaus Foerster. 2022. Adversarial Cheap Talk. In *Decision Awareness in Reinforcement Learning Workshop at ICML 2022*.
- [27] Andrei Lupu and Doina Precup. 2020. Gifting in Multi-Agent Reinforcement Learning. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (Auckland, New Zealand) (AAMAS ’20)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 789–797.
- [28] William H. Press and Freeman J. Dyson. 2012. Iterated Prisoner’s Dilemma contains strategies that dominate any evolutionary opponent. *Proceedings of the National Academy of Sciences* 109, 26 (2012), 10409–10413.
- [29] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. 2018. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *International conference on machine learning*. PMLR, 4295–4304.
- [30] Tim Salimans, Jonathan Ho, Xi Chen, and Ilya Sutskever. 2017. Evolution Strategies as a Scalable Alternative to Reinforcement Learning. arXiv preprint arXiv:1703.03864.
- [31] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. arXiv preprint arXiv:1707.06347.
- [32] L. S. Shapley. 1953. Stochastic Games. *Proceedings of the National Academy of Sciences* 39, 10 (1953), 1095–1100.
- [33] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Vedavyas Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy P. Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. 2016. Mastering the game of Go with deep neural networks and tree search. *Nat.* 529, 7587 (2016), 484–489.

- [34] Glenn H Snyder. 1971. "Prisoner's Dilemma" and "Chicken" Models in International Politics. *International Studies Quarterly* 15, 1 (1971), 66–103.
- [35] Alexander Vezhnevets, Yuhuai Wu, Maria Eckstein, Rémi Leblond, and Joel Z Leibo. 2020. Options as responses: Grounding behavioural hierarchies in multi-agent reinforcement learning. In *International Conference on Machine Learning*. PMLR.
- [36] Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh, Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, Laurent Sifre, Trevor Cai, John P. Agapiou, Max Jaderberg, Alexander Sasha Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dalibard, David Budden, Yury Sulsky, James Molloy, Tom Le Paine, Çağlar Gülçehre, Ziyu Wang, Tobias Pfaff, Yuhuai Wu, Roman Ring, Dani Yogatama, Dario Wünsch, Katrina McKinney, Oliver Smith, Tom Schaul, Timothy P. Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps, and David Silver. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nat.* 575, 7782 (2019), 350–354.
- [37] Timon Willi, Akbir Khan, Newton Kwan, Mikayel Samvelyan, Chris Lu, and Jakob Foerster. 2023. Pax: Multi-Agent Learning in JAX. <https://github.com/ucl-dark/pax>.
- [38] Timon Willi, Alistair Letcher, Johannes Treutlein, and Jakob N. Foerster. 2022. COLA: Consistent Learning with Opponent-Learning Awareness. In *International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 162)*. 23804–23831.
- [39] Zhe Wu, Kai Li, Enmin Zhao, Hang Xu, Meng Zhang, Haobo Fu, Bo An, and Junliang Xing. 2021. L2E: Learning to Exploit Your Opponent. arXiv preprint arXiv:2102.09381.
- [40] Yuyu Yuan, Ting Guo, Pengqian Zhao, and Hongpu Jiang. 2022. Adherence Improves Cooperation in Sequential Social Dilemmas. *Applied Sciences* 12, 16 (2022), 8004.
- [41] Stephen Zhao, Chris Lu, Roger Baker Grosse, and Jakob Nicolaus Foerster. 2022. Proximal Learning With Opponent-Learning Awareness. *arXiv preprint arXiv:2210.10125* (2022).