

# A Meta-Game Evaluation Framework for Deep Multiagent Reinforcement Learning

Zun Li<sup>1</sup>, Michael P. Wellman<sup>2</sup>

<sup>1</sup>Google DeepMind

<sup>2</sup>University of Michigan

lizun@google.com, wellman@umich.edu

## Abstract

Evaluating deep multiagent reinforcement learning (MARL) algorithms is complicated by stochasticity in training and sensitivity of agent performance to the behavior of other agents. We propose a meta-game evaluation framework for deep MARL, by framing each MARL algorithm as a *meta-strategy*, and repeatedly sampling normal-form empirical games over combinations of meta-strategies resulting from different random seeds. Each empirical game captures both *self-play* and *cross-play* factors across seeds. These empirical games provide the basis for constructing a sampling distribution, using bootstrapping, over a variety of game analysis statistics. We use this approach to evaluate state-of-the-art deep MARL algorithms on a class of negotiation games. From statistics on individual payoffs, social welfare, and empirical best-response graphs, we uncover strategic relationships among self-play, population-based, model-free, and model-based MARL methods. We also investigate the effect of run-time search as a *meta-strategy operator*, and find via meta-game analysis that the search version of a meta-strategy generally leads to improved performance.

## 1 Introduction

Evaluating complex AI algorithms requires careful attention to stochastic factors. Deep reinforcement learning (RL) algorithms in particular are subject to randomness within the algorithm and the operational environment, and variations with choices of hyperparameters and initial conditions. It is conventional to address these uncertainties in part by aggregating results across multiple runs [Bellemare *et al.*, 2013]. Deep multiagent RL (MARL) algorithms present these issues, and additional challenges due to agent interactions. Evaluating performance against humans has been one source of compelling demonstrations [Silver *et al.*, 2016; Vinyals *et al.*, 2019; Perolat *et al.*, 2022], but this approach is limited by the range of tasks for which human expertise exists, and the cost of engaging it when it is available. Generally speaking, we lack evaluation protocols for comparing different MARL methods in a statistically principled way.

In purely adversarial (*i.e.*, two-player zero-sum) environments, distance to Nash equilibrium may be a sufficient metric [Brown *et al.*, 2020; Schmid *et al.*, 2023], as all equilibria are interchangeably optimal. More generally, where there are multiple equilibria or where we do not necessarily expect equilibrium behavior, the metrics for MARL performance may be less clear. In collaborative domains, global team return is the common objective [Foerster *et al.*, 2018], however complex learning dynamics may lead agents using the same MARL algorithm to equilibria of distinct machine conventions in different runs [Hu *et al.*, 2020].

We seek an approach to evaluating deep MARL in general-sum domains. We propose a *meta-game* evaluation framework (§5), which frames MARL algorithms as *meta-strategies*: mapping games and random seeds to joint policies. We sample seed combinations to generate meta-game instances, from which we compute evaluation metrics of interest based on game-theoretic solution concepts. Through resampling and bootstrap techniques, we generate a statistical characterization of algorithm performance in these games. Our contributions:

- The meta-game evaluation framework.
- A new search algorithm for imperfect information games, *Gumbel Information-Set Monte-Carlo Tree Search* (§6), based on the recent development of Gumbel AlphaZero [Danilhelka *et al.*, 2022].
- Extensive experiments on state-of-the-art MARL algorithms (§7.2) on a class of negotiation games (§7.1), illustrating the framework and providing new evidence regarding the algorithms studied.

## 2 Related Work

While evaluating single-agent deep RL algorithms is well-studied [Henderson *et al.*, 2018; Agarwal *et al.*, 2021], there are relatively few works that consider evaluation principles for deep MARL [Gronauer and Diepold, 2022]. Typically, agents are evaluated against a selected set of background opponents or emergent behaviors in certain contexts [Lowe *et al.*, 2017; Leibo *et al.*, 2021]. Other work has defined evaluation protocols for cooperative settings [Papoudakis *et al.*, 2020; Gorsane *et al.*, 2022], where a global team reward is well-defined. Our scenario falls into the *agent-vs-agent* cate-

gory of Balduzzi *et al.* [2018], who argue for the necessity of comprehensively considering the possible joint interactions.

Kiekintveld and Wellman [2008] employed a concept of meta-games for evaluating general game-playing methods. In their setting, meta-strategies map normal-form game specifications to strategy choices. Treutlein *et al.* [2021] construct *label-free coordination* (LFC) games among instances of a cooperative MARL algorithm, in order to study zero-shot coordination [Hu *et al.*, 2020]. LFC games are a kind of meta-game in our sense, as they take dynamic game descriptions (Dec-POMDPs) as input.

Bootstrapping is a non-parametric statistical approach that constructs sampling distributions for any specified statistic by resampling from the original dataset [Davison and Hinkley, 1997]. Bootstrapping techniques have been applied in machine learning methods such as aggregating decision-tree models [Breiman, 1996]. Widenbeck *et al.* [2014] applied the bootstrap for statistical analysis of game-theoretic models estimated from simulations.

### 3 Game Theory Preliminaries

A *normal-form representation* of a *game*  $\mathcal{G}$  consists of a *player set*  $\mathcal{N} = \{1, \dots, N\}$ , and for each player  $i \in \mathcal{N}$  a *pure strategy space*  $\Pi_i$  and a *utility function*  $u_i : \Pi \rightarrow \mathbb{R}$ .  $\Pi = \Pi_1 \times \dots \times \Pi_N$  is the *joint pure strategy space*. A *mixed strategy*  $\sigma_i \in \Delta(\Pi_i)$  for player  $i$  defines a probability distribution over that player’s pure strategy space. Player  $i$ ’s *payoff* for choosing  $\sigma_i$  while others play  $\sigma_{-i} = (\sigma_j)_{j \neq i}$  is given by expectation over the respective mixtures:  $u_i(\sigma) = \mathbb{E}_{\pi_i \sim \sigma_i, \pi_{-i} \sim \sigma_{-i}}[u(\pi_i, \pi_{-i})]$ .

The *regret* of player  $i$  who plays  $\sigma_i$  when others are playing  $\sigma_{-i}$  is  $\text{REGRET}_i(\sigma_i, \sigma_{-i}) = \max_{\pi'_i \in \Pi_i} u_i(\pi'_i, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i})$ . A *Nash equilibrium* (NE) is a mixed strategy profile  $\sigma$  such that nobody has positive regret:  $\forall i. \text{REGRET}_i(\sigma_i, \sigma_{-i}) = 0$ . As a measure of approximation to NE, we define  $\text{SUMREGRET}(\sigma) = \sum_i \text{REGRET}_i(\sigma_i, \sigma_{-i})$ .

If a game is known to be symmetric, we can reduce its normal-form description complexity. Formally, in a *symmetric game*, players share the same strategy space  $\Pi_1 = \dots = \Pi_N = \Pi$ , and utility function  $u_1 = \dots = u_N = u$ . Furthermore, each player’s utility is permutation-invariant to other players’ strategies. For symmetric games we overload  $\Pi$  to refer to the common individual strategy space, rather than the joint space. In a *symmetric profile*, every player adopts the same (generally mixed) strategy. Solutions in symmetric profiles are guaranteed to exist in relevant settings [Nash, 1951; Cheng *et al.*, 2004; Hefti, 2017], and are generally preferred absent any basis for breaking symmetry [Kreps, 1990].

Also given symmetry, let  $u(\sigma', \sigma)$  be the expected payoff of a player when it chooses a strategy  $\sigma' \in \Delta(\Pi)$  while the other  $N - 1$  play the same mixed strategy  $\sigma \in \Delta(\Pi)$ . We thus have  $\text{REGRET}(\sigma', \sigma) = \max_{\pi' \in \Pi} u(\pi', \sigma) - u(\sigma', \sigma)$ . A symmetric NE strategy  $\sigma$  satisfies  $\text{REGRET}(\sigma, \sigma) = 0$ .

An *extensive-form game* (EFG) representation goes beyond normal form to include temporal and information structure. Effectively, an EFG defines a dynamical system (usually visualized by a tree) with a world state or history  $h$  that is not necessarily fully observable by every player. At each

state  $h$ , starting with the initial state  $h_0$ , there is a single player  $\tau(h) \in \mathcal{N} \cup \{c\}$ , where  $c$  is the *chance player*, designated to select an action  $a \in \mathcal{A}(h)$ . The chance player chooses according to a fixed random policy. Player  $i \in \mathcal{N}$  chooses according to its *information state/set*,  $s_i(h)$ , which comprises the information that  $i$  has observed at  $h$ . Often  $s_i(h)$  is represented by a concatenation of public and private action-observation histories. Following action  $a$ , the world transits to  $h' = ha$  until a terminal state  $z \in \mathcal{Z}$  is reached. Then each player  $i \in \mathcal{N}$  receives a utility  $u_i(z)$  as a function of  $z$ . A *behavioral strategy* or *policy*,  $\pi_i$ , of player  $i$  is a mapping from  $i$ ’s infostates to distributions over actions. For EFGs we overload  $\Pi_i$  to refer to the set of behavioral strategies of player  $i$ . A joint behavioral strategy profile  $\pi$  thus induces a probability distribution over the terminal states and we define  $u_i(\pi)$  as the expected payoff for player  $i$  under this distribution. In this paper, we consider EFGs with *perfect recall*, that is, information sets  $s_i(h)$  distinguish all actions  $i$  had taken to reach  $h$ . A consequence is that any mixed strategy  $\sigma_i \in \Delta(\Pi_i)$  is payoff-equivalent to some behavioral strategy  $\pi_i \in \Pi_i$  [Aumann, 1964].

### 4 Multiagent Training Algorithms

We define a *multiagent training algorithm* (MATA)  $\mathcal{M}$  as a stochastic procedure that produces a policy profile  $\pi = \mathcal{M}(\mathcal{G}, \Theta, \omega)$  for an EFG. In general, the input EFG  $\mathcal{G}$  cannot be tractably represented as an explicit game tree. Instead, we assume the game is given in the form of a black-box simulator that the algorithm can exercise by submitting actions and receiving observations and rewards.  $\Theta$  is the set of hyperparameters of the algorithm, and  $\omega$  is a random seed. If  $\mathcal{G}$  is symmetric, then it is often natural to constrain the output  $\pi$  to be likewise symmetric (*i.e.*, single policy to be played by all). More generally,  $\pi$  is a policy profile. For the AlphaZero MATA, for example,  $\mathcal{G}$  could be represented by a Go simulator, and  $\Theta$  would include the learning rate schedule and neural architecture. The output profile  $\pi$  specifies Go-playing policies for white and black, respectively.

A MATA is effectively a form of *meta-strategy*: a procedure that given a game  $\mathcal{G}$ , generates a strategy profile for  $\mathcal{G}$ . We can employ the MATA to play from the perspective of any particular player  $i$ , simply by selecting the  $i^{\text{th}}$  element  $\pi(i)$  of the output profile.

A key issue for analysis of MATAs is uncertainty in strategy generation. It has been well observed (*e.g.*, by Hu *et al.* [2020]) that a MATA with the same  $\mathcal{G}$  and  $\Theta$  but different  $\omega$  may generate policies with vastly different strategic behaviors. For example, in a negotiation game, different runs of a MATA may lead to strategies that adopt distinct offering conventions. In the present work, we assume the hyperparameters  $\Theta^m$  for each MATA  $\mathcal{M}^m$  have been fixed, so the uncertainty in behavior of a training algorithm is fully captured by the random seeds. Note that there is always discretion about what one considers a distinct MATA  $\mathcal{M}$  versus a parametric variation, so it is possible to bring choice among hyperparameter settings within the scope of our analysis framework.

## 5 Meta-Game Evaluation Framework

Given  $M$  different MATAs  $\{(\mathcal{M}^1, \Theta^1), \dots, (\mathcal{M}^M, \Theta^M)\}$  with associated hyperparameters, how can we evaluate their relative performance with respect to a given game  $\mathcal{G}$ ? Viewing the MATAs simply as game-solvers, we could focus on measures of their effectiveness in deriving a solution—for example, time and accuracy of convergence to Nash equilibrium. Viewing the MATAs’ role as generating strategies to play a game, however, requires a different focus that considers the interaction with other strategy generators. This is particularly salient for games that have a multiplicity of solutions (the general case), or for which the operable solution concept may be open to question. Consequently, we propose to analyze competing MATAs by framing their interaction as itself a game. As MATAs are *meta-strategies*, we refer to this approach as a *meta-game* evaluation framework.

As noted above, we are particularly concerned with uncertainty in the results of multiagent training. In analysis of the MATA meta-game, therefore, we aim to characterize the implications of this uncertainty in probabilistic form.

### 5.1 Empirical Game-Theoretic Analysis

Our approach employs *empirical game-theoretic analysis* (EGTA), a methodology for reasoning about games through agent-based simulation [Wellman *et al.*, 2024]. EGTA aligns with our assumption that the game of interest  $\mathcal{G}$  is defined by a black-box simulator. In the typical framing, EGTA operates by estimating an *empirical game model* in normal form over an enumerated strategy set. The enumerated strategies are a small selection from the full strategy space of the extensive game represented by the simulator.

We employ EGTA with respect to an  $N$ -player game of interest  $\mathcal{G}$ , and strategies defined by the output of MATAs. The meta-game is likewise over  $N$  players, and is symmetric regardless of whether  $\mathcal{G}$  is symmetric or not. Let  $\hat{\pi}^m$  denote the output from MATA  $m$ , for instance  $\hat{\pi}^m = \mathcal{M}^m(\mathcal{G}, \Theta^m, \omega)$ : the result from running the MATA on  $\mathcal{G}$  for a particular random seed. We also allow that  $\hat{\pi}^m$  be an aggregate of policy profiles from multiple random seeds. From these MATA-generated profiles  $\hat{I} = \{\hat{\pi}^1, \dots, \hat{\pi}^M\}$ , we construct an empirical meta-game  $\mathcal{MG}(\hat{I})$  over policy space  $\hat{I}$  as follows.

If the base game  $\mathcal{G}$  is symmetric, then the  $\hat{\pi}^m$  are single-player policies, and we estimate the meta-game utility function by the standard EGTA approach of simulating profiles over  $\hat{I}$ . If  $\mathcal{G}$  is not symmetric, then each  $\hat{\pi} \in \hat{I}$  is an  $N$ -player profile for  $\mathcal{G}$ . Simulating a meta-game profile  $(\hat{\pi}_1, \dots, \hat{\pi}_N)$  of these base-game profiles entails first assigning the  $\hat{\pi}$  to players, according to a random permutation *perm* drawn uniformly from the  $N!$  possibilities. Then it simulates a play where player  $i$  of  $\mathcal{MG}$  plays  $\hat{\pi}_i(\text{perm}(i))$  as if it is player *perm*( $i$ ) in  $\mathcal{G}$ . This construction also corresponds to an EFG beginning with a root chance node that uniformly chooses among  $N!$  different outcomes, each followed by a copy of the original EFG with player indices permuted.

Symmetry of the meta-game reflects a view that, for multiagent training, developing effective strategies for each of the player positions is equally important. If this were not the case, or if one wished to perform an analysis of the differential ef-

fectiveness of various MATAs from the perspectives of different players or roles, a non-symmetric (or role-symmetric) meta-game model could be constructed instead.

### 5.2 Meta-Game Evaluation Procedure

Just as single-agent RL algorithms are statistically evaluated by return performance across different random seeds, we can analyze strategic properties among MATAs across different *combinations* of seeds. Let  $\mathbf{X}$  denote statistics characterizing the strategic properties of interest (discussed below). Given an  $N$ -player game  $\mathcal{G}$  and parametrized MATAs  $\{(\mathcal{M}^1, \Theta^1), \dots, (\mathcal{M}^M, \Theta^M)\}$ , our evaluation procedure comprises the following steps:

1. Select a finite set of seeds  $\Omega^m$  for each MATA  $m$ . Generate  $\hat{\pi}^m(\omega) = \mathcal{M}^m(\mathcal{G}, \Theta^m, \omega)$  for each  $\omega \in \Omega^m$ .
2. For each  $m$ , uniformly sample  $|\Omega^m|$  seeds from  $\Omega^m$  with replacement, yielding the sequence  $(\omega_1, \dots, \omega_{|\Omega^m|})$ . Let  $\hat{\pi}^m$  be a profile that is payoff-equivalent to a uniform mixture over the multiset  $\{\hat{\pi}^m(\omega_1), \dots, \hat{\pi}^m(\omega_{|\Omega^m|})\}$ .
3. Given  $\hat{I} = \{\hat{\pi}^1, \dots, \hat{\pi}^M\}$ , estimate the symmetric empirical meta-game  $\mathcal{MG}(\hat{I})$  as described in §5.1.
4. Compute the statistics-of-interest  $\mathbf{X}$  from  $\mathcal{MG}(\hat{I})$ .
5. Repeat Steps 2 through 4. Estimated profile payoffs should be memoized for reuse across iterations. Obtain an empirical distribution of  $\mathbf{X}$  and report statistical properties of  $\mathbf{X}$ .

One way to understand this evaluation procedure is to view each MATA as a *mixed strategy*, selecting policies uniformly over the possible random seeds. The “ground-truth” meta-game represents an expectation over the results of this randomization. The empirical meta-game estimates this from finite samples  $\Omega^m$ . By resampling from these seeds at hand, Step 2 constructs multiple empirical games among MATAs, from which we construct sampling distributions over  $\mathbf{X}$  using bootstrapping.

There are a variety of choices for statistics  $\mathbf{X}$  to gather. The only requirement is that  $\mathbf{X}$  can be computed from the information in a normal-form game model. For example, one possible metric on MATA performance is *uniform-score*: average payoff against a uniform distribution over other MATAs. Such scores have been employed in a variety of contexts, however putting equal weight on the possible counterparts is questionable, as they are not equally relevant [Balduzzi *et al.*, 2018].

An alternative proposed by Jordan *et al.* [2007] is *NE-regret*:  $\text{REGRET}(\pi, \sigma^*)$ , where  $\sigma^*$  is a symmetric mixed equilibrium of  $\mathcal{MG}(\hat{I})$ . The motivation for this measure is that it focuses on behavior against rational opponents. Performance against obviously flawed opponents should carry much less weight, as they are less likely to be encountered in practice, all else equal.<sup>1</sup> For games with multiple equilibria,

<sup>1</sup>If however there is some basis to expect opponents who are flawed or boundedly rational in some particular way, then by all means it would make sense to measure regret with respect to a solution concept capturing that basis.

however, NE-regret is sensitive to the choice of solutions  $\sigma^*$ . This sensitivity is inherent to situations with multiple equilibria, but it can still be helpful to adopt a focal equilibrium to reduce ambiguity in analysis. Balduzzi *et al.* [2018] proposed **Nash averaging**, which is essentially NE-regret with respect to the max-entropy equilibrium. The intuition for preferring to evaluate with respect to higher-entropy solutions is that they reflect diversity, and thus reward robustness to a wide range of rational opponents.

### 5.3 Max-Entropy Nash Equilibrium

Computing max-entropy NE is hard in general, but practically feasible to approximate for bi-matrix games of modest size. We adapt the mixed-integer programming formulation of Sandholm *et al.* [2005]:

$$\begin{aligned}
 & \min_{\sigma^*} \sum_{\pi \in \hat{\Pi}} \sigma^*(\pi) \log \sigma^*(\pi) \\
 & \text{s.t. } \forall \pi \in \hat{\Pi}. \quad u_\pi = \sum_{\pi' \in \hat{\Pi}} \sigma^*(\pi') u(\pi, \pi'), \\
 & \quad u^* \geq u_\pi, \quad u^* - u_\pi \leq U b_\pi, \quad \sigma^*(\pi) \leq 1 - b_\pi, \\
 & \quad \sigma^*(\pi) \geq 0, \quad \sum_{\pi} \sigma^*(\pi) = 1, \quad b_\pi \in \{0, 1\},
 \end{aligned} \tag{1}$$

with real variables  $\sigma^*(\pi)$ ,  $u_\pi$ , binary variables  $b_\pi$  for each strategy  $\pi$ , and an equilibrium payoff real variable  $u^*$ .  $U$  is the maximum difference across payoffs. The variables  $b_\pi$  indicate whether strategy  $\pi$  is outside the equilibrium support. That is,  $b_\pi = 1$  iff  $\sigma^*(\pi) = 0$ . Otherwise  $u_\pi = u^*$ : strategies in the support have the same payoff. The details of our implementation are in App. A.<sup>2</sup>

Although max-entropy NE is not generally unique, it narrows the possibilities considerably compared to unconstrained (approximate or exact) NE.

## 6 Search as a Meta-Strategy Operator

Many MATAs produce a policy network  $\mathbf{p}$  that maps directly from an infostate to a distribution over actions in a forward pass for every player. Recent work has found that leveraging computation at run-time and adding search to  $\mathbf{p}$  can improve performance in large EFG domains [Silver *et al.*, 2018; Brown *et al.*, 2020; Schmid *et al.*, 2023]. As a case study for our meta-game evaluation framework, we apply it to investigate the effect of search as a general policy improver.

Toward that end, we propose a heuristic search method for large EFGs based on information-set MCTS (IS-MCTS) [Cowling *et al.*, 2012] and Gumbel AlphaZero [Danilhelka *et al.*, 2022]. Alg. 1 presents the procedure, **Gumbel IS-MCTS**, in detail. Parameterized by a policy net  $\mathbf{p}$  and a value net  $\mathbf{v}$ , Gumbel IS-MCTS conducts multiple passes over the game-tree guided by  $\mathbf{v}$  and  $\mathbf{p}$  at an input infostate  $s$ , and outputs an action  $a$  for decision-making. We can apply this procedure to a variety of underlying MATAs, as a **meta-strategy operator**: transforming  $\mathcal{M}$  to  $\mathcal{M}'$  (with additional hyperparameters like simulation budget). The meta-strategy  $\mathcal{M}'$  in effect

<sup>2</sup>Appendices can be found in an arxiv version of this paper [Li and Wellman, 2024].

---

### Algorithm 1 Gumbel IS-MCTS

---

```

1: function Gumbel-Search( $s, \mathbf{v}, \mathbf{p}$ )
2:    $\forall (s, a). R(s, a) \leftarrow 0, C(s, a) \leftarrow 0$ 
3:    $\forall a \in \mathcal{A}(s). \text{sample } g(a) \sim \text{Gumbel}(0), \hat{\mathcal{A}} \leftarrow \mathcal{A}(s)$ 
4:   repeat
5:     Sample a world state:  $h \sim \text{Pr}(h \mid s, \mathbf{p})$ 
6:     while do
7:       if  $h$  is terminal then
8:          $\mathbf{r} \leftarrow$  payoffs of players Break
9:       else if  $i \triangleq \tau(h)$  is chance then
10:         $a \leftarrow$  sample according to chance
11:       else if  $s_i(h)$  not in search tree then
12:        Add  $s_i(h)$  to search tree
13:         $\mathbf{r} \leftarrow \mathbf{v}(s_i(h)).$  Break
14:       else if  $s_i(h)$  is root node  $s$  then
15:         $a, \hat{\mathcal{A}} \leftarrow$  one step of sequential halving (Alg. 3)
16:        based on  $GS(s, a)$  and remaining actions in  $\hat{\mathcal{A}}$ 
17:       else
18:        Select  $a$  according to (3) in App. B.2.
19:       end if
20:        $h \leftarrow ha$ 
21:     end while
22:     for  $(s_i, a)$  in this trajectory do
23:       Increment  $R(s_i, a)$  by  $\mathbf{r}_i, C(s_i, a)$  by 1.
24:     end for
25:   until  $\text{num\_sim}$  simulations done
26:   return Action  $a$  that remains in  $\hat{\mathcal{A}}$ 
27: end function

```

---

adds run-time search to the output policy  $\mathbf{p}$  of  $\mathcal{M}$ . Unlike AlphaZero—which uses the same MCTS method for training and run-time, with meta-strategy operators we can explore a variety of MATAs as training-time methods which produce  $\mathbf{v}$  and  $\mathbf{p}$  for search at test-time [Sokota *et al.*, 2024] (details in §7.2).

Just like MCTS, IS-MCTS incrementally builds and traverses a search tree and aggregates statistics such as visit counts  $C(s, a)$  and aggregated values  $R(s, a)$  for visited  $(s, a)$  pairs. During each simulation of the search (line 5), a world state is sampled from a posterior belief  $\text{Pr}(h \mid s, \mathbf{p})$  assuming the opponents played according to  $\mathbf{p}$  prior to  $s$ . In our test domains,  $\text{Pr}(h \mid s, \mathbf{p})$  can be computed exactly via Bayes’s rule, where in larger domains, using particle filtering [Silver and Veness, 2010] or deep generative models [Hu *et al.*, 2021; Li *et al.*, 2023] to approximate  $\text{Pr}(h \mid s, \mathbf{p})$  are possible. Further technical details are provided in App. B.

The key feature of Gumbel IS-MCTS is how it selects actions at the search nodes. At the beginning of the search (line 3), a Gumbel random variable,  $g(a)$ , is sampled i.i.d. for each legal action  $a$  of the root, for later use in action selection. At the root (line 15), the algorithm treats each legal action as an arm of a stochastic bandit, and uses a sequential-halving algorithm [Pepels *et al.*, 2014] (Alg. 3) to distribute the simulation budget. Sequential-halving algorithms usually are designed for minimizing the *simple regret* [Bubeck *et al.*, 2009], which is the regret at the last-iteration action recommendation. By contrast, UCB-style algorithms are usu-




Figure 1: Example start of sequential bargaining game instance.

ally designed for minimizing the *accumulated regret* during an online learning process. For a game-playing search algorithm, minimizing simple regret makes more sense in terms of producing a single optimal action at a decision point.

We assign to each arm  $a$  a **Gumbel score**  $GS(s, a) = g(a) + \text{logit } \mathbf{p}(s, a) + G(\hat{q}(s, a))$ . The second term is the logit of  $a$  produced by  $\mathbf{p}$ , and the third term is a monotone transformation of the action value  $\hat{q}(s, a)$ , which is estimated by  $R(s, a)$ ,  $C(s, a)$ , and  $\mathbf{v}$ . The intuition is that a high  $\hat{q}(s, a)$  value indicates a direction for policy improvement. Indeed, the improved policy  $Imp(\mathbf{p})(s, a) \triangleq \text{SoftMax}(\text{logit } \mathbf{p}(s, a) + G(\hat{q}(s, a)))$  provably achieves higher expected values [Danihelka *et al.*, 2022, App. B]. The forms of  $G$  and  $\hat{q}$  are detailed in App. B.1.

Adding Gumbel noise  $g(a)$  implements the ‘‘Gumbel top-K-trick’’: deterministically selecting the top  $K$  actions according to  $GS(s, a)$  is equivalent to sampling  $K$  actions from  $Imp(\mathbf{p})(s, a)$  without replacement [Huijben *et al.*, 2022]. The Gumbel score induces a low-variance non-deterministic action selection of the root node during the sequential halving process, which encourages exploration while distributing the simulation budget toward actions likely to yield higher expected values.

At a non-root node (line 17), an action is selected to minimize the discrepancy between  $Imp(\mathbf{p})$  and the produced visited frequency (details in App. B.2). At the end of the search, Gumbel IS-MCTS outputs the action  $a$  that survives the sequential halving procedure.

## 7 Evaluation Study

### 7.1 Domain: Alternating Negotiation

We test our MATAs and evaluation framework in the context of a two-player negotiation game. The particular version we study, called ‘‘Deal-or-No-Deal’’ (DoND) by Lewis *et al.* [2017], is a family of two-player negotiation games where the players alternate proposals to divide a pool of resources: three different goods available in various amounts. At the start of the game, quantities of goods in the pool and players’ private valuations for units of each good are sampled from a known distribution. The valuation distribution is constrained to fix a constant value for both players of the aggregate pool:  $\mathbf{c} \cdot \mathbf{w}_1 = \mathbf{c} \cdot \mathbf{w}_2 = 10$ . For example, in Fig. 1, the pool of goods can be represented as  $\mathbf{c} = [1, 2, 3]$ . The negotiation proceeds through alternating proposals for dividing the resources. In this example, player 1 first proposes  $\mathbf{o}_1 = [1, 2, 0]$  for itself which entails  $\mathbf{o}_2 = \mathbf{c} - \mathbf{o}_1 = [0, 0, 3]$  for player 2. Player 2 rejects and proposes a counter-offer, to which 1 may either agree or continue bargaining. For a particular game parametrization  $\text{Barg}(T, \varepsilon, \gamma)$ , the game ends if either (1) a deal is made, (2) a maximum number of rounds  $T$  is reached,

or (3) chance decides to terminate the game, which happens at every round with probability  $\varepsilon$ . If an agreement  $(\mathbf{o}_1, \mathbf{o}_2)$  was reached at the  $t^{\text{th}}$  round, player  $i$  receives payoff  $\gamma^t \mathbf{w}_i \cdot \mathbf{o}_i$ . Otherwise both players receive zero payoff.

DoND is a family of challenging general-sum environments with imperfect information. Players have a common interest in reaching a deal (and quickly, for  $\gamma < 1$ ), and the different ways of dividing the pool have different total value. In our studies, we sample configurations uniformly from a database of 6796 published by Lewis *et al.* [2017]. This leads to prohibitive complexity: a game with  $T = 10$  has  $O(10^{11})$  infosets for every player, which is intractable for enumeration (detailed in App. G), and grows exponentially with  $T$ .

### 7.2 Benchmark Algorithms

Our meta-game evaluation considers the following  $M = 17$  MATAs. These algorithms represent a comprehensive set of state-of-the-art MARL algorithms, including methods using self-play based training, population based training, and model-free and model-based approaches. Details are below and in App. D.

#### Independent/Multiagent PPO (IDPPO/MAPPO)

Both IDPPO and MAPPO train policy and value nets using self-play trajectories by minimizing the trust-region clipped loss [Schulman *et al.*, 2017]. Value nets are trained by minimizing L2 loss from the targets produced by GAEs. In IDPPO each player maintains its own policy and value nets whereas in MAPPO [Yu *et al.*, 2022] players share the same neural nets.

#### Regularized Nash Dynamics (R-NaD) and NFSP

Both R-NaD [Perolat *et al.*, 2021] and Neural Fictitious Self-Play (NFSP) [Heinrich and Silver, 2016] are self-play model-free MARL algorithms originally designed for purely adversarial settings. R-NaD has recently shown success in producing human-level agents in Stratego [Perolat *et al.*, 2022], where it iteratively trains policy nets by minimizing NeuRD loss [Hennes *et al.*, 2020] and value nets using the V-Trace estimator [Espenholt *et al.*, 2018] on a sequence of regularized games. NFSP mimics the classic fictitious play algorithm [Heinrich and Silver, 2016] by alternating between training a supervised learning net that summarizes historical plays and training a DQN policy that serves as a best response.

#### Policy Space Response Oracles (PSRO) and FCP

PSRO [Lanctot *et al.*, 2017] and Fictitious Co-Play (FCP) [Strouse *et al.*, 2021] are two population-based MARL algorithms. PSRO is an EGTA method that iteratively adds policies that are best responses to the current solution. We use max-entropy Nash as the solution concept (solving an asymmetric version of (1)), and PPO as the best-response method. We consider two ways of extracting the final agents: (i) PSRO: using the final-iteration solution and (ii) PSRO-LAST: using the final-iteration best-response policies. FCP has notably demonstrated success in collaborating with humans from scratch [Strouse *et al.*, 2021]. FCP builds a population by picking policies of different skill levels on multiple self-play runs, and trains final best responses against such

populations. Our implementation of FCP uses IDPPO to generate self-play runs, picks the agents based on social welfare, and uses PPO as the best-response method.

### Gumbel Search and Vanilla AlphaZero-style Search

We include the following variants of Gumbel IS-MCTS: (i) G-Search: Using Gumbel IS-MCTS for both training the networks (Alg. 4) at training time and conducting search at test time. (ii) G-Search-PN: Using search at training time, but only policy net without search at test time. (iii) G-Search-IDPPO: Using IDPPO to train the policy and value nets, based on which Gumbel IS-MCTS executes test-time search. (iv) G-Search-MAPPO and (v) G-Search-R-NaD are similarly defined. In addition, we implement an extension of the original AlphaZero-style MCTS to IS-MCTS. The differences between this search method (Alg. 2 in App. C) and Gumbel IS-MCTS are in the mechanisms for exploration (e.g., Dirichlet noise) and action selection (e.g., PUCT) within the search tree. We include (vi) VA-Search: Using this search for both training time and test time, and (vii) VA-Search-PN: similarly defined as (ii).

### Heuristic Strategies

We further include the following baseline bargaining strategies: (i) **Uniform**, which uniformly samples among all legal actions at each decision point. (ii) **Tough**, which never agrees, and always proposes uniformly among offers that maximize its own payoff. (iii) **Soft**, which always agrees, or uniformly proposes among all offers if it is the first-mover.

## 7.3 Experimental Setup

All methods are implemented within the programming model of OpenSpiel [Lanctot *et al.*, 2019]. Each MATA that involves neural training uses models of approximately equal size (e.g., number of layers, hidden nodes). The neural net inputs are likewise the same across all MATAs, and include complete history information (i.e., losslessly represent infosets) for the DoND game. For procedures that involve self-play or RL training we employ roughly the same number of training trajectories. For all search methods we use 200 simulations per call. We fine-tuned the learning rates based on SUMREGRET performance. Most of the MATAs reliably produce approximate equilibria. Details are in App. E.

We set  $|\Omega^m| = 10$ ,  $m \in \{1, \dots, M\}$ . To simplify the evaluation procedure, we pre-compute an empirical payoff matrix covering every policy pair:  $(\hat{\pi}_1^m(\omega), \hat{\pi}_2^{m'}(\omega'))$ , for  $m, m' \in \{1, \dots, M\}$ ,  $\omega$  and  $\omega'$  among the seeds sampled for  $\mathcal{M}^m$  and  $\mathcal{M}^{m'}$ , respectively. For each payoff entry we run  $2 \times 10^4$  simulations to compute the expected payoff. We then sample and analyze  $10^6 M \times M$  empirical meta-games from this matrix, per Steps 2–4 of §5.2, to obtain distributions over the statistics of interest.

## 7.4 Results

We test all 17 MATAs on two DoND instances:  $\text{Barg}(10, 0, 1)$  and  $\text{Barg}(30, 0.125, 0.935)$ . The two settings are qualitatively different. The latter includes discounting and per-round ending probability, incentivizing the players to find and agree on a good deal in early rounds. Without these factors, the

first game reduces to an ultimatum-game-like environment. Meta-game statistics are reported in Tables 1 and 2.

In addition to NE-Regret and uniform scores, we also report a statistic we term the *NE-Nash-Bargaining-Score* (NE-NBS). The NE-NBS for MATA  $m$  in an empirical game is defined as the utility-product between  $\pi^m$  and an opponent strategy  $\sigma^*$ :  $u(\pi^m, \sigma^*) \cdot u(\sigma^*, \pi^m)$ . Here  $\sigma^*$  is a max-entropy Nash computed by (1). This statistic is intended to measure the effectiveness of an agent in achieving high social welfare and fairness with a rational opponent.

From the tables we can see that while there is positive correlation between NE-Regret Scores and Uniform Scores, NE-Regret Scores are better at identifying the most robust MATAs: the rankings of the top four MATAs in NE-regret scores are the same for both environments, and from the empirical distribution plots in App. F we can see G-Search-R-NaD and G-Search consistently produce strategies that are selected in a max-entropy equilibrium support with high probability. By contrast uniform-score comparisons can be distorted by non-salient success against weak strategies. It is also interesting to note positive correlation between NE-Regret and NE-NBS. This reflects the non-zero-sum nature of these environments.

The performance of search-based agents is especially noteworthy. Three of the top four MATAs employ search at test time. All search methods are stronger than their policy net counterparts, confirming that IS-MCTS is a policy-improver in imperfect information games. We hypothesize that this is because a search-based algorithm is usually a better responder to its policy network counterpart, while both behave strategically similarly against other methods. The relative strength among search-based algorithms mirrors that of their policy-net counterparts. In both environments, G-Search and VA-Search are among the top four. This suggests that using search at training time is compatible with the same search at test time, and produces better policy and value nets compared with other methods.

We notice some algorithms behave qualitatively differently in these two environments. We found FCP consistently produces collaborative strategies with rather high agreement probability. This may explain its better performance in the second environment, which encourages settling a deal in early rounds. Likewise, Soft achieves higher individual scores and NE-NBS in the second environment. By contrast, while Tough achieves better individual performance than Soft in the ultimatum game, it is still worse than all learning methods, as shown by its terrible NE-NBS values for both environments.

Another interesting view is provided by *best-response graphs*. In a best-response graph for a game instance, strategies are vertices, and there is a directed edge  $(m_1 \rightarrow m_2)$  iff  $m_2 = \text{argmax}_{m'} u(\pi^{m'}, \pi^{m_1})$ . We generate an aggregate graph for a MATA meta-game by recording the frequency of such edges in the sampled empirical games. These are shown for the two game settings in Fig. 2. The correspondences between the vertex labels and MATAs can be found in the ‘‘Label’’ column of the tables.

The empirical best-response graphs support several observations. First, many MATA vertices such as IDPPO(label 6) and MAPPO(7) have self-edges with non-negligible frequen-

MATA Name	Label	NE-Regret-Score	Uniform-Score	NE-NBS	MATA Name	Label	NE-Regret-Score	Uniform-Score	NE-NBS
G-Search-R-NaD	5	<b>0.002±0.020</b>	6.158±0.040	44.010±0.794	MAPPO	7	0.721±0.059	5.602±0.035	32.832±0.818
G-Search	1	0.010±0.046	<b>6.215±0.015</b>	<b>44.345±0.496</b>	IDPPO	6	0.766±0.133	5.645±0.050	30.240±2.051
R-NaD	11	0.045±0.043	5.977±0.027	42.008±0.738	NFSP	8	0.806±0.050	5.766±0.014	36.534±0.721
VA-Search	15	0.092±0.023	6.074±0.029	40.836±0.551	VA-Search-PN	16	0.843±0.071	5.570±0.026	36.540±0.717
G-Search-MAPPO	3	0.354±0.047	6.038±0.027	37.808±0.962	FCP	0	1.116±0.096	5.804±0.013	38.665±1.092
PSRO-LAST	10	0.414±0.051	5.890±0.026	39.187±0.762	Tough	13	1.801±0.093	4.017±0.031	6.581±0.137
PSRO	9	0.417±0.055	5.904±0.022	38.288±0.873	Soft	12	3.189±0.098	3.773±0.013	27.023±0.855
G-Search-IDPPO	2	0.480±0.062	6.088±0.031	36.385±1.317	Uniform	14	4.274±0.062	2.391±0.004	11.980±0.494
G-Search-PN	4	0.575±0.048	5.811±0.013	42.174±0.493					

 Table 1: Results for  $\text{Barg}(10, 0, 1)$ , with 95% confidence intervals. MATAs are listed in increasing order of NE-Regret.

MATA Name	Label	NE-Regret-Score	Uniform-Score	NE-NBS	MATA Name	Label	NE-Regret-Score	Uniform-Score	NE-NBS
G-Search-R-NaD	5	<b>0.000±0.001</b>	<b>6.147±0.013</b>	43.681±0.497	PSRO	9	0.488±0.056	5.688±0.052	39.851±0.728
G-Search	1	0.005±0.033	6.078±0.019	43.198±0.545	G-Search-MAPPO	3	0.729±0.095	5.656±0.056	32.983±0.994
R-NaD	11	0.045±0.010	6.121±0.012	<b>44.072±0.417</b>	IDPPO	6	0.749±0.059	5.504±0.042	36.350±0.685
VA-Search	15	0.075±0.023	6.058±0.014	42.605±0.610	PSRO-LAST	10	0.906±0.213	5.391±0.161	35.723±2.119
G-Search-PN	4	0.294±0.019	5.833±0.016	42.285±0.448	MAPPO	7	1.135±0.144	5.306±0.086	31.878±1.378
FCP	0	0.360±0.026	5.864±0.013	43.321±0.265	Soft	12	1.918±0.018	4.352±0.011	35.990±0.267
VA-Search-PN	15	0.412±0.033	5.668±0.014	40.314±0.325	Uniform	14	4.161±0.031	2.370±0.007	12.562±0.211
NFSP	8	0.421±0.011	5.786±0.014	39.151±0.481	Tough	13	4.542±0.082	2.685±0.025	2.270±0.097
G-Search-IDPPO	2	0.450±0.038	5.820±0.019	36.997±0.432					


 Table 2: Results for  $\text{Barg}(30, 0.125, 0.935)$ , with 95% confidence intervals. MATAs are listed in increasing order of NE-Regret.

cies. This is consistent with previous observations that self-play MARL algorithms are likely to produce agents that overfit to themselves [Hu *et al.*, 2020]. However, self-edges do not necessarily suggest poor generalization performances. For example, G-Search-NaD(5) has a self-edge with a high probability in both environments, but still performs the best. Self-edges persist even after applying search as policy-improver at test time. For example, there are self-edges for G-Search-IDPPO(2) and G-Search-MAPPO(3). This may suggest that there are certain strategic correlations between a search agent and its policy net counterpart. Interestingly, we also observe certain self-edge probabilities for PSRO(9) and PSRO-LAST(10), which are normally regarded as population-based training methods. It could be that the equilibria selected at each iteration introduce correlations between different players. By contrast, NFSP(8) and FCP(0) are two MATAs without strong self-edges; both use a uniform distribution for opponent modeling, which may reduce the correlation between player positions.

Edges across MATAs illuminate strategic interaction structure. Tough(13) is unsurprisingly the best response to Soft(12). G-Search-NaD(5) is a strong attractor in both graphs, as G-Search(1) is in the first. Sometimes a search-based MATA is a best response to its policy-net counterpart, as illustrated by R-NaD(11) to G-Search-R-NaD(5), and VA-Search-PN(16) to VA-Search(15). This could also explain why MATAs such as G-Search-PN(4) and VA-Search-PN(16) do not have apparent self-edges.

## 8 Conclusion

We propose a meta-game evaluation framework for MARL in general-sum environments. Our approach is analogous to the evaluation process for single-agent RL, effectively aggregating the strategic analysis procedures across possible worlds defined by different seed combinations. We illustrated the method by constructing a meta-game model over a compre-


 Figure 2: Empirical best-response graphs for  $\text{Barg}(10, 0, 1)$  and  $\text{Barg}(30, 0.125, 0.935)$ .

hensive set of multiagent training algorithms using deep RL on a class of negotiation games. The meta-game analysis evaluated the algorithms in multiple ways, most prominently through max entropy NE-regret ranking and the structure of best-response graphs. Bootstrap statistics provide a basis for assessing uncertainty in evaluation results.

Experimental results support several interesting observations about Gumbel IS-MCTS as a meta-strategy operator, especially regarding the value of search at both training time and test time, and on patterns of strategic interactions among the algorithms.

A key feature of our evaluation framework is the flexibility to investigate a variety of statistics of interest through a carefully structured process. Future work will focus on understanding the robustness of alternative measures that can be employed for algorithm assessment through meta-games, and developing effective tools for statistical analysis.

## Acknowledgments

This work was supported in part by funding from the US Army Research Office (MURI grant W911NF-18-1-0208), and a grant from Open Philanthropy.

## References

- [Agarwal *et al.*, 2021] Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, et al. Deep reinforcement learning at the edge of the statistical precipice. In *35th Int'l Conf. on Neural Information Processing Systems*, pages 29304–29320, 2021.
- [Aumann, 1964] Robert J. Aumann. Mixed and behavior strategies in infinite extensive game. In M. Dresher, L. S. Shapley, and A. W. Tucker, editors, *Advances in Game Theory*, Annals of Mathematics Studies, pages 627–650. Princeton University Press, 1964.
- [Balduzzi *et al.*, 2018] David Balduzzi, Karl Tuyls, Julien Perolat, and Thore Graepel. Re-evaluating evaluation. In *32nd Int'l Conf. on Neural Information Processing Systems*, 2018.
- [Bellemare *et al.*, 2013] Marc G. Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- [Breiman, 1996] Leo Breiman. Bagging predictors. *Machine Learning*, 24:123–140, 1996.
- [Brown *et al.*, 2020] Noam Brown, Anton Bakhtin, Adam Lerer, and Qucheng Gong. Combining deep reinforcement learning and search for imperfect-information games. In *34th Int'l Conf. on Neural Information Processing Systems*, pages 17057–17069, 2020.
- [Bubeck *et al.*, 2009] Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. Pure exploration in multi-armed bandits problems. In *20th Int'l Conf. on Algorithmic Learning Theory*, pages 23–37, 2009.
- [Cheng *et al.*, 2004] Shih-Fen Cheng, Daniel M. Reeves, Yevgeniy Vorobeychik, and Michael P. Wellman. Notes on equilibria in symmetric games. In *AAMAS-04 Workshop on Game-Theoretic and Decision-Theoretic Agents*, 2004.
- [Cowling *et al.*, 2012] Peter I. Cowling, Edward J. Powley, and Daniel Whitehouse. Information set Monte Carlo tree search. *IEEE Transactions on Computational Intelligence and AI in Games*, 4:120–143, 2012.
- [Danihelka *et al.*, 2022] Ivo Danihelka, Arthur Guez, Julian Schrittwieser, and David Silver. Policy improvement by planning with Gumbel. In *10th Int'l Conf. on Learning Representations*, 2022.
- [Davison and Hinkley, 1997] Anthony Christopher Davison and David Victor Hinkley. *Bootstrap Methods and Their Application*. Cambridge University Press, 1997.
- [Espenholt *et al.*, 2018] Lasse Espenholt, Hubert Soyer, Remi Munos, et al. IMPALA: Scalable distributed deep-RL with importance weighted actor-learner architectures. In *35th Int'l Conf. on Machine Learning*, pages 1407–1416, 2018.
- [Foerster *et al.*, 2018] Jakob Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. Counterfactual multi-agent policy gradients. In *32nd AAAI Conf. on Artificial Intelligence*, 2018.
- [Gorsane *et al.*, 2022] Rihab Gorsane, Omayma Mahjoub, Ruan John de Kock, Roland Dubb, Siddarth Singh, and Arnu Pretorius. Towards a standardised performance evaluation protocol for cooperative MARL. In *36th Int'l Conf. on Neural Information Processing Systems*, pages 5510–5521, 2022.
- [Gronauer and Diepold, 2022] Sven Gronauer and Klaus Diepold. Multi-agent deep reinforcement learning: A survey. *Artificial Intelligence Review*, 55:895–943, 2022.
- [Hefti, 2017] Andreas Hefti. Equilibria in symmetric games: Theory and applications. *Theoretical Economics*, 12(3):979–1002, 2017.
- [Heinrich and Silver, 2016] Johannes Heinrich and David Silver. Deep reinforcement learning from self-play in imperfect-information games. *arXiv preprint arXiv:1603.01121*, 2016.
- [Henderson *et al.*, 2018] Peter Henderson, Riashat Islam, Philip Bachman, et al. Deep reinforcement learning that matters. In *32nd AAAI Conf. on Artificial Intelligence*, 2018.
- [Hennes *et al.*, 2020] Daniel Hennes, Dustin Morrill, Shayegan Omidshafiei, et al. Neural replicator dynamics: Multiagent learning via hedging policy gradients. In *19th Int'l Conf. on Autonomous Agents and Multi-Agent Systems*, pages 492–501, 2020.
- [Hu *et al.*, 2020] Hengyuan Hu, Adam Lerer, Alex Peysakhovich, and Jakob Foerster. “Other-play” for zero-shot coordination. In *37th Int'l Conf. on Machine Learning*, pages 4399–4410, 2020.
- [Hu *et al.*, 2021] Hengyuan Hu, Adam Lerer, Noam Brown, and Jakob Foerster. Learned belief search: Efficiently improving policies in partially observable settings. *arXiv preprint arXiv:2106.09086*, 2021.
- [Huijben *et al.*, 2022] Iris A. M. Huijben, Wouter Kool, Max B. Paulus, and Ruud J. G. Van Sloun. A review of the Gumbel-max trick and its extensions for discrete stochasticity in machine learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):1353–1371, 2022.
- [Jordan *et al.*, 2007] Patrick R. Jordan, Christopher Kiekintveld, and Michael P. Wellman. Empirical game-theoretic analysis of the TAC supply chain game. In *6th Int'l Conf. on Autonomous Agents and Multi-Agent Systems*, pages 1188–1195, 2007.
- [Kiekintveld and Wellman, 2008] Christopher Kiekintveld and Michael P. Wellman. Selecting strategies using empirical game models: An experimental analysis of meta-strategies. In *7th Int'l Conf. on Autonomous Agents and Multi-Agent Systems*, pages 1095–1101, 2008.
- [Kreps, 1990] David M. Kreps. *Game Theory and Economic Modelling*. Oxford University Press, 1990.
- [Lanctot *et al.*, 2017] Marc Lanctot, Vinicius Zambaldi, Audrunas Gruslys, et al. A unified game-theoretic approach to multiagent reinforcement learning. In *31st Int'l Conf.*



- on *Neural Information Processing Systems*, pages 4190–4203, 2017.
- [Lanctot *et al.*, 2019] Marc Lanctot, Edward Lockhart, Jean-Baptiste Lespiau, et al. Openspiel: A framework for reinforcement learning in games. *arXiv preprint arXiv:1908.09453*, 2019.
- [Leibo *et al.*, 2021] Joel Z. Leibo, Edgar A. Dueñez-Guzman, Alexander Vezhnevets, et al. Scalable evaluation of multi-agent reinforcement learning with Melting Pot. In *38th Int’l Conf. on Machine Learning*, pages 6187–6199, 2021.
- [Lewis *et al.*, 2017] Mike Lewis, Denis Yarats, Yann N. Dauphin, Devi Parikh, and Dhruv Batra. Deal or no deal? End-to-end learning for negotiation dialogues. In *Conference on Empirical Methods in Natural Language Processing*, 2017.
- [Li and Wellman, 2024] Zun Li and Michael P. Wellman. A meta-game evaluation framework for deep multiagent reinforcement learning. *arXiv preprint arXiv:2405.00243*, 2024.
- [Li *et al.*, 2023] Zun Li, Marc Lanctot, Kevin R. McKee, et al. Combining tree-search, generative models, and Nash bargaining concepts in game-theoretic reinforcement learning. *arXiv preprint arXiv:2302.00797*, 2023.
- [Lowe *et al.*, 2017] Ryan Lowe, Yi I. Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In *31st Int’l Conf. on Neural Information Processing Systems*, 2017.
- [Nash, 1951] John Nash. Non-cooperative games. *Annals of Mathematics*, pages 286–295, 1951.
- [Papoudakis *et al.*, 2020] Georgios Papoudakis, Filippos Christianos, Lukas Schäfer, and Stefano V. Albrecht. Comparative evaluation of cooperative multi-agent deep reinforcement learning algorithms. *arXiv preprint arXiv:2006.07869*, 2020.
- [Pepels *et al.*, 2014] Tom Pepels, Tristan Cazenave, Mark H. M. Winands, and Marc Lanctot. Minimizing simple and cumulative regret in monte-carlo tree search. In *ECAI-14 Workshop on Computer Games*, pages 1–15, 2014.
- [Perolat *et al.*, 2021] Julien Perolat, Remi Munos, Jean-Baptiste Lespiau, Shayegan Omidshafiei, et al. From Poincaré recurrence to convergence in imperfect information games: Finding equilibrium via regularization. In *38th Int’l Conf. on Machine Learning*, pages 8525–8535, 2021.
- [Perolat *et al.*, 2022] Julien Perolat, Bart De Vylder, Daniel Hennes, et al. Mastering the game of Stratego with model-free multiagent reinforcement learning. *Science*, 378(6623):990–996, 2022.
- [Sandholm *et al.*, 2005] Tuomas Sandholm, Andrew Gilpin, and Vincent Conitzer. Mixed-integer programming methods for finding Nash equilibria. In *19th AAAI Conf. on Artificial Intelligence*, pages 495–501, 2005.
- [Schmid *et al.*, 2023] Martin Schmid, Matej Moravcik, Neil Burch, and Rudolf Kadlec others. Student of games: A unified learning algorithm for both perfect and imperfect information games. *Science Advances*, 2023.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Silver and Veness, 2010] David Silver and Joel Veness. Monte-Carlo planning in large POMDPs. In *24th Int’l Conf. on Neural Information Processing Systems*, volume 23, 2010.
- [Silver *et al.*, 2016] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, et al. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- [Silver *et al.*, 2018] David Silver, Thomas Hubert, Julian Schrittwieser, et al. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419):1140–1144, 2018.
- [Sokota *et al.*, 2024] Samuel Sokota, Gabriele Farina, David J Wu, Hengyuan Hu, Kevin A Wang, J. Zico Kolter, and Noam Brown. The update equivalence framework for decision-time planning. In *Thirteenth International Conference on Learning Representation*, 2024.
- [Strouse *et al.*, 2021] D. J. Strouse, Kevin McKee, Matt Botvinick, Edward Hughes, and Richard Everett. Collaborating with humans without human data. In *35th Int’l Conf. on Neural Information Processing Systems*, pages 14502–14515, 2021.
- [Treutlein *et al.*, 2021] Johannes Treutlein, Michael Dennis, Caspar Oesterheld, and Jakob Foerster. A new formalism, method and open issues for zero-shot coordination. In *38th Int’l Conf. on Machine Learning*, pages 10413–10423, 2021.
- [Vinyals *et al.*, 2019] Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, et al. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- [Wellman *et al.*, 2024] Michael P. Wellman, Karl Tuyls, and Amy Greenwald. Empirical game-theoretic analysis: A survey. Technical Report 2403.04018, arXiv, 2024.
- [Wiedenbeck *et al.*, 2014] Bryce Wiedenbeck, Ben-Alexander Cassell, and Michael P. Wellman. Bootstrap statistics for empirical games. In *13th Int’l Conf. on Autonomous Agents and Multi-Agent Systems*, pages 597–604, 2014.
- [Yu *et al.*, 2022] Chao Yu, Akash Velu, Eugene Vinitsky, Jiayuan Gao, Yu Wang, Alexandre Bayen, and Yi Wu. The surprising effectiveness of PPO in cooperative multi-agent games. In *36th Int’l Conf. on Neural Information Processing Systems*, pages 24611–24624, 2022.