

Deep (Predictive) Discounted Counterfactual Regret Minimization

Hang Xu^{1,2}, Kai Li^{1,2,*}, Haobo Fu⁶, Qiang Fu⁶, Junliang Xing⁵, Jian Cheng^{1,3,4}

¹C²DL, Institute of Automation, Chinese Academy of Sciences

²School of Artificial Intelligence, University of Chinese Academy of Sciences

³AiRiA

⁴Maicro.ai

⁵Tsinghua University

⁶Tencent AI Lab

{xuhang2020, kai.li, jian.cheng}@ia.ac.cn, {haobofu, leonfu}@tencent.com, jlxing@tsinghua.edu.cn

Abstract

Counterfactual regret minimization (CFR) is a family of algorithms for effectively solving imperfect-information games. To enhance CFR’s applicability in large games, researchers use neural networks to approximate its behavior. However, existing methods are mainly based on vanilla CFR and struggle to effectively integrate more advanced CFR variants. In this work, we propose an efficient model-free neural CFR algorithm, overcoming the limitations of existing methods in approximating advanced CFR variants. At each iteration, it collects variance-reduced sampled advantages based on a value network, fits cumulative advantages by bootstrapping, and applies discounting and clipping operations to simulate the update mechanisms of advanced CFR variants. Experimental results show that, compared with model-free neural algorithms, it exhibits faster convergence in typical imperfect-information games and demonstrates stronger adversarial performance in a large poker game.

Code — <https://github.com/rpSebastian/DeepPDCFR>

Extended version — <https://arxiv.org/abs/2511.08174>

Introduction

Imperfect-information games (IIGs) serve as a foundational framework for modeling strategic interactions among multiple players where certain information remains hidden. Addressing these games poses significant challenges, as it requires reasoning under uncertainty about opponents’ private information. Such hidden information is pervasive in real-world scenarios, such as negotiation (Gratch, Nazari, and Johnson 2016), security (Lisy, Davis, and Bowling 2016), medical treatment (Sandholm 2015), and recreational games (Brown and Sandholm 2019b), making research on IIGs theoretically and practically crucial. The primary goal in solving IIGs is to compute an (approximate) Nash equilibrium (NE) (Nash 1950)—a strategy profile where no player can gain by unilaterally altering its strategy.

Similar to most research on solving IIGs, we focus on learning an NE in two-player zero-sum IIGs. We also assume the model-free setting, where the algorithm does not

have an exact simulator of the game and only samples episodes from the game. When the perfect game model is available, the family of counterfactual regret minimization (CFR) algorithms (Zinkevich et al. 2007; Tammelin 2014; Brown and Sandholm 2019a; Farina, Kroer, and Sandholm 2021) is one of the most successful approaches for computing an NE, which iteratively minimizes the cumulative counterfactual regrets of both players so that the average strategy profile converges to an NE in two-player zero-sum IIGs. Due to its robust theoretical foundation and strong empirical performance, CFR and its variants have driven several significant advancements in this field (Bowling et al. 2015; Moravčík et al. 2017; Brown and Sandholm 2018, 2019b).

When lacking a game model, outcome-sampling Monte Carlo CFR (OS-MCCFR) (Lanctot et al. 2009) has been proposed to approximate the counterfactual regrets in each iteration by sampling episodes from the game. To further scale the algorithm to large IIGs, many novel neural CFR variants have been developed. OS-DeepCFR (Brown et al. 2019) employs function approximation with deep neural networks to approximate the cumulative counterfactual regrets instead of tabular storage. DREAM (Steinberger, Lerer, and Brown 2020), which is built upon variance-reduced MC-CFR (Schmid et al. 2019), employs a learned value function as a baseline to reduce the high variance in estimating cumulative counterfactual regrets. ESCHER (McAleer et al. 2023) uses a history value function and a fixed sampling strategy for the updating player to avoid using importance sampling.

Although these neural approaches have greatly accelerated CFR in large IIGs, they primarily focus on approximating the behavior of vanilla CFR or the variant LinearCFR (Brown and Sandholm 2019a). Besides, they rely on a large replay buffer to store experiences, which is used to refit the cumulative counterfactual regrets each iteration. Recent advancements in tabular CFR have demonstrated that novel CFR variants can achieve significantly faster convergence compared to vanilla CFR and LinearCFR. To reduce the cost of picking wrong actions, CFR+ (Tammelin 2014) clips negative cumulative counterfactual regrets in each iteration, Discounted CFR (DCFR) (Farina, Kroer, and Sandholm 2021) discounts the cumulative counterfactual regrets in each iteration, and DCFR+ (Hang et al. 2022) combines

*Corresponding author.

the key insights of CFR+ and DCFR to achieve faster convergence. Predictive CFR+ (PCFR+) (Farina, Kroer, and Sandholm 2021) leverages the predictability of the counterfactual regrets in each iteration to accelerate the convergence speed. PDCFR+ (Hang et al. 2024) integrates PCFR+ and DCFR in a principled manner, showcasing faster convergence in non-poker IIGs. These tabular variants provide a promising avenue for neural CFR to further unlock its potential in achieving faster convergence.

To this end, we propose novel model-free neural CFR variants, Variance Reduction Deep DCFR+ (VR-DeepDCFR+) and Variance Reduction Deep PDCFR+ (VR-DeepPDCFR+). These algorithms employ deep neural networks to approximate the behavior of DCFR+ and PDCFR+, respectively, while leveraging learned baseline functions to mitigate the high variance caused by episode sampling.

Approximating advanced CFR variants presents a significant challenge, as these tabular CFR variants update cumulative counterfactual regrets by bootstrapping from the previous iteration. This makes it impossible to directly approximate them using samples from all iterations stored in the replay buffer, as is typically done when approximating vanilla CFR or LinearCFR. Furthermore, approximating the counterfactual regrets in each iteration is particularly challenging since their computation depends on expectation values weighted by the opponent’s reach probabilities. These unnormalized values introduce substantial difficulties for neural networks to approximate effectively. Advantages, however, have been shown to be effectively approximated and generalized using neural networks (Schulman et al. 2017), and can be interpreted as a form of weighted counterfactual regrets (Srinivasan et al. 2018). Therefore, we directly approximate cumulative advantages during each iteration by using the samples collected in the current iteration while bootstrapping from the results of the previous iteration.

Moreover, under the model-free setting, the algorithm samples only a single action per state, resulting in high variance in the collected samples during each iteration. To mitigate this issue, we introduce an auxiliary value network inspired by DREAM (Steinberger, Lerer, and Brown 2020) to reduce the variance induced by episode sampling. Building upon the approximated cumulative advantages, we integrate the novel tabular variants DCFR+ and PDCFR+ into neural CFR. For DCFR+, we apply discounting and clipping to the cumulative advantages at each iteration, enabling efficient action reuse and controlling the potential overgrowth of cumulative advantages over iterations. For PDCFR+, we further introduce an additional network to fit the advantages at each iteration, which is then used to predict the next advantages and compute the new strategy. Experimental results demonstrate that these algorithms achieve competitive performance compared to other model-free neural methods.

Preliminaries

Notations

Extensive-form games (Osborne and Rubinstein 1994) provide a tree-based formalism widely used to describe IIGs. These games involve a finite set $\mathcal{N} = \{1, 2, \dots, N\}$ of **play-**

ers, along with a special player c called **chance** following a fixed, known stochastic strategy. A **history** h represents a sequence of all actions taken by players, including any private information available only to specific players. The set of all possible histories forms \mathcal{H} , while $\mathcal{Z} \subseteq \mathcal{H}$ denotes **terminal histories** where no further actions are possible. The set of terminal histories that can be reached from a history h is denoted by $\mathcal{Z}[h] = \{z \in \mathcal{Z} : h \sqsubseteq z\}$, where the relationship $g \sqsubseteq h$ indicates that g is equal to or a **prefix** of h . At any given history h , players choose from the **actions** available, represented by $\mathcal{A}(h) = \{a : ha \in \mathcal{H}\}$. The player making the decision at history h is denoted by $\mathcal{P}(h)$. Each player $i \in \mathcal{N}$ is associated with a **utility function** $u_i(z) : \mathcal{Z} \rightarrow \mathbb{R}$, which assigns a utility to each terminal history $z \in \mathcal{Z}$.

In IIGs, the lack of information is modeled using **information sets** \mathcal{I}_i for each player $i \in \mathcal{N}$, where $h, h' \in \mathcal{I}$ indicates that player i cannot distinguish between them based on the information available. For instance, in poker, histories in the same information set differ only in opponents’ private cards. Hence, $\mathcal{A}(\mathcal{I}) = \mathcal{A}(h)$ and $\mathcal{P}(\mathcal{I}) = \mathcal{P}(h)$ for any $h \in \mathcal{I}$. The set of all terminal histories that can be reached from an information set \mathcal{I} is denoted by $\mathcal{Z}[\mathcal{I}] = \bigcup_{h \in \mathcal{I}} \mathcal{Z}[h]$, and $z[\mathcal{I}]$ is the unique history $h \in \mathcal{I}$ such that $h \sqsubseteq z$.

A **strategy** $\sigma_i(\mathcal{I})$ assigns a probability distribution over actions $\mathcal{A}(\mathcal{I})$ available to player i in information set \mathcal{I} , where $\sigma_i(\mathcal{I}, a)$ represents the probability of player i choosing action a in \mathcal{I} . Similarly, the strategies must be consistent across all histories in an information set. Thus, for any $h_1, h_2 \in \mathcal{I}$, we have $\sigma_i(\mathcal{I}) = \sigma_i(h_1) = \sigma_i(h_2)$. A **strategy profile** $\sigma = \{\sigma_i \mid \sigma_i \in \Sigma_i, i \in \mathcal{N}\}$ specifies the strategies for all players, where Σ_i denotes the set of all possible strategies for player i , and σ_{-i} refers to the strategies of other players.

The **history reach probability** $\pi^\sigma(h)$ is the joint probability of reaching history h under strategy profile σ , computed as $\pi^\sigma(h) = \prod_{h' \sqsubseteq h} \sigma_{\mathcal{P}(h')}(h', a)$. It factorizes as $\pi^\sigma(h) = \pi_i^\sigma(h) \pi_{-i}^\sigma(h)$, where $\pi_i^\sigma(h)$ is player i ’s contribution, and $\pi_{-i}^\sigma(h)$ is the contributions of all other players. The **information set reach probability** $\pi^\sigma(\mathcal{I})$ is defined as $\pi^\sigma(\mathcal{I}) = \sum_{h \in \mathcal{I}} \pi^\sigma(h)$, and the **interval information set reach probability** from h to h' is defined as $\pi^\sigma(h, h') = \pi^\sigma(h) / \pi^\sigma(h')$ if $h' \sqsubseteq h$. $\pi_i^\sigma(\mathcal{I}), \pi_{-i}^\sigma(\mathcal{I}), \pi_i^\sigma(h, h'), \pi_{-i}^\sigma(h, h')$ are defined similarly.

The **expected utility** $u_i(\sigma_i, \sigma_{-i})$ for player i denotes her utility when playing σ_i against opponents’ strategy σ_{-i} . Formally, $u_i(\sigma_i, \sigma_{-i}) = \sum_{z \in \mathcal{Z}} \pi^\sigma(z) u_i(z)$. At the level of an information set \mathcal{I} , the expected utility for taking action a is

$$u_i^\sigma(\mathcal{I}, a) = \frac{\sum_{h \in \mathcal{I}} \pi^\sigma(h) \sum_{z \in \mathcal{Z}[ha]} \pi^\sigma(ha, z) u_i(z)}{\sum_{h \in \mathcal{I}} \pi^\sigma(h)},$$

and for the whole information set: $u_i^\sigma(I) = \sum_{a \in \mathcal{A}(I)} \sigma_i(I, a) u_i^\sigma(I, a)$. The **advantage** $A_i^\sigma(I, a) = u_i^\sigma(I, a) - u_i^\sigma(I)$ then quantifies the utility gain of playing action a instead of the current strategy $\sigma_i(I, a)$.

Best Response and Nash Equilibrium

The **best response** to σ_{-i} is any strategy $\text{BR}(\sigma_{-i})$ that maximizes the expected utility, satisfying $u_i(\text{BR}(\sigma_{-i}), \sigma_{-i}) = \max_{\sigma'_i \in \Sigma_i} u_i(\sigma'_i, \sigma_{-i})$. An **NE** is a strategy profile $\sigma^* =$

$(\sigma_i^*, \sigma_{-i}^*)$ where each player plays a best response to the others, ensuring $\forall i \in \mathcal{N}, u_i(\sigma_i^*, \sigma_{-i}^*) = \max_{\sigma_i' \in \Sigma_i} u_i(\sigma_i', \sigma_{-i}^*)$. The **exploitability** of a strategy σ_i is defined as $e_i(\sigma_i) = u_i(\sigma_i^*, \sigma_{-i}^*) - u_i(\sigma_i, \text{BR}(\sigma_i))$. In an ϵ -NE, no player's exploitability exceeds ϵ . The exploitability of a strategy profile σ , given by $e(\sigma) = \sum_{i \in \mathcal{N}} e_i(\sigma_i)/|\mathcal{N}|$, represents the approximation error relative to the NE.

Counterfactual Regret Minimization

Counterfactual Regret Minimization (CFR) (Zinkevich et al. 2007) frequently uses **counterfactual value** $v_i^\sigma(I, a)$, which is the expected utility of an action a at an information set $I \in \mathcal{I}_i$ for player i , weighted by the probability that i reaches I . Formally, $v_i^\sigma(I, a) = \sum_{z \in Z[I]} \pi_{-i}^\sigma(z[I]) \pi^\sigma(z[I]a, z) u_i(z)$, and $v_i^\sigma(I) = \sum_{a \in \mathcal{A}(I)} \sigma_i(I, a) v_i^\sigma(I, a)$. Starting from a uniform strategy σ^1 , CFR traverses the game tree in each iteration t to compute the **instantaneous counterfactual regret** $r_i^t(I, a) = v_i^{\sigma^t}(I, a) - v_i^{\sigma^t}(I)$, which relates to the advantages as $r_i^t(I, a) = A_i^t(I, a) \pi_{-i}^{\sigma^t}(I)$ (Srinivasan et al. 2018), accumulates it into **cumulative counterfactual regret** $R_i^t(I, a) = \sum_{k=1}^t r_i^k(I, a)$, and updates strategies by regret-matching (Hart and Mas-Colell 2000): $\sigma_i^{t+1}(I, a) = \frac{\max(0, R_i^t(I, a))}{\sum_{a' \in \mathcal{A}(I)} \max(0, R_i^t(I, a'))}$, using the uniform strategy when all regrets are non-positive. The **average strategy** $\bar{\sigma}^t$ converges to an NE and is computed via **cumulative strategy** $C_i^t(I, a)$:

$$C_i^t(I, a) = \sum_{k=1}^t \left(\pi_i^{\sigma^k}(I) \sigma_i^k(I, a) \right), \bar{\sigma}_i^t(I, a) = \frac{C_i^t(I, a)}{\sum_{a' \in \mathcal{A}(I)} C_i^t(I, a')}.$$

Tabular CFR Variants

Since the birth of CFR, numerous variants have been proposed to accelerate convergence. **CFR+** (Tammelin 2014; Bowling et al. 2015) improves convergence by: (1) clipping negative cumulative counterfactual regrets: $R_i^t(I, a) = \max(R_i^{t-1}(I, a) + r_i^t(I, a), 0)$; (2) using a linear weighted average strategy: $C_i^t(I, a) = C_i^{t-1}(I, a) + t \pi_i^{\sigma^t}(I) \sigma_i^t(I, a)$; (3) applying alternating updates. **DCFR** (Brown and Sandholm 2019a) further introduces discounting:

$$R_i^t(I, a) = \begin{cases} R_i^{t-1}(I, a) \frac{(t-1)^\alpha}{(t-1)^\alpha + 1} + r_i^t(I, a), & \text{if } R_i^{t-1}(I, a) > 0 \\ R_i^{t-1}(I, a) \frac{(t-1)^\beta}{(t-1)^\beta + 1} + r_i^t(I, a), & \text{otherwise,} \end{cases}$$

$$C_i^t(I, a) = C_i^{t-1}(I, a) \left(\frac{t-1}{t} \right)^\gamma + \pi_i^{\sigma^t}(I) \sigma_i^t(I, a).$$

LinearCFR is a special case of DCFR, with updates $R_i^t(I, a) = R_i^{t-1}(I, a) + t r_i^t(I, a)$, $C_i^t(I, a) = C_i^{t-1}(I, a) + t \pi_i^{\sigma^t}(I) \sigma_i^t(I, a)$. **DCFR+** (Hang et al. 2022, 2024) integrates CFR+ and DCFR: $R_i^t(I, a) = \max\left(R_i^{t-1}(I, a) \frac{(t-1)^\alpha}{(t-1)^\alpha + 1} + r_i^t(I, a), 0\right)$. **PCFR+** (Farina, Kroer, and Sandholm 2021) follows CFR+ for updating cumulative counterfactual regrets and predicts the next iteration's cumulative counterfactual regrets as $\tilde{R}_i^{t+1}(I, a) = \max(R_i^t(I, a) + \tilde{r}_i^{t+1}(I, a), 0)$, where the prediction of instantaneous counterfactual regrets $\tilde{r}_i^{t+1}(I, a)$ is assumed to change slowly and is set to $r_i^t(I, a)$. Then PCFR+ uses

$\tilde{R}_i^{t+1}(I, a)$ in regret-matching to compute the next strategy. **PDCFR+** (Hang et al. 2024) updates cumulative counterfactual regrets like DCFR+ and predicts the next iteration's cumulative counterfactual regrets as $\tilde{R}_i^{t+1}(I, a) = \max(R_i^t(I, a) \frac{t^\alpha}{t^\alpha + 1} + \tilde{r}_i^{t+1}(I, a), 0)$. For the cumulative strategy, DCFR+, PCFR+, and PDCFR+ follow the same formula as DCFR but differ in their choices of γ .

Monte Carlo CFR

The tabular CFR variants improve convergence but require full game tree traversals and tabular storage, which are infeasible in large games. To reduce time complexity, **Monte Carlo CFR** (MCCFR) (Lanctot et al. 2009) estimates instantaneous counterfactual regrets by sampling portions of the game tree. MCCFR includes **external sampling** (ES), which explores all actions for one player while sampling actions for others, and **outcome sampling** (OS), which samples actions for all players along a single episode. This work focuses on the model-free variant OS-MCCFR, which learns directly from sampled episodes without a perfect simulator. At iteration t , episodes are sampled using a **sampling strategy** ξ^t with $\xi_i^t(I, a) = \epsilon/|\mathcal{A}(I)| + (1-\epsilon)\sigma_i^t(I, a)$ and $\xi_{-i}^t(I, a) = \sigma_{-i}^t(I, a)$. The **sampled counterfactual value** $\hat{v}_i^{\sigma^t}(I, a | z)$ uses importance sampling to remain unbiased:

$$\hat{v}_i^{\sigma^t}(I, a | z) = \frac{\pi_{-i}^{\sigma^t}(z[I]) \pi^{\sigma^t}(z[I]a, z) u_i(z)}{\pi^{\xi^t}(z)} = \frac{\pi^{\sigma^t}(z[I]a, z) u_i(z)}{\pi_i^{\xi^t}(z[I]) \pi^{\xi^t}(z[I], z)}.$$

The **sampled instantaneous counterfactual regrets** is $\hat{r}_i^t(I, a) = \hat{v}_i^{\sigma^t}(I, a | z) - \hat{v}_i^{\sigma^t}(I | z)$, where $\hat{v}_i^{\sigma^t}(I | z) = \sum_{a \in \mathcal{A}(I)} \sigma_i^t(I, a) \hat{v}_i^{\sigma^t}(I, a | z)$. They are unbiased estimators of $r_i^t(I, a)$. Similarly, the **sampled strategy** $\hat{\sigma}_i^t(I, a | z)$ is defined as $\sigma_i^t(I, a) \pi_i^{\sigma^t}(I) / \pi^{\xi^t}(z)$.

DeepCFR

DeepCFR (Brown et al. 2019) uses neural networks to approximate CFR, avoiding tabular storage of cumulative counterfactual regrets and cumulative strategies. In each iteration, it performs K partial traversals, stores the sampled instantaneous counterfactual regrets in a reservoir buffer, and trains a network from scratch to predict cumulative counterfactual regrets. The expectation value of action a in an information set I after T iterations is $R_i^T(I, a) / \sum_{t=1}^T \pi^{\xi^t}(I)$, where the denominator accounts for sampling bias canceled out during regret matching. A second buffer stores strategies across iterations to approximate the average strategy. For the sampling strategy, DeepCFR adopts ES for better performance but relies on a perfect simulator, while remaining compatible with OS. Since LinearCFR's update can be rewritten as $R_i^t(I, a) = \sum_{t=1}^T t r_i^t(I, a)$, DeepCFR can approximate LinearCFR by weighing samples in iteration t by t , yielding improved performance.

Related Work

Learning an NE in IIGs by combining deep reinforcement learning and game theory algorithms has gained considerable attention in recent years. These approaches can generally be divided into three main categories. (1) Policy-Space

Algorithm 1: Training procedures for VR-DeepDCFR+ and VR-DeepPDCFR+

Input: total iterations T , traversal times K , parameters α, γ , exploration coefficient ϵ .

Initialize each player’s cumulative advantage network $R(I, a | \theta_i^0)$ with parameters θ_i^0 ;

Initialize each player’s instantaneous advantage network $r(I, a | \phi_i^0)$ with parameters ϕ_i^0 ;

Initialize history value network $Q(h, a | w^0)$ and history value buffer \mathcal{B}_Q ;

Initialize reservoir-sampled strategy buffer \mathcal{B}_Π and each player’s advantage buffer $\mathcal{B}_{V,i}$;

for CFR iteration $t = 1$ to T **do**

 clear each player’s advantage buffer $\mathcal{B}_{V,i}$;

foreach player i **do**

for traversal $k = 1$ to K **do**

\perp Traverse $(\emptyset, i, \mathcal{B}_{V,i}, \mathcal{B}_\Pi, \mathcal{B}_Q, t)$, \triangleright (Algorithm 2);

 Train θ_i^t on loss $\mathcal{L}(\theta_i^t) = \mathbb{E}_{(I, \bar{r}) \sim \mathcal{B}_{V,i}} \left[\sum_{a \in \mathcal{A}(I)} \left(\max(R(I, a | \theta_i^{t-1}), 0) \frac{(t-1)^\alpha}{(t-1)^{\alpha+1}} + \bar{r}(I, a) - R(I, a | \theta_i^t) \right)^2 \right]$;

 For VR-DeepPDCFR+, Train ϕ_i^t on loss $\mathcal{L}(\phi_i^t) = \mathbb{E}_{(I, \bar{r}) \sim \mathcal{B}_{V,i}} \left[\sum_{a \in \mathcal{A}(I)} (\bar{r}(I, a) - r(I, a | \phi_i^t))^2 \right]$;

 Train w^t on loss $\mathcal{L}(w^t) = \mathbb{E}_{(t, h, a, \hat{u}, h', I', i) \sim \mathcal{B}_Q} \left[\left(\hat{u} + \sum_{a' \in \mathcal{A}(h')} \sigma_i^{t+1}(I', a') Q(h', a' | w^t) - Q(h, a | w^t) \right)^2 \right]$;

Train ψ on loss $\mathcal{L}(\psi) = \mathbb{E}_{(I, t, \sigma^t) \sim \mathcal{B}_\Pi} \left[\left(\frac{t}{T} \right)^\gamma \sum_{a \in \mathcal{A}(I)} (\sigma^t(I, a) - \Pi(I, a | \psi))^2 \right]$;

Output: The average strategy network $\Pi(I, a | \psi)$.

Response Oracle (PSRO) (Lanctot et al. 2017) and its variants maintain a population of strategies and iteratively compute the best response to a meta-strategy. Neural Fictitious Self-Play (NFSP) (Heinrich and Silver 2016) is a special case of PSRO, where the meta-strategy is the uniform distribution over past strategies. While these methods are effective and scalable, they rely on computationally expensive approximate best response calculations, and their convergence speed is often related to the size of the strategy space. (2) Numerous neural CFR methods have been developed to approximate the behavior of tabular CFR (Brown et al. 2019; Li et al. 2020; Gruslys et al. 2020; Steinberger, Lerer, and Brown 2020; Liu, Li, and Togelius 2022; Meng et al. 2023; McAleer et al. 2023). Our approach builds on this line of work by approximating novel tabular CFR variants to further accelerate convergence. (3) Another research direction involves modifying policy gradient algorithms to enable convergence to an NE (Srinivasan et al. 2018; Lockhart et al. 2019; Hennes et al. 2020; Fu et al. 2022). However, their performance is sensitive to hyperparameters.

Deep (Predictive) Discounted CFR

Fitting Cumulative Advantages by Bootstrapping

DeepCFR can naturally integrate with LinearCFR but struggles to approximate the behaviors of more advanced CFR variants like DCFR+ and PDCFR+, which rely on cumulative counterfactual regrets from the previous iteration. A straightforward approach involves approximating instantaneous counterfactual regrets at each iteration and bootstrapping on the estimated cumulative counterfactual regrets from the prior iteration. However, counterfactual values are expected utilities weighted by opponents’ reach probabilities, which diminish significantly over long episodes. These reach probabilities vary widely across information

sets, making it challenging for networks to effectively learn values across diverse orders of magnitude (Van Hasselt et al. 2016). Furthermore, as demonstrated in Theorem 1 (with all proofs provided in Appendix A), the expected estimation are scaled by sampling reach probabilities. This results in the expectation of the estimated cumulative counterfactual regrets taking the form $\sum_{t=1}^T \left[r_i^t(I, a) / \pi^{\xi^t}(I) \right]$. Since denominators change across iterations, it leads to deviation from CFR’s behavior.

Theorem 1 *By using outcome sampling to collect data $(I, \hat{r}_i^t(I))$ into a buffer \mathcal{B}_i for player i in iteration t , and training a neural network $r(I, a | \phi_i^t)$ on loss $\mathcal{L}(\phi_i^t) = \mathbb{E}_{(I, \hat{r}_i^t(I)) \sim \mathcal{B}_i} \left[\sum_{a \in \mathcal{A}(I)} (\hat{r}_i^t(I, a) - r(I, a | \phi_i^t))^2 \right]$, the expected target value of $r(I, a | \phi_i^t)$ for any sampled information set I is given by:*

$$\mathbb{E}_{z \sim \xi^t} [\hat{r}_i^t(I, a) | z \in Z_I] = \frac{r_i^t(I, a)}{\pi^{\xi^t}(I)}.$$

To address these challenges, we adjust the calculation of the sampled counterfactual value as $\check{v}_i^{\sigma^t}(I, a | z) = \frac{\pi^{\sigma^t}(z|I)u_i(z)}{\pi^{\xi^t}(z|I)}$, and $\check{v}_i^{\sigma^t}(I | z)$ and $\check{r}_i^t(I, a)$ are defined similarly. As demonstrated in Theorem 2, these expectations correspond to the advantages of the information set I . Advancements in deep reinforcement learning have shown that neural networks excel at predicting and generalizing advantages, even in complex environments with large state spaces (Schulman et al. 2017). By computing the cumulative advantages $\check{R}^t(I, a) = \sum_{k=1}^t A_i^{\sigma^t}(I, a)$ instead of cumulative counterfactual regrets, the process can be interpreted as a type of weighted CFR since $r_i^t(I, a) = \pi_{-i}^{\sigma^t}(I) A_i^{\sigma^t}(I, a)$ (Srinivasan et al. 2018).

Algorithm 2: CFR Traversal with Outcome Sampling for VR-DeepDCFR+ and VR-DeepPDCFR+

Function `Traverse` ($h, i, \mathcal{B}_{V,i}, \mathcal{B}_{\Pi}, \mathcal{B}_Q, t$):

Input: History h , traversing player i , advantage buffer $\mathcal{B}_{V,i}$, strategy buffer \mathcal{B}_{Π} , history value buffer \mathcal{B}_Q , iteration t .
if h is terminal **then**
| **return** the utility of player i ;
For VR-DeepDCFR+, compute strategy $\sigma^t(I, a)$ from $R(I, a | \theta_{\mathcal{P}(h)}^{t-1})$ using regret matching;
For VR-DeepPDCFR+, compute strategy $\sigma^t(I, a)$ from the predicted cumulative advantages
 $\max \left(R(I, a | \theta_{\mathcal{P}(h)}^{t-1}), 0 \right) \frac{(t-1)^\alpha}{(t-1)^{\alpha+1}} + r(I, a | \phi_{\mathcal{P}(h)}^{t-1})$ using regret matching;
for $a \in \mathcal{A}(h)$ **do**
| $\xi^t(I, a) \leftarrow \epsilon_{|\mathcal{A}(h)|} + (1 - \epsilon)\sigma^t(I, a)$ if $\mathcal{P}(h) = i$ else $\sigma^t(I, a)$;
 $\hat{a} \sim \xi^t(I), h' \leftarrow h\hat{a}$;
while $\mathcal{P}(h')$ is chance **do**
| $a \sim \sigma(h'), h' \leftarrow h'a$;
 $\bar{v}(I' | z) \leftarrow \text{Traverse}(h', i, \mathcal{B}_{V,i}, \mathcal{B}_{\Pi}, \mathcal{B}_Q, t)$;
for $a \in \mathcal{A}(h)$ **do**
| $\bar{v}(I, a | z) \leftarrow Q_i(h, a | w^{t-1}) + \frac{\bar{v}(I'|z) - Q_i(h, a | w^{t-1})}{\xi^t(I, a)}$ if $a = \hat{a}$ else $Q_i(h, a | w^{t-1})$
 $\bar{v}(I | z) \leftarrow \sum_{a \in \mathcal{A}(h)} \sigma^t(I, a) \bar{v}(I, a | z)$;
if $\mathcal{P}(h) = i$ **then**
| **for** $a \in \mathcal{A}(h)$ **do**
| | $\bar{r}(I, a) \leftarrow \bar{v}(I, a) - \sum_{a' \in \mathcal{A}(I)} \sigma^t(I, a') \bar{v}(I, a')$;
| | Insert (I, \bar{r}) into the advantage buffer $\mathcal{B}_{V,i}$;
else
| Insert $(I, t, \sigma^t(I))$ into the strategy buffer \mathcal{B}_{Π} ;
Insert $(t, h, \hat{a}, \hat{u}, h', I', i)$ into the history value buffer \mathcal{B}_Q ;
return $\bar{v}(I | z)$;

Theorem 2 By using outcome sampling to collect data $(I, \check{r}_i^t(I))$ into a buffer \mathcal{B}_i for player i in iteration t , and training a neural network $r(I, a | \phi_i^t)$ on loss $\mathcal{L}(\phi_i^t) = \mathbb{E}_{(I, \check{r}_i^t(I)) \sim \mathcal{B}_i} \left[\sum_{a \in \mathcal{A}(I)} (\check{r}_i^t(I, a) - r(I, a | \phi_i^t))^2 \right]$, the expected target value of $r(I, a | \phi_i^t)$ for any sampled information set I is given by:

$$\mathbb{E}_{z \sim \xi^t} [\check{r}_i^t(I, a) | z \in Z_I] = \frac{r_i^t(I, a)}{\pi_{-i}^{\xi^t}(I)} = A_i^{\sigma^t}(I, a).$$

To address the high variance introduced by the importance sampling term $\pi_i^{\sigma^t}(I) / \pi^{\xi^t}(z)$ in the sampled strategy $\hat{\sigma}_i^t(I, a | z)$, which complicates network training, we directly use the strategy $\sigma_i^t(I, a)$ as the sampled strategy $\check{\sigma}_i^t(I, a | z)$. However, when it is player i 's turn to collect data, we save the sampled strategy $\check{\sigma}_{-i}^t(I, a | z)$ to the buffer for the opponent player $-i$. The expectation of the cumulative strategy for player $-i$ is given by $\sum_{t=1}^T \mathbb{E}_{z \in \xi^t} [\check{\sigma}_{-i}^t(I, a | z)] = \sum_{t=1}^T \pi_{-i}^{\xi^t}(I) \pi_{-i}^{\sigma^t}(I) \sigma_{-i}^t(I, a)$. This can be interpreted as a form of weighted cumulative strategy, where $\pi_{-i}^{\xi^t}(I)$ acts as the weight in iteration t .

We now describe the training procedures for the cumulative advantage and average strategy network. Similar to OS-DeepCFR, outcome sampling is used to sample K episodes

per iteration, and these experiences are added into replay buffers. For player i , the replay buffer $\mathcal{B}_{V,i}$ stores information sets I and advantage estimates $\check{r}(I, a)$, which are used to train a cumulative advantage network $R(I, a | \theta_i^t)$ to approximate cumulative advantages at a given information set. Unlike OS-DeepCFR, where the replay buffer retains samples from all iterations, the replay buffer is cleared at the start of each iteration. The network $R(I, a | \theta_i^t)$ is trained by bootstrapping according to the loss

$$\mathcal{L}(\theta_i^t) = \mathbb{E}_{(I, \check{r}) \sim \mathcal{B}_{V,i}} \left[\sum_{a \in \mathcal{A}(I)} (R(I, a | \theta_i^{t-1}) + \check{r}(I, a) - R(I, a | \theta_i^t))^2 \right].$$

Another buffer \mathcal{B}_{Π} with reservoir sampling stores information sets I , iteration numbers t , and sampling strategies $\check{\sigma}^t$. It is used to train an average network $\Pi(I, a | \psi)$ that approximates the average strategy over all iterations. The network $\Pi(I, a | \psi)$ is optimized using the following loss:

$$\mathcal{L}(\psi) = \mathbb{E}_{(I, t, \check{\sigma}^t) \sim \mathcal{B}_{\Pi}} \left[\sum_{a \in \mathcal{A}(I)} (\check{\sigma}^t(I, a) - \Pi(I, a | \psi))^2 \right].$$

Approximating Advanced CFR Variants

The estimated cumulative advantages pave the way for approximating the behaviors of DCFR+ and PDCFR+. Both methods update cumulative counterfactual regrets by applying discounting and clipping operations. This results in the

loss for the network $R(I, a | \theta_i^t)$:

$$\mathcal{L}(\theta_i^t) = \mathbb{E}_{(I, \tilde{r}) \sim \mathcal{B}_{V,i}} \left[\sum_{a \in \mathcal{A}(I)} \left(\max(R(I, a | \theta_i^{t-1}), 0) \frac{t-1}{t} + \tilde{r}(I, a) - R(I, a | \theta_i^t) \right)^2 \right],$$

where the sequence of discounting and clipping is adjusted to facilitate sampling-based approximation of the expectation $\tilde{r}(I, a)$. Since PDCFR+ relies on predicting the instantaneous counterfactual regrets for the next iteration to compute a new strategy, an instantaneous advantage network $r(I, a | \phi_i^t)$ is trained for each player i . This network estimates the instantaneous advantages in iteration t using samples from the replay buffer $\mathcal{B}_{V,i}$, with the loss

$$\mathcal{L}(\phi_i^t) = \mathbb{E}_{(I, \tilde{r}) \sim \mathcal{B}_{V,i}} \left[\sum_{a \in \mathcal{A}(I)} (\tilde{r}(I, a) - r(I, a | \phi_i^t))^2 \right].$$

The instantaneous advantage network is then used to predict the cumulative advantages for the next iteration as $\max(R(I, a | \theta_i^t), 0) \frac{t}{t\alpha+1} + r(I, a | \phi_i^t)$. For the average strategy, the cumulative strategy can be expressed as $C_i^t(I, a) = C_i^{t-1}(I, a) + t^{-\gamma} \pi_i^{\sigma^t}(I) \sigma_i^t(I, a)$. So we can train the average strategy network with the loss:

$$\mathcal{L}(\psi) = \mathbb{E}_{(I, t, \sigma^t) \sim \mathcal{B}_\Pi} \left[\left(\frac{t}{T} \right)^\gamma \sum_{a \in \mathcal{A}(I)} (\sigma^t(I, a) - \Pi(I, a | \psi))^2 \right],$$

where T is the total number of iterations.

Variance Reduction Based on Baseline Functions

To mitigate the high variance caused by episode sampling, prior works such as DREAM (Steinberger, Lerer, and Brown 2020) and ESCHER (McAleer et al. 2023) incorporate value functions at history nodes. DREAM uses value functions as baseline functions for each action, and constructing an unbiased estimator of counterfactual regrets, while ESCHER directly computes counterfactual regrets using value functions. Since ESCHER requires accurate value estimation and thus incurs significant training time, we adopt the variance reduction approach of DREAM.

For each episode z in iteration t , we extract and store a set of experience tuples $(t, h, \hat{a}, \hat{u}, h', I', i)$ in the history value buffer \mathcal{B}_Q . Each tuple represents player i taking action \hat{a} at history node h , transitioning to node h' associated with information set I' , and player 1 receiving utility \hat{u} (where $\hat{u} = u_1(h')$ if h' is a terminal history, and $\hat{u} = 0$ otherwise). The network $Q(h, a | w^t)$ estimates the value of each action for player 1 at every history node under the strategy σ^{t+1} . It is trained with the loss

$$\mathcal{L}(w^t) = \mathbb{E}_{(t, h, \hat{a}, \hat{u}, h', I', i) \sim \mathcal{B}_Q} \left[\left(\hat{u} + \sum_{a' \in \mathcal{A}(h')} \sigma_i^{t+1}(I', a') Q(h', a' | w^t) - Q(h, \hat{a} | w^t) \right)^2 \right].$$

The network is trained in an off-policy manner, eliminating the need to sample new episodes under σ^{t+1} , thereby improving learning efficiency. Since we assume the game is two-player zero-sum, we have $Q_1(h, a | w^t) = Q(h, a | w)$ and $Q_2(h, a | w^t) = -Q(h, a | w)$. It is used to compute baseline-adjusted sampled value

$$\bar{v}_i^{\sigma^t}(I, a | z) = \begin{cases} Q_i(h, a | w^{t-1}) + \frac{\bar{v}_i^{\sigma^t}(I' | z) - Q_i(h, a | w^{t-1})}{\xi^t(I, a)} & \text{if } a = \hat{a} \\ Q_i(h, a | w^{t-1}) & \text{otherwise,} \end{cases}$$

$$\bar{v}_i^{\sigma^t}(I | z) = \begin{cases} u_i(z) & \text{if } h = z \\ \sum_{a \in \mathcal{A}(h)} \sigma_i^t(I, a) \bar{v}_i^{\sigma^t}(I, a | z) & \text{otherwise.} \end{cases}$$

We then replace the sampled advantages $\tilde{r}_i^t(I, a | z)$ with baseline-adjusted sampled advantages $\bar{r}_i^t(I, a | z) = \bar{v}_i^{\sigma^t}(I, a | z) - \bar{v}_i^{\sigma^t}(I | z)$, which serve as unbiased estimators (Schmid et al. 2019). Algorithm 1 outlines the complete training procedure of our algorithms.

Experiments

In this section, we evaluate the performance of VR-DeepDCFR+ and VR-DeepPDCFR+ through extensive experiments. We first demonstrate their empirical convergence toward the NE across eight widely used IIGs in the research community. We provide detailed descriptions of the games in Appendix B. Exploitability is used as the performance metric to showcase the convergence speeds. We then conduct experiments on a large poker game. Given the large size of the game, we assess performance by playing against five agents with different styles, using average rewards as the performance metric. All testing games are sourced from OpenSpiel (Lanctot et al. 2020).

We compare our methods against five model-free neural methods: NFSP, q-based policy gradient (QPG) / regret policy gradient (RPG) (Srinivasan et al. 2018), OS-DeepCFR and DREAM. All methods use similar network architectures with three hidden layers with 64 neurons each. We normalize the utilities received by all algorithms in each game to the range $[-1, 1]$. The implementation of NFSP, QPG, and RPG are sourced from OpenSpiel, while OS-DeepCFR is adapted from ES-DeepCFR in OpenSpiel. Hyperparameters for NFSP and OS-DeepCFR follow the OpenSpiel reproduction report (Walton and Lisy 2021), while those for QPG/RPG are from (Farina and Sandholm 2021). For DREAM, we adopt the base hyperparameters of OS-DeepCFR and use a circular buffer of size 1,000,000 for the history value network. VR-DeepDCFR+ and VR-DeepPDCFR+ share same common hyperparameters with DREAM, with specific settings of $\alpha=2, \gamma=2$ for DeepDCFR+ and $\alpha=2.3, \gamma=2$ for VR-DeepPDCFR+. All algorithms use the same hyperparameters across all games. Detailed configurations are provided in Appendix C.

Convergence to Equilibrium

We run each algorithm four times with different random seeds, and the results are shown in Figure 1. In all plots, the x-axis is the number of episodes sampled by each algorithm, and the y-axis is exploitability shown on a log scale. The shaded area represents 95% confidence intervals over four random seeds. The two policy gradient algorithms QPG and RPG converge to an exploitability of 0.01 in *Kuhn Poker*, but perform poorly in more complex games, consistent with findings in the work (Farina and Sandholm 2021). Neural CFR variants generally converge faster than NFSP, mainly due to CFR's theoretically superior convergence rate. Compared to OS-DeepCFR, the algorithms proposed in this work converge faster in most games, demonstrating the effectiveness of approximating cumulative advantages. Cumulative advantages exhibit less variance than cumulative counterfactual regrets, leading to more stable neural network training and faster convergence of the average strategy.

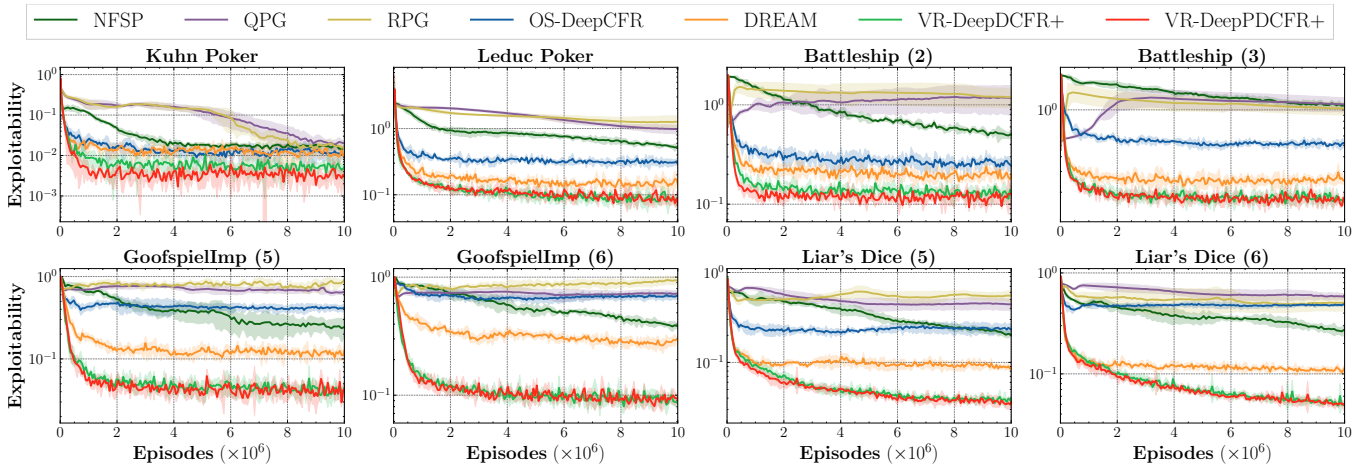


Figure 1: Convergence results of seven model-free neural algorithms on eight testing games.

Moreover, VR-DeepDCFR+ and VR-DeepPDCFR+ inherit the convergence advantages of DCFR+ and PDCFR+ over vanilla CFR and LinearCFR. Experimental results show that VR-DeepDCFR+ and VR-DeepPDCFR+ outperform OS-DeepCFR and DREAM in convergence speed across most games, highlighting the benefit of integrating neural networks with advanced CFR variants. The running time of VR-DeepDCFR+ is roughly the same as that of DREAM, since the main difference lies in their loss formulations, which incur negligible overhead. Moreover, by using bootstrapping instead of retraining from scratch, it actually requires fewer total training steps.

Head-to-Head Evaluation

We evaluate the algorithms on the large poker game flop hold'em poker (*FHP*) by playing 20,000 matches against five rule-based agents with different styles. Each agent estimates hand strength at decision points and follows predefined rules reflecting different degrees of aggressiveness, tightness, or bluffing. These agents simulate diverse exploit scenarios, allowing a multidimensional assessment of strategy robustness (Li and Miikkulainen 2018). Please refer to Appendix D for details. We use the average reward over these matches as the performance metric.

Given the poor performance of NFSP, QPG, and RPG in typical IIGs, we focus on comparing four neural CFR variants. We increase the number of sampled episodes to 10^8 , buffer size to 10^7 , neurons per layer to 128, while keeping other hyperparameters unchanged. The results are shown in Figure 2. Final average rewards are -7.8 ± 1.4 chips for OS-DeepCFR, -2.0 ± 3.1 for DREAM, 11.6 ± 1.2 for VR-DeepDCFR+, and 11.3 ± 0.9 for VR-DeepPDCFR+. Among the four methods, VR-DeepDCFR+ and VR-DeepPDCFR+ consistently outperform various rule-based agents with different styles. Notably, in professional Texas Hold'em matches, an average reward of five chips per hand is considered a significant skill gap (Moravčík et al. 2017). Therefore, compared to other neural CFR variants, the proposed methods demonstrate higher learning efficiency and superior performance in the large poker game.




Figure 2: Head-to-head evaluation results of four neural CFR variants on *FHP*.

Ablation Studies

The proposed algorithms consist of three key components: bootstrapped cumulative advantages estimation, approximating advanced CFR variants, and baseline-based variance reduction. To evaluate the impact of each component, we use VR-DeepPDCFR+ as the base method and test performance on four IIGs after removing each module individually. Experimental results show that all three components contribute to improved learning efficiency and overall performance. Please refer to Appendix E for details.

Conclusions and Future Research

This work proposes two novel model-free neural CFR variants, VR-DeepDCFR+ and VR-DeepPDCFR+, for learning an NE in two-player zero-sum IIGs. In each iteration, the algorithms collect variance-reduced sampled advantages using a history value network, bootstrap cumulative advantages, and apply discounting and clipping to simulate the behaviors of advanced tabular CFR variants DCFR+ and PDCFR+. Experimental results demonstrate that our methods achieve faster convergence across eight widely used IIGs and obtain higher average rewards against various rule-based agents in a large poker game compared to other model-free neural algorithms. Several promising directions remain for future work. One potential avenue is to improve the prediction of instantaneous advantages in VR-DeepPDCFR+, possibly by leveraging recurrent neural networks to capture temporal dependencies more effectively (Sychrovský et al. 2024).

Acknowledgments

This work is supported in part by the National Key R&D Program of China (No. 2025ZD0122000), the Natural Science Foundation of China (Nos. 62222606 and 61902402), the Key Research and Development Program of Jiangsu Province (No. BE2023016), and the CCF-Baidu Open Fund.

References

- Bowling, M.; Burch, N.; Johanson, M.; and Tammelin, O. 2015. Heads-up limit hold'em poker is solved. *Science*, 347(6218): 145–149.
- Brown, N.; Lerer, A.; Gross, S.; and Sandholm, T. 2019. Deep counterfactual regret minimization. In *International Conference on Machine Learning*, 793–802.
- Brown, N.; and Sandholm, T. 2018. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374): 418–424.
- Brown, N.; and Sandholm, T. 2019a. Solving imperfect-information games via discounted regret minimization. In *AAAI Conference on Artificial Intelligence*, 1829–1836.
- Brown, N.; and Sandholm, T. 2019b. Superhuman AI for multiplayer poker. *Science*, 365(6456): 885–890.
- Farina, G.; Kroer, C.; and Sandholm, T. 2021. Faster game solving via predictive Blackwell approachability: Connecting regret matching and mirror descent. In *AAAI Conference on Artificial Intelligence*, 5363–5371.
- Farina, G.; and Sandholm, T. 2021. Model-free online learning in unknown sequential decision making problems and games. In *AAAI Conference on Artificial Intelligence*, 5381–5390.
- Fu, H.; Liu, W.; Wu, S.; Wang, Y.; Yang, T.; Li, K.; Xing, J.; Li, B.; Ma, B.; Fu, Q.; et al. 2022. Actor-critic policy optimization in a large-scale imperfect-information game. In *International Conference on Learning Representations*, 1–28.
- Gratch, J.; Nazari, Z.; and Johnson, E. 2016. The misrepresentation game: How to win at negotiation while seeming like a nice guy. In *International Conference on Autonomous Agents and Multiagent Systems*, 728–737.
- Gruslys, A.; Lanctot, M.; Munos, R.; Timbers, F.; Schmid, M.; Pérolat, J.; Morrill, D.; Zambaldi, V.; Lespiau, J.-B.; Schultz, J.; Azar, M. G.; Bowling, M.; and Tuyls, K. 2020. The Advantage Regret-Matching Actor-Critic. arXiv:2008.12234.
- Hang, X.; Kai, L.; Bingyun, L.; Haobo, F.; Qiang, F.; Junliang, X.; and Cheng, J. 2024. Minimizing Weighted Counterfactual Regret with Optimistic Online Mirror Descent. In *International Joint Conference on Artificial Intelligence*, 5272–5280.
- Hang, X.; Kai, L.; Haobo, F.; Qiang, F.; and Junliang, X. 2022. AutoCFR: Learning to Design Counterfactual Regret Minimization Algorithms. In *AAAI Conference on Artificial Intelligence*, 5244–5251.
- Hart, S.; and Mas-Colell, A. 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5): 1127–1150.
- Heinrich, J.; and Silver, D. 2016. Deep Reinforcement Learning from Self-Play in Imperfect-Information Games. arXiv:1603.01121.
- Hennes, D.; Morrill, D.; Omidshafiei, S.; Munos, R.; Pérolat, J.; Lanctot, M.; Gruslys, A.; Lespiau, J.-B.; Parmas, P.; Duéñez-Guzmán, E.; et al. 2020. Neural replicator dynamics: Multiagent learning via hedging policy gradients. In *International Conference on Autonomous Agents and Multiagent Systems*, 492–501.
- Lanctot, M.; Lockhart, E.; Lespiau, J.-B.; Zambaldi, V.; Upadhyay, S.; Pérolat, J.; Srinivasan, S.; Timbers, F.; Tuyls, K.; Omidshafiei, S.; Hennes, D.; Morrill, D.; Muller, P.; Ewalds, T.; Faulkner, R.; Kramár, J.; Vylder, B. D.; Saeta, B.; Bradbury, J.; Ding, D.; Borgeaud, S.; Lai, M.; Schrittwieser, J.; Anthony, T.; Hughes, E.; Danihelka, I.; and Ryan-Davis, J. 2020. OpenSpiel: A Framework for Reinforcement Learning in Games. arXiv:1908.09453.
- Lanctot, M.; Waugh, K.; Zinkevich, M.; and Bowling, M. 2009. Monte Carlo sampling for regret minimization in extensive games. In *Advances in Neural Information Processing Systems*, 1078–1086.
- Lanctot, M.; Zambaldi, V.; Gruslys, A.; Lazaridou, A.; Tuyls, K.; Pérolat, J.; Silver, D.; and Graepel, T. 2017. A unified game-theoretic approach to multiagent reinforcement learning. In *Advances in Neural Information Processing Systems*, 4193–4206.
- Li, H.; Hu, K.; Zhang, S.; Qi, Y.; and Song, L. 2020. Double Neural Counterfactual Regret Minimization. In *International Conference on Learning Representations*, 1–20.
- Li, X.; and Miikkulainen, R. 2018. Opponent modeling and exploitation in poker using evolved recurrent neural networks. In *Genetic and Evolutionary Computation Conference*, 189–196.
- Lisy, V.; Davis, T.; and Bowling, M. 2016. Counterfactual regret minimization in sequential security games. In *AAAI Conference on Artificial Intelligence*, 544–550.
- Liu, W.; Li, B.; and Togelius, J. 2022. Model-free neural counterfactual regret minimization with bootstrap learning. *IEEE Transactions on Games*, 15(3): 315–325.
- Lockhart, E.; Lanctot, M.; Pérolat, J.; Lespiau, J.-B.; Morrill, D.; Timbers, F.; and Tuyls, K. 2019. Computing approximate equilibria in sequential adversarial games by exploitability descent. In *International Joint Conference on Artificial Intelligence*, 464–470.
- McAleer, S.; Farina, G.; Lanctot, M.; and Sandholm, T. 2023. ESCHER: Eschewing importance sampling in games by computing a history value function to estimate regret. In *International Conference on Learning Representations*, 1–22.
- Meng, L.; Ge, Z.; Tian, P.; An, B.; and Gao, Y. 2023. An efficient deep reinforcement learning algorithm for solving imperfect information extensive-form games. In *AAAI Conference on Artificial Intelligence*, 5823–5831.
- Moravčík, M.; Schmid, M.; Burch, N.; Lisý, V.; Morrill, D.; Bard, N.; Davis, T.; Waugh, K.; Johanson, M.; and Bowling, M. 2017. DeepStack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, 356(6337): 508–513.

- Nash, J. J. F. 1950. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences of the United States of America*, 36(1): 48–49.
- Osborne, M. J.; and Rubinstein, A. 1994. *A course in game theory*. MIT press.
- Sandholm, T. 2015. Steering evolution strategically: Computational game theory and opponent exploitation for treatment planning, drug design, and synthetic biology. In *AAAI Conference on Artificial Intelligence*, 4057–4061.
- Schmid, M.; Burch, N.; Lanctot, M.; Moravcik, M.; Kadlec, R.; and Bowling, M. 2019. Variance reduction in Monte Carlo counterfactual regret minimization (VR-MCCFR) for extensive form games using baselines. In *AAAI Conference on Artificial Intelligence*, 2157–2164.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. arXiv:1707.06347.
- Srinivasan, S.; Lanctot, M.; Zambaldi, V.; Pérolat, J.; Tuyls, K.; Munos, R.; and Bowling, M. 2018. Actor-critic policy optimization in partially observable multiagent environments. In *Advances in Neural Information Processing Systems*, 3426–3439.
- Steinberger, E.; Lerer, A.; and Brown, N. 2020. DREAM: Deep Regret minimization with Advantage baselines and Model-free learning. arXiv:2006.10410.
- Sychrovský, D.; Šustr, M.; Davoodi, E.; Bowling, M.; Lanctot, M.; and Schmid, M. 2024. Learning not to regret. In *AAAI Conference on Artificial Intelligence*, 15202–15210.
- Tammelin, O. 2014. Solving Large Imperfect Information Games Using CFR+. arXiv:1407.5042.
- Van Hasselt, H. P.; Guez, A.; Hessel, M.; Mnih, V.; and Silver, D. 2016. Learning values across many orders of magnitude. In *Advances in Neural Information Processing Systems*, 4294–4302.
- Walton, M.; and Lisy, V. 2021. Multi-agent Reinforcement Learning in OpenSpiel: A Reproduction Report. arXiv:2103.00187.
- Zinkevich, M.; Johanson, M.; Bowling, M.; and Piccione, C. 2007. Regret minimization in games with incomplete information. In *Advances in Neural Information Processing Systems*, 1729–1736.