
Simplex Neural Population Learning: Any-Mixture Bayes-Optimality in Symmetric Zero-sum Games

Siqi Liu^{1,2} Marc Lanctot² Luke Marris^{1,2} Nicolas Heess²

Abstract

Learning to play optimally against any mixture over a diverse set of strategies is of important practical interests in competitive games. In this paper, we propose simplex-NeuPL that satisfies two desiderata *simultaneously*: i) learning a population of strategically diverse basis policies, represented by a single conditional network; ii) using the same network, learn best-responses to *any* mixture over the simplex of basis policies. We show that the resulting conditional policies incorporate prior information about their opponents effectively, enabling near optimal returns against arbitrary mixture policies in a game with tractable best-responses. We verify that such policies behave Bayes-optimally under uncertainty and offer insights in using this flexibility at test time. Finally, we offer evidence that learning best-responses to any mixture policies is an effective auxiliary task for strategic exploration, which, by itself, can lead to more performant populations.

1. Introduction

How could we train agents to perform optimally against arbitrary mixtures over diverse opponent policies? Population learning offers one potential answer: generate a diverse set of opponents and train the agent to respond to mixtures of opponents over the population. The question then becomes how the population is generated and what properties it should have. In two-player zero-sum games, there is a well-known solution to this problem based on game-theoretic foundations: a Nash equilibrium distribution (NE, Nash (1951)) over a population of policies maximizes an agent’s worst-case return against all possible opponent policies. Despite its theoretical appeal, searching the entire policy space quickly becomes intractable for

most games. To this end, empirical game-theoretic analysis (EGTA, Wellman (2006)) proposed to study strategic exploration in games by investigating empirical (meta-)games, where each player considers only a small subset of possible policies. Policy-Space Response Oracles (PSRO, Lanctot et al. (2017)), further proposed a general, iterative framework towards constructing such empirical games. At each iteration, the policy population incorporates a new basis policy that is trained to best-respond to a mixture over its predecessors, following a meta-strategy solver (MSS). Importantly, when the best-response operator is exact, certain meta-strategy solvers produce meta-strategies known to converge to an NE of the game.

One property of the NE target distribution is that it optimizes a safe objective: it maximizes the expected payoff in the *worst-case*, with the assumption that the opponents would play minimax-optimally. This assumption, however, rarely holds in practice — real-world agents could play arbitrarily far from NE, a phenomenon frequently observed among human players (Wright & Leyton-Brown, 2017), due to inadequate training or simply, to the overwhelming complexity of the game. This translates to the unfortunate situation where NE, though unexploitable, often leads to sub-optimal decision making at test time. The flexibility for players to express subjective beliefs over the opponent and to play optimally, based on such beliefs is thus of interests. We refer to this ability to play optimally against any mixture over a diverse set of policies as *any-mixture optimality*. Indeed, skilled human players are observed to resort to such flexibility when competing in games, adjusting their behaviours based on assumptions about their opponents so as to play optimally, if their assumptions prove correct (King-Casas et al., 2005; Schlicht et al., 2010).

Unfortunately, existing population learning algorithms such as PSRO precisely lack such flexibility. The choice of MSS not only controls the strategic diversity of the resulting population, but also, restricts the set of basis policies that can be executed at test time. In particular, the output of population learning is a set of best-responses to specific mixture policies, enumerated by the MSS at each iteration. Consequently, a player can only play optimally against a few sets of opponents, or forgo optimality entirely and execute

¹University College London, UK ²DeepMind, UK. Correspondence to: Siqi Liu <liusiqi@google.com>.

the NE mixture policy so as to be assured of safety, in expectation, over many games. At its extreme, a player cannot guarantee to play optimally even when the opponent uses the same population of policies and publicly declares their strategy in advance, nor can they play optimally if they wish to consider all strategies equally likely *a priori* without unduly ruling out any opponent strategy.

Our goal is therefore to extend game-theoretic population learning algorithms so as to offer *any-mixture optimality* at test time. To this end, we interpret PSRO geometrically as iteratively expanding a population simplex whose vertices correspond to the set of basis policies, each best-responding to a point within the simplex from the previous iteration (Figure 1). To instead learn best-responses to *all* points within the population simplex, we further generalise recent work on Neural Population Learning (NeuPL, Liu et al. (2022)), a general framework that incorporates principled population learning algorithms, using scalable and efficient representation for the population of policies via a single conditional neural network. The result is thus a simple extension that not only retains the efficiency and game-theoretic properties of NeuPL, but also yields a conditional policy that behaves optimally against *arbitrary* mixtures over the policy population (Section 3). Additionally, we recognize best-response solving across the population simplex as optimising a continuum of Bayes-optimal objectives (Humplik et al., 2019; Ortega et al., 2019) and demonstrate properties of the resulting policies typically associated with Bayes-optimality. In particular, we show that the resulting conditional policies effectively incorporate prior information about their opponents so as to achieve near optimal returns against arbitrary mixtures policies in a game with tractable best-response solutions (Section 4.1). We further compare different choices of policies at test time in a more complex, partially-observed, spatiotemporal strategy game and show that executing the NE mixture policy can be far from optimal whereas executing an *uninformed* policy that considers all opponent strategies equally likely *a priori* can be highly effective (Section 4.2). Lastly, we show that simplex-NeuPL is not only critical in providing *any-mixture optimality*, but also, facilitates strategic exploration by promoting transfer across best-responses to the continuum of mixture policies, leading to more performant populations at no extra costs (Section 4.3).

2. Background

2.1. Partially-Observed Stochastic Games (POSG)

Stochastic games (Shapley, 1953) generalise the basic formalism of Markov Decision Processes (MDPs) to multiple players. To model partial observability, we define a symmetric zero-sum partially-observed stochastic game (Hansen et al., 2004) by (S, O, X, A, P, R) where S defines the

state space, O the observation space and $X : S \rightarrow O \times O$ the observation function that returns partial views of the state for both players. Let $P : S \times A \times A \rightarrow \Pr(S)$ be the state transition distribution given a state and joint actions, $R : S \rightarrow \mathbb{R} \times \mathbb{R}$ the reward function defining rewards for both players in state s_t , denoted $R(s_t) = (r_t, \bar{r}_t)$. In state s_t , players act according to policies conditioned on their respective observation histories $(\pi(jo_{1:t}), \pi^{\theta}(jo_{1:t}))$. In practice, observation history can be represented as fixed-size embedding with the use of a learned recurrent neural network. Player π achieves an expected return of $J(\pi, \pi^{\theta}) = \mathbb{E}_{\pi, \pi^{\theta}}[\sum_t r_t]$ against π^{θ} . A game is said to be *symmetric* if the expected return of a policy is only dependent on the policy played by the other player, rather than the identity or order of the player. A policy π is said to best-respond to π^{θ} if $\partial_{\pi} J(\pi, \pi^{\theta}) = 0$. We note $\pi \in \text{BR}(\pi^{\theta})$, if a best-response policy against π^{θ} can be computed tractably. In practice, an exact best-response operator may be intractable computationally and we define approximate best-response (ABR) operators as $\hat{\pi} \in \text{ABR}(\pi, \pi^{\theta})$ such that $J(\hat{\pi}, \pi^{\theta}) \geq J(\pi, \pi^{\theta})$. In other words, an approximate best-response operator produces a policy $\hat{\pi}$ that performs at least as well as π against π^{θ} . It’s worth noting that we focus on POSG instead of Normal-form Games (NFGs) as our focus is on developing Bayes-optimal policies that can benefit from information gathered through sequential interactions.

2.2. Population Learning

Population Learning defines an iterative procedure for strategic exploration in games. In particular, we consider the formalism of Policy-Space Response Oracles (PSRO, Lanctot et al. (2017)) which combined EGTA with deep reinforcement learning. Given a symmetric zero-sum, partially-observed, stochastic game where each player has access to the same set of N policies $\Pi := \{\pi_i\}_{i=0}^{N-1}$, we define a normal-form empirical (meta-)game where players’ i -th action corresponds to executing policy π_i for an episode. A probability assignment $\sigma \in \mathbb{R}^N$ over the policy population therefore defines a meta-game mixture strategy, or a mixture policy in the underlying game, with $\sigma \in \mathbb{R}^{j \times j-1}$ representing the space of j j -dimensional distributions, or the volume of a $(j \times j - 1)$ -simplex. When executing a meta-game mixture strategy, an action of the meta-game, or a policy in the underlying game, is sampled at the start of each episode, following σ . The definition of (approximate) best-response readily extends to mixture policies, with $J(\pi, \sigma) = \mathbb{E}_{\pi, \sigma}[\sum_t r_t]$. We further define the empirical payoff matrix $U \in \mathbb{R}^{j \times j \times j \times j} = \text{EVAL}(\sigma)$ with $U_{ij} := J(\pi_i, \pi_j)$ the payoff of the i -th meta-game pure-strategy when playing against the j -th. We further recall the definition of meta-strategy solver (MSS) $f : \mathbb{R}^{j \times j \times j \times j} \rightarrow \mathbb{R}^{j \times j-1}$ which derives a meta-game mixture

strategy from the empirical payoff matrix U . PSRO thus defines the following iterative procedure: at the i -th iteration, $\pi_i = \text{ABR}(\pi, \sigma_{i-1})$ is introduced to the policy population, with $\sigma_{i-1} = f(U)$ and π a randomly initialised policy. Starting from an arbitrary initial population π_0 (typically a singleton $f(\pi_0)$), PSRO proceeds iteratively, until the ABR operator fails to achieve a strictly positive payoff at that iteration. Finally, we note that while any MSS can be used, specific ones offer appealing properties. In particular, when NE is used as the meta-strategy solution, PSRO is known to converge to a NE of the *full* game. We refer to this implementation as PSRO-NASH for short.

Neural Population Learning NeuPL (Liu et al., 2022) differs from the iterative population learning procedure described thus far in two ways. First, the population of policies is represented using a shared conditional network $\sigma_i = f(j_{t-1}, \sigma_i); \sigma_i \in \mathcal{G}$, with $\mathcal{G} = \{f_{\sigma_i} \mid \sigma_i \in \mathcal{G}\}$ representing the adjacency matrix of an interaction graph (Garnelo et al., 2021) or equivalently, a set of meta-game mixture strategies. Second, the optimisation of the policy population proceeds concurrently, with policy (j_{t-1}, σ_i) maximising its expected returns against a mixture policy π_i , defined over the neural population itself. We note that π_i corresponds to a *single* neural network which is shared across all policies within the neural population. Extending MSS from the iterative case, NeuPL updates the sequence of meta-strategies to best-respond to concurrently, using a meta-graph solver (MGS) F , with $F(U)$. Importantly, NeuPL is known to converge to a NE when a NE meta-strategy solver is applied iteratively, with $\sigma_i = \text{SOLVE-NE}(U_{<i, <i})$. This replicates the strategic exploration dynamics of PSRO-NASH. While any MGS could be used in NeuPL, we restrict our discussions to PSRO (i.e. lower-triangular \mathcal{G}) for the rest of this work given its game-theoretic properties. Finally, we note that in contrast to the iterative setting, the *effective* population size is defined by the number of unique rows within \mathcal{G} , with $j \text{UNIQUEROWS}(\mathcal{G}) \leq N$. The effective population size is therefore driven by the MGS used, adjusted dynamically through time based on the empirical payoff matrix U .

3. Methods

Population Simplex We now introduce the *population simplex* to serve as a geometric interpretation for population learning and motivates our proposed extension. A population simplex is defined by a set of N policies $\sigma_i = f_{\pi_i} g_{i=0}^{N-1}$ where each policy corresponds to a vertex of a simplex Δ^{N-1} . Their mixture policies thus span the volume of the simplex with σ a barycentric coordinate within the simplex. Analogous to PSRO, the set of vertices are selected iteratively: i) given σ , a MSS selects a coordinate $\sigma_{i-1} = f(\text{EVAL}(\sigma))$ and; ii) given σ_{i-1} , a BR op-

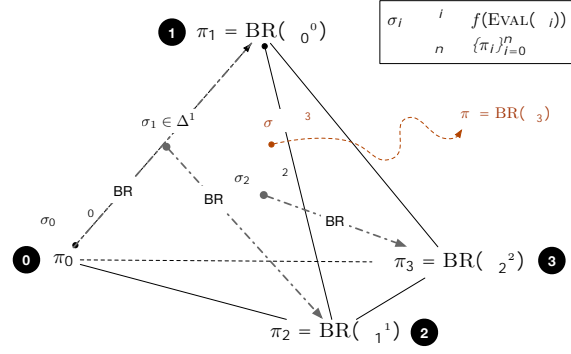


Figure 1. An iteratively expanding population where π_i best-responds to σ_{i-1} . Dashed arrows correspond to best-response operations. The solid lines are edges of the simplex and the vertices correspond to the set of basis policies. Each point $\sigma \in \Delta^3$ realises a mixture policy σ and admits its best-response policy $\text{BR}(\sigma)$ (shown in brown).

erator proposes a new vertex $\pi^0 = \text{BR}(\sigma)$ that joins the existing population simplex to form an extended simplex Δ^{j+1} . At each iteration, PSRO expands an Δ^{j-1} to Δ^j if $J(\pi^0, \sigma) > 0$ with $\sigma = \text{EVAL}(\sigma)$, if not, this iterative process terminates. This process is visualised in Figure 1, developing a sequence of policies iteratively starting from $\sigma = f_{\pi_0} g$, forming a population 3-simplex.

This geometric interpretation of population clarifies the interplay between the MSS and the BR operator — for a given population simplex, one could in principle compute a best-response to each point within the simplex, developing infinitely many candidate policies that can be added to the population. Nevertheless, such procedure has been infeasible computationally, as best-response solving often comes at significant computational cost even for a single policy. A feasible solution is thus to rely on a meta-strategy solver. A MSS proposes a specific point within the simplex worth best-responding to, directing computational resources efficiently. This process forgoes optimal returns for all but a few select points for which the population offers best-responses, but yields population-level desiderata such as convergence to the NE (McMahan et al., 2003) or maximal exploration of the policy space (Balduzzi et al., 2019).

Simplex Neural Population Learning Our proposal, is therefore to generalise NeuPL to additionally and concurrently optimise best-responses to *all* mixtures within the population simplex. Specifically, we utilize NeuPL to produce a set of basis policies and recognize that \mathcal{G} within the population simplex, we can optimise (j_{t-1}, σ) to maximise its expected returns against the mixture policy π_i . This leads to Algorithm 1 where in addition to optimising the discrete set of conditional policies $\sigma_i = f_{\pi_i} g_{i=1}^N$ as in NeuPL, we also optimise best-responses to any mixture

policies σ , with probability ϵ , where the opponent prior σ is sampled according to a symmetric Dirichlet distribution with equal concentration α assigned to each unique policies (i.e. $\text{UNIQUEROWS}(\cdot)$) of the neural population. In other words, we sample mixture opponent policies uniformly over the population simplex, with support over the set of unique policies in the population simplex. We denote the concentration parameters as α_{\leq} to indicate that $j \text{UNIQUEROWS}(\cdot) j \leq N$ and $\text{UNIF}(\cdot)$ refers to sampling one distribution from a set of probability distributions uniformly at random. This procedure is illustrated in Algorithm 1.

At convergence, simplex-NeuPL leads to a conditional policy from which one may construct not only all mixture policies within the population simplex, but also, their Bayes-optimal responses $(j \circ t, \sigma)$. In subsequent analyses, we refer to $(j \circ t, \sigma)$ as the *uninformed* policy with σ the uniform distribution (i.e. uninformative opponent prior) and $(j \circ t, \sigma)$ the *informed* policy as it is conditioned on the (often privileged) opponent sampling distribution.

Algorithm 1 Simplex Neural Population Learning

```

1:  $(j \circ t, \sigma) \leftarrow \text{CONDNEUPL}(j, t)$   $\triangleright$  Conditional neural population net.
2:  $F := f_{\sigma_i} g_{i=0}^{N-1}$   $\triangleright$  Initial interaction graph.
3:  $F : \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}^N \times \mathbb{R}^N$   $\triangleright$  Meta-graph solver.
4: while continuing do
5:  $\sigma \leftarrow \text{CONDNEUPL}(j, t, \sigma)$   $\triangleright$  Neural population.
6: simplex-sampling  $\sigma \leftarrow \text{BERN}(\epsilon)$   $\triangleright$  With prob.  $\epsilon$ .
7: if simplex-sampling then
8:  $\sigma \leftarrow \text{DIRICHLET}(\alpha_{\leq})$   $\triangleright$  From simplex.
9: else
10:  $\sigma \leftarrow \text{UNIF}(\text{UNIQUEROWS}(\cdot))$   $\triangleright$  From MGS.
11:  $(j \circ t, \sigma) \leftarrow \text{ABR}((j \circ t, \sigma), \sigma)$ 
12:  $U \leftarrow \text{EVAL}(\cdot; \cdot)$   $\triangleright$  (Optional) if  $F$  adaptive.
13:  $F(U)$   $\triangleright$  (Optional) if  $F$  adaptive.
14: return  $(j \circ t, \sigma)$ 
    
```

4. Results

We experiment with simplex-NeuPL across two domains. First, we study the imperfect-information game of *goofspiel* which remains amenable to analytical posterior inference and exact best-response solving (Section 4.1). Second, we explore the partially-observed, spatiotemporal strategy game of *running-with-scissors*, where information-seeking actions and observation history representation are critical in inferring opponent strategies and therefore, winning the game. The policy space of the latter is significantly larger and we seek to understand the trade-offs involved in the choice of policies at test time as well as the effect of unseen opponents on implicit posterior inference (Section 4.2). Throughout all experiments, we follow PSRO-NASH where

an off-the-shelf Nash solver is used at each iteration of the meta-game solving, as in Liu et al. (2022). The specific implementation of $F_{\text{PSRO-N}}$ used as well as further details of our specific experimental setup are described in Appendix B.

4.1. Goofspiel

The game of *goofspiel* is a symmetric zero-sum bidding card game where players spend bid cards to collect points from a deck of point cards. In particular, we focus on the imperfect information variant of this game, with 5 point cards, revealed deterministically in descending order¹. Players do not observe the bidding card played by its opponent, but only the win-loss history of each point card. This game has long been subject to game theoretic analyses with well-known strategic cycles (Ross, 1971; Rhoads & Bartholdi, 2012). An example game is shown in Figure 2 (Left) where player 2 wins the game by conceding the highest value point card but guarantees wins of all remaining point cards.

In this section, we empirically investigate the effect of simplex-NeuPL in this domain. First, we show that simplex-NeuPL preserves game-theoretic strategic exploration as in Liu et al. (2022). This is expected, as any-mixture optimality implies that the resulting conditional policy $(j \circ t, \sigma)$ can best-respond to the subset of mixture policies $\sigma = f_{\sigma_i} g_{i=0}^{N-1}$ recommended by the meta-graph solver. Second, we verify that more generally, the resulting *informed* policy approaches Bayes-optimality facing *any* mixture policies. This is a novel and important property for generalisation, as it allows for incorporating a wide range of prior beliefs at test time, as opposed to sampling from a small set of best-response policies enumerated during training. Lastly, we show that the resulting model exhibits implicit posterior inference over opponent identities through interaction. This echoes prior works showing that meta-learning over a range of tasks induces Bayes-optimal behaviours (Mikulik et al., 2020; Ortega et al., 2019).

Game-Theoretic Strategic Exploration Strategic exploration in EGTA is typically expressed as iteratively solving for best-responses to mixture strategies of the induced empirical game, which are a subset of all mixture policies over the population of policies. Figure 2 (Middle) illustrates an example neural population learned by simplex-NeuPL, where each policy is optimised to best-respond to the NE over its predecessors as in PSRO-NASH. Specifically, the i -th policy $(j \circ t, \sigma_i)$ is optimised to best-respond to a mixture over $\sigma_{<i} = f_{\sigma_j} g_{j=1}^{i-1}$, following the meta-strategy NE solver given their pairwise payoff matrix $U_{<i, <j}$. The initial policy of this policy population is fixed to play bid cards in random order with a known best-response that plays bid cards in descending order, spending

¹Implementation from OpenSpiel, see Appendix A.1.1




Figure 2. (Left) an example game of *goofspiel* showing 5 point cards revealed in descending order and the two players each playing their bidding cards in (hidden) order. In this game player 1 wins the first point card but loses all subsequent points and ends up losing the game; (Middle) a neural population of strategically diverse policies optimised by simplex-NeuPL, following PSRO-NASH; (Right) average return obtained by the exact best-response policy (blue), informed policy $\theta(jo_t, \sigma)$ (red), uninformed policy $\theta(jo_t, \sigma)$ (cyan) and the empirical NE mixture policy (orange) evaluated against 6 sets of opponent mixture policies $\mathcal{F}_{\theta, \sigma}^{i, k}; \sigma_{i, k} \sim \text{DIR}(\alpha_k) \mathcal{G}_{i=1}^{256}; \alpha_k \mathcal{G}_{k=1}^6$. The k -th opponent set features mixture policies whose mixture distributions are of a certain level of entropy (denoted $H(\sigma)$), sampled from a symmetric Dirichlet distribution of concentration α_k . Each point shows the average return against one of the opponent sets.

bid cards matching the point card at each turn (Ross, 1971). Indeed, (jo_t, σ_1) solely focuses on best-responding to the initial random policy and implements this deterministic point-matching policy. In turn, (jo_t, σ_2) seeks to best-respond to (jo_t, σ_1) , sacrificing the highest value point card in exchange for all remaining points. This recovers the known optimal policy against the point-matching. We further illustrate the set of policies implemented by the neural population in Appendix A.1.2. Extending NeuPL, we confirm that simplex-NeuPL similarly accommodates principled population learning algorithms such as PSRO-NASH, exploring the policy space of the game strategically.

Any-Mixture Bayes-Optimality We now verify empirically that our proposed extension leads to a conditional policy that can best-respond to *any* mixture policies supported by the neural population, by simply conditioning on the prior distribution over opponent identities. To this end, we sample arbitrary prior distributions σ over the simplex from symmetric Dirichlet distributions and evaluate the expected returns achieved by our method and several baselines against the same mixture policies $\mathcal{F}_{\theta, \sigma}$. We compare **i)** the *informed* policy (jo_t, σ) , conditioned on the true prior σ ; **ii)** the *uninformed* policy (jo_t, σ) , conditioned on an uninformative uniform prior σ ; **iii)** an exact best-response policy solved analytically as well as **iv)** the empirical NE mixture policy $\pi_{\text{NE}}^{\text{NE}}$ with $\sigma_{\text{NE}} = \text{SOLVE-NE}(U)$. Figure 2 (Right) illustrates the result of this comparison, categorised by the levels of uncertainty present in the sampled priors. Notably, the conditional policy (jo_t, σ) performs optimally against sampled mixture policies over the simplex, with its expected return approaching that of the exact best-response solution. Interestingly, we show that the uninformed policy performs markedly worse though the gap narrows as the true priors themselves become less informative with increased entropy. Last but not the least,

we note that $\pi_{\text{NE}}^{\text{NE}}$ achieved similar returns as (jo_t, σ) . This makes intuitive sense, as neither policy incorporates prior belief over the opponent distribution. Further details on the experimental setup, including the sampling of prior distributions and exact best-response solving are described in Appendix A.1.3.

Implicit Bayesian Inference of Opponent Strategies For the instance of *goofspiel* that we are considering it is possible to compute the posterior distributions over opponent identities analytically, given a prior distribution, a set of policies and an observation history. Consider a policy population $\pi = \mathcal{F}_{\theta, \sigma}(jo_t, \sigma_i) \mathcal{G}_{i=0}^N$ with $o_t = (a_t, w_t)$ where a_t denotes the private bid card played by the player at time t , w_t the publicly observed binary win-loss of the previous point card, σ_i the identity of the player’s policy, σ_j the opponent identity at play, $\Pr(\sigma_j)$ the prior over the opponent identities, a_t^j the unobserved action taken by the opponent at time t . The posterior distribution over the opponent policy σ_j can be computed by

$$\Pr(\sigma_j | jo_t) = \frac{\Pr(o_t | \sigma_j) \Pr(\sigma_j)}{\sum_{\hat{\sigma}_j} \Pr(o_t | \hat{\sigma}_j) \Pr(\hat{\sigma}_j)}$$

with:

$$\Pr(o_t | \sigma_j) = \sum_{a_{<t}^j} \left[\prod_{k=0}^{t-1} \pi_{\theta}(a_k^j | o_{<k}^j, \sigma_j) \Pr(w_{k+1} | a_k, a_k^j) \right]$$

where the sum is over all legal unobserved opponent action sequences $a_{<t}^j$. Specifically, the set of legal action sequences corresponds to all $5!$ permutations of the bidding cards. Note that as the underlying state transition function is deterministic, $\Pr(w_{k+1} | a_k, a_k^j)$ reduces to the binary indicator on the consistency between the pair of bidding card at step k and the publicly observed win-loss at step $k + 1$. Intuitively, the analytical posterior distribution considers the likelihood of all possible unobserved opponent action sequences that are consistent with the public win-loss history.

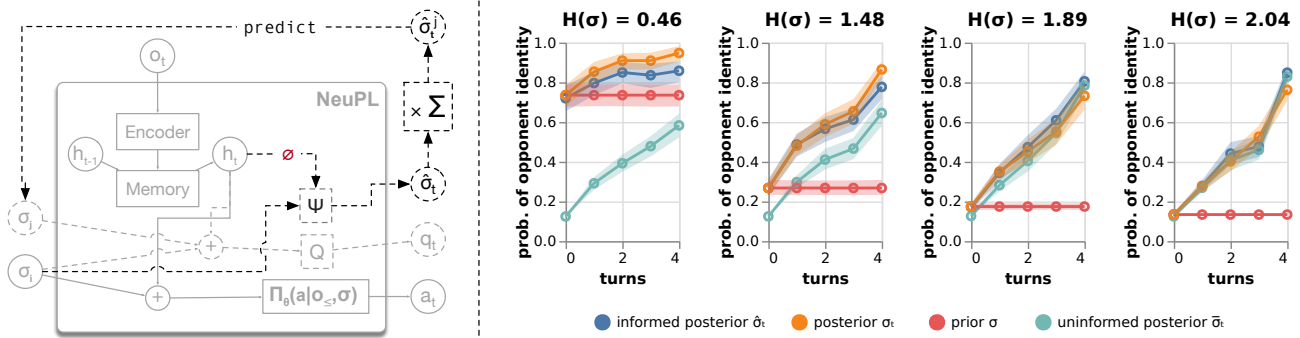


Figure 3. (Left) The network architecture of simplex-NeuPL with a posterior readout head. The baseline NeuPL architecture is shown in gray as in Liu et al. (2022) and is used unmodified in this work; (Right) The probability assigned to the true opponent identity under different distributions, including σ the prior distribution, σ_t the analytical posterior distribution, $\hat{\sigma}_t$ the “implicit” posterior distribution inferred by the posterior readout head and $\bar{\sigma}_t$ the “implicit” posterior distribution inferred with an uninformative uniform prior. The probability assignment is shown across “turns” (i.e. no card has been played at turn 0).

To verify that the observation history representation implicitly performs inference of the opponent identity, we introduced a posterior readout network $\hat{\sigma}_t = (\sigma_j, ?(h_t))$ that is optimised to infer the true opponent identity σ_j given a prior distribution σ_j , as shown in Figure 3 (Left). The readout head is parameterised by a MLP network and outputs a probability assignment $\hat{\sigma}_t$ over the set of meta-game strategies Σ . We report the output $\hat{\sigma}_t$ as the “implicit” posterior distribution at time t . Note that the readout head is optimised with a stop-gradient operator (shown as $?$) and has no influence on the representation of the policy network. The readout head is optimised alongside the agent during training, with an auxiliary regression loss.

Figure 3 (Right) illustrates the Bayesian posterior update that occurs implicitly through interaction within an episode. In particular, we report the probability assigned to the ground truth opponent identity under several distributions, grouped by different levels of entropy in the priors sampled from symmetric Dirichlet distributions. We make the following observations. First, both the informed implicit posterior and analytical posterior assign identical probability to the opponent identity as the prior distribution at turn 0, implying that the posterior readout correctly incorporates prior information about the opponent before any interaction. Similarly, the uninformed implicit posterior, $\bar{\sigma}_0 := (\sigma, ?(h_0))$ reliably assigns a probability consistent with that of a uniform distribution. As the entropy of the prior distribution increases and becomes less informative, all distributions converge to the same probability assignment at turn 0. Second, the implicit posterior distribution $\hat{\sigma}_t$ closely follows that of the analytical posterior σ_t and tends to correctly assign higher probability to the true opponent policy as the policy gathers more evidence about its opponent. We note that the implicit posterior need not exactly reproduce

its analytical counterpart, as an accurate prediction of the opponent identity is only needed if doing so improves the expected return of the policy. Lastly, we note that remarkably, an uninformed policy quickly catches up to its informed counterpart purely through interaction with the opponent, making increasingly accurate posterior inference about the opponent at play. We emphasise that representation of observation history, predictive of opponent identities, is solely a result of population learning. We visualise such implicit posterior inference in action in Appendix A.1.4.

4.2. Running-with-Scissors

The game of *running-with-scissors* extends *rock-paper-scissors* to the spatiotemporal setting with two competing players collecting *rock*, *paper* and *scissors* such that one’s inventory would compare favorably to its opponent’s during the final confrontation. The game is partially-observed, with each player observing a small 4x4 grid in front of itself at each time. The two players confront each other when the episode terminates after 500 timesteps, or when one player *tags* its opponent during a close encounter. This game has been studied in-depth in prior works (Vezhnevets et al., 2020; Liu et al., 2022), revealing a range of interesting strategies and the importance of inferring opponent strategies through interaction in this game. In this section, we hope to understand the empirical *test-time* benefits of any-mixture optimality as enabled by simplex-NeuPL. Further descriptions of the game is available in Appendix A.2.1.

We first study the scenario where both players have access to the same population of policies and illustrate that executing the NE mixture policy π^{NE} is far from optimal when evaluated against arbitrarily sampled mixture policies. Instead, executing an uninformed policy that infer opponent strat-




Figure 4. (Left) the meta-graph representing the sequence of opponent mixtures to best-respond following NE; (Middle) the pairwise payoff matrix between the population of 8 policies and their NE mixture σ^{NE} ; (Right) the expected payoffs of the NE mixture policy, uniform mixture policy, the informed policy $\theta(jo_t, \sigma)$ and the uninformed policy $\theta(jo_t, \sigma)$ evaluated against mixture opponent policies σ with σ from Dirichlet distributions of different levels of concentration.

egy dynamically through interaction could be surprisingly effective. We then turn to the setting where the opponent, or “column-player”, has access to a distinct, concealed population of policies. We investigate the mechanism of implicit posterior inference as implemented by the uninformed policy, showing that out-of-distribution opponent policies can be embedded in terms of strategically similar policies in one’s own policy population.

Test-time Policy Selection Figure 4 (Left, Middle) visualise a neural population of 8 policies, optimised via simplex-NeuPL following PSRO-NASH—the population represents a sequence of iterative best-responses, exhibiting several strategic cycles and does not admit a single dominant policy, as evidenced by its NE equilibrium σ_{NE} . Given this population of policies, Figure 4 (Right) compares the expected returns of several polices that can be constructed using the same conditional network $\theta(jo_t, \sigma)$, including the NE mixture policy, the uniform mixture policy, $\theta(jo_t, \sigma)$ the policy informed of the (privileged) opponent mixture distribution and $\theta(jo_t, \sigma)$, its uninformed counterpart. We make several observations. First, both uniform mixture policy and NE mixture policy significantly under-performed their adaptive counterparts, as they must sample and commit to a specific best-response policy without interaction that may or may not perform well against the specific opponent at play. Nevertheless, executing the NE mixture policy remains preferable, as it cannot lose to any mixture policies supported by the policy population and achieves a positive return in expectation against arbitrarily sampled mixture opponents. The uniform-conditioned policy performs surprisingly well. In contrast to Figure 2, the significant gap between σ_{NE} and the uninformed policy reflects the nature of this temporally-extended game: unlike *goofspiel* where each action has direct implication on the final return, *running-with-scissors* allows players to interact

without committing to a specific strategy early in games. The ability to infer opponent identities through interaction is thus attractive in real-world games, many of which afford extended spatiotemporal structure.

Held-out Opponent Representation We have thus far restricted ourselves to the simplified setting where the “column-player” has access to the same population of policies as the “row-player”. A more realistic setting, however, is one where the “column-player” has its own, concealed population of policies. In this setting, we study the behaviour of posterior inference as implemented by the uninformed row-player policy $\theta_r(jo_t, \sigma)$, facing held-out column-player policies $f_c(jo_t, \sigma_f)g_{i=0}^N$.

Figure 5 (Left) visualise the pairwise strategic divergence matrix D between the two populations of policies that share the same initial policy but are otherwise optimised independently, following PSRO-NASH. To measure behavioural similarity, we evaluate the expected Jensen-Shannon divergence between pairs of policies with element $D_{ij} = \mathbb{E}_{o_t, \tau^i} [D_{\text{JS}}[\theta_r(jo_t, \sigma_i^r) | \theta_c(jo_t, \sigma_j^c)]]$ where τ^i is the observation history distribution following $\theta_r(jo_t, \sigma_i^r)$, playing against the uniform mixture policy of the column-player. We highlight the following observations on the two populations of policies. First, both populations developed the three pure-resource policies, echoing (Liu et al., 2022), as reflected by the comparatively reduced strategic divergence along the diagonal elements. This trend becomes less pronounced beyond the initial pure-resource policies, as the sequence of best-responses diverge across the two populations due to approximation in best-response solving. Figure 5 (Right) illustrates the effect of inferred implicit posterior δ_t in the presence of unknown opponents: for the j -th column-player policy, the δ_t -weighted divergence $\delta_t^T D_{[\cdot, j]}$ decreases early on (shown in solid), as the

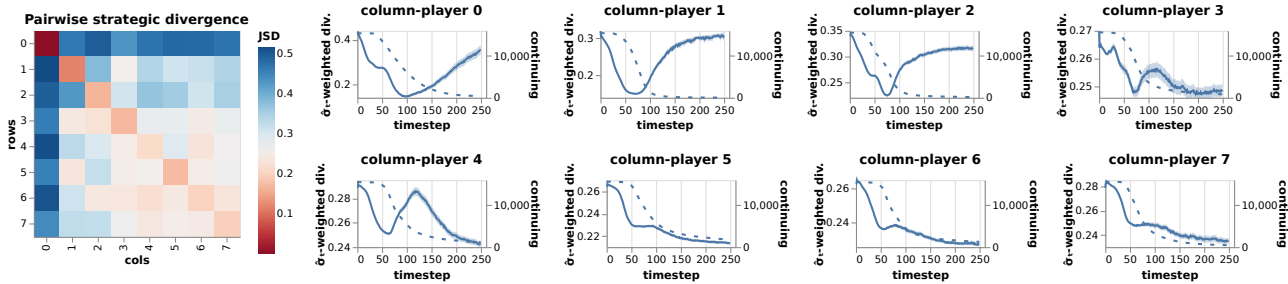


Figure 5. (Left) Pairwise strategic divergence between row and column player’s policies measured in Jensen-Shannon divergence; (Right) Implicit posterior weighted strategic divergence inferred by an uninformed row-player policy when facing each column-player policy through time (solid) and the number of continuing episodes remaining due to early termination (dashed).

row-player policy represents the column-player policies implicitly in terms of policies similar to that of its own policy population. As an increasing number of episodes early-terminate (due to players tagging each other, shown in dashed), the policy tend to struggle to identify the opposing strategies in the small number of continuing episodes.

4.3. Ablation Studies

We investigate the impact of simplex-sampling in Algorithm 1. In particular, we are interested in the two extremes: compared to NeuPL, is simplex-sampling necessary for generalisation to any-mixture optimality; and conversely, is simplex-sampling, by itself, sufficient for strategic exploration. To answer these questions, we compare the performance achieved by the uninformed policy

$(j_{\theta}(\cdot, \sigma))$ when different simplex-sampling frequency ϵ is used when training neural populations of up to 8 policies, evaluated against a uniform mixture of 8 held-out opponents in *running-with-scissors*. Figure 6 (Left) reveals the importance of simplex-sampling in this setting. In particular, the uninformed policy fails to generalise optimally without sampling from the simplex during training. At the other extreme, sampling from the simplex *alone* significantly underperforms, too. We hypothesise that this is due to insufficient best-response learning, with few samples contributing towards learning strategically relevant policies following the MSS. Figure 6 (Right) echoes this observation by evaluating populations of policies in relative terms (measured in Relative Population Performance, see Appendix A.2.2) against the same held-out opponent population. Interestingly, it shows that simplex-sampling can improve population-level performance by concurrently optimising a generalised class of best-response policies. This result corresponds to a generalised statement on knowledge transfer discussed in Liu et al. (2022) — while NeuPL enabled transfers across policies each best-responding to a specific mixture policy, simplex-NeuPL encourages transfers across best-responses to *any* mixtures.

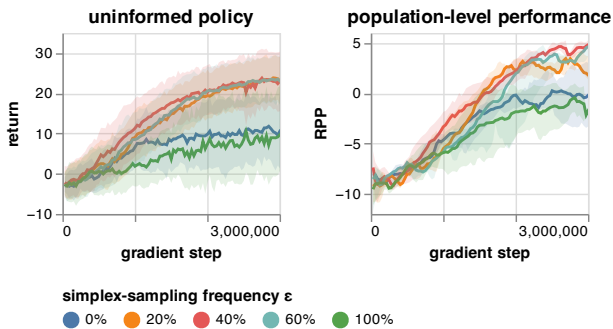


Figure 6. (Left) expected returns of uninformed policies $\theta(j_{\theta}(\cdot, \sigma))$ evaluated against a held-out uniform mixture of opponent policies; (Right) relative population performance evaluated against the same held-out opponent population. Each configuration is repeated over 3 trials.

5. Related Works

The geometric interpretation of mixed-strategies representation as points within a simplex dates as far back as 1951 when Nash Equilibria were shown to exist in finite games, using Brouwer’s fixed-point theorem defined over any closed, bounded convex sets (Nash, 1951). Our work proposes to consider a more dynamic view of policy population simplex — one that is iteratively expanded with vertices corresponding to best-responses to specific points of the population simplex at the preceding iteration. Building on recent works on efficient representation of best-responses to mixed-strategies (Liu et al., 2022), our proposal renders it computationally feasible to learn best-responses to *any* points within the population simplex.

From the perspective of iterative game-solving, several prior works can be synergistic with simplex-NeuPL, too. For instance, performing well against *any* mixtures over a set of policies has been previously studied in Smith et al.

(2020a;b), where approximate best-responses to mixture policies are constructed by combining Q-values of best-responses to individual mixture components. While the combined policy does not optimise a Bayes-optimal objective directly, such policies may prove beneficial when used as behaviour priors so as to accelerate best-response learning to any-mixture opponents. Motivated by similar observations as ours, Wu et al. (2021) proposed to study few-shot adaptation to diverse opponents via gradient-based meta-learning. By comparison, our proposed method yields a policy that adapts Bayes-optimally to a range of opponent priors, without further test-time gradient-based adaptation. While our work builds on prior works that suggest best-responding to mixed-strategies according to specific MSSs based on game-theoretic solution concepts (Lancot et al., 2017; McMahan et al., 2003), other approaches have been recently proposed to use a learned MSS optimised for alternative population-level objectives (Feng et al., 2021). McAleer et al. (2022) further proposed to consider an *expanded* restricted game such that the population as a whole is guaranteed to exhibit monotonically decreasing exploitability across best-response iterations. We leave further investigation in combining these ideas with simplex-NeuPL, a general tool towards learning any-mixture best-responses, to future works.

On the other end of the spectrum, a rich body of prior work have also been dedicated towards understanding the role of representing uncertainties from the perspective of a single agent, similar to the role of opponent prior conditioning in our work. Such uncertainties may arise due to partial-observability of the underlying environment dynamics (Zintgraf et al., 2019) or due to the presence of other agents in the environment (Vezhnevets et al., 2020; Raileanu et al., 2018; Zheng et al., 2018). A common tool that cuts across these works is the explicit representation of latent variables designed to explain the variations in the observing player’s observation history, including those caused by other interacting players. These works convincingly showed the need to reason about uncertainties in the environment explicitly, especially when interacting with other agents under partial-observability. Our learning method for each policy in the population is more closely aligned with the framework of Bayesian multi-task RL, where the policy learns to infer the underlying environment dynamics simply by optimising its expected returns on a distribution over tasks, without an explicitly designed latent variables representing such uncertainties. Consistent with prior empirical and theoretical studies (Humplik et al., 2019; Ortega et al., 2019; Mikulik et al., 2020), we show that the learned observation history representation is predictive of the opponent identity, simply as a result of return maximization, without an auxiliary prediction task that encourages it to do so.

6. Conclusions

In this work, we interpret population learning geometrically and recognizes its connections to any-mixture Bayes-optimality. By learning best-responses to the entire population simplex, we obtain a conditional policy that can not only execute arbitrary policies within the simplex, but also, their Bayes-optimal responses. Empirically, we show that the resulting conditional policies are capable of incorporating a wide range of prior beliefs about the opponent, yielding near optimal returns against arbitrary mixture policies. Importantly, we show that for real-world games, an uninformed policy can be surprisingly effective, as it exploits the temporal structure of the game and optimally trades-off exploration and exploitation under uncertainty.

Acknowledgements

We would like to thank Stephen McAleer, Ian Gemp and Karl Tuyls for helpful comments and suggestions on an early draft of this work. John Schultz and Edward Lockhart for their support and continued contributions to the OpenSpiel library which made various game-theoretic analysis possible in this work. Finally, we would like to thank the anonymous ICML reviewers for their thoughtful and constructive feedback on our work which help improved our submission significantly.

References

- Abdolmaleki, A., Springenberg, J. T., Tassa, Y., Munos, R., Heess, N., and Riedmiller, M. Maximum a posteriori policy optimisation. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=S1ANxQWob>.
- Balduzzi, D., Garnelo, M., Bachrach, Y., Czarnecki, W., Perolat, J., Jaderberg, M., and Graepel, T. Open-ended learning in symmetric zero-sum games. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 434–443. PMLR, 09–15 Jun 2019. URL <http://proceedings.mlr.press/v97/balduzzi19a.html>.
- Feng, X., Slumbers, O., Wan, Z., Liu, B., McAleer, S., Wen, Y., Wang, J., and Yang, Y. Neural auto-curricula in two-player zero-sum games. *Advances in Neural Information Processing Systems*, 34, 2021.
- Garnelo, M., Czarnecki, W. M., Liu, S., Tirumala, D., Oh, J., Gidel, G., van Hasselt, H., and Balduzzi, D. Pick your battles: Interaction graphs as population-level objectives for strategic diversity. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent*

- Systems*, AAMAS '21, pp. 1501–1503, Richland, SC, 2021. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450383073.
- Hansen, E. A., Bernstein, D. S., and Zilberstein, S. Dynamic programming for partially observable stochastic games. In *Proceedings of the 19th National Conference on Artificial Intelligence*, AAAI'04, pp. 709–715. AAAI Press, 2004. ISBN 0262511835.
- Humplik, J., Galashov, A., Hasenclever, L., Ortega, P. A., Teh, Y. W., and Heess, N. Meta reinforcement learning as task inference. *arXiv preprint arXiv:1905.06424*, 2019.
- King-Casas, B., Tomlin, D., Anen, C., Camerer, C. F., Quartz, S. R., and Montague, P. R. Getting to know you: Reputation and trust in a two-person economic exchange. *Science*, 308(5718):78–83, 2005. doi: 10.1126/science.1108062. URL <https://www.sciencemag.org/doi/abs/10.1126/science.1108062>.
- Lanctot, M., Zambaldi, V., Gruslys, A., Lazaridou, A., Tuyls, K., Perolat, J., Silver, D., and Graepel, T. A unified game-theoretic approach to multiagent reinforcement learning. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/3323fe11e9595c09af38fe67567a9394-Paper.pdf>.
- 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. OpenSpiel: A framework for reinforcement learning in games. *CoRR*, abs/1908.09453, 2019. URL <http://arxiv.org/abs/1908.09453>.
- Liu, S., Marris, L., Hennes, D., Merel, J., Heess, N., and Graepel, T. NeuPL: Neural population learning. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=MI_X3fJkl_1.
- McAleer, S., Wang, K., Lanier, J., Lanctot, M., Baldi, P., Sandholm, T., and Fox, R. Anytime PSRO for two-player zero-sum games. In *Reinforcement Learning in Games workshop (RLG @ AAAI)*, 2022.
- McMahan, H. B., Gordon, G. J., and Blum, A. Planning in the presence of cost functions controlled by an adversary. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning*, ICML'03, pp. 536–543. AAAI Press, 2003. ISBN 978-1-57735-189-4.
- Mikulik, V., Delétang, G., McGrath, T., Genewein, T., Martic, M., Legg, S., and Ortega, P. Meta-trained agents implement bayes-optimal agents. *Advances in neural information processing systems*, 33:18691–18703, 2020.
- Nash, J. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951. ISSN 0003486X. URL <http://www.jstor.org/stable/1969529>.
- Ortega, P. A., Wang, J. X., Rowland, M., Genewein, T., Kurth-Nelson, Z., Pascanu, R., Heess, N., Veness, J., Pritzel, A., Sprechmann, P., Jayakumar, S. M., McGrath, T., Miller, K., Azar, M., Osband, I., Rabinowitz, N., György, A., Chiappa, S., Osindero, S., Teh, Y. W., van Hasselt, H., de Freitas, N., Botvinick, M., and Legg, S. Meta-learning of sequential strategies. *arXiv:1905.03030 [cs, stat]*, 2019. URL <http://arxiv.org/abs/1905.03030>.
- Raileanu, R., Denton, E., Szlam, A., and Fergus, R. Modeling others using oneself in multi-agent reinforcement learning. In *International conference on machine learning*, pp. 4257–4266. PMLR, 2018.
- Rhoads, G. C. and Bartholdi, L. Computer solution to the game of pure strategy. *Games*, 3(4):150–156, 2012.
- Ross, S. M. Goofspiel—the game of pure strategy. *Journal of Applied Probability*, 8(3):621–625, 1971.
- Schlicht, E. J., Shimojo, S., Camerer, C. F., Battaglia, P., and Nakayama, K. Human wagering behavior depends on opponents' faces. *PloS one*, 5(7):e11663, 2010.
- Shapley, L. S. Stochastic games. *Proceedings of the National Academy of Sciences*, 39(10):1095–1100, 1953. ISSN 0027-8424. doi: 10.1073/pnas.39.10.1095. URL <https://www.pnas.org/content/39/10/1095>. Publisher: National Academy of Sciences. eprint: <https://www.pnas.org/content/39/10/1095.full.pdf>.
- Shoham, Y. and Leyton-Brown, K. Multiagent systems: Algorithmic, game-theoretic, and logical foundations. In *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*, chapter 4. Cambridge University Press, 2008.
- Smith, M., Anthony, T., and Wellman, M. Iterative empirical game solving via single policy best response. In *International Conference on Learning Representations*, 2020a.

- Smith, M. O., Anthony, T., Wang, Y., and Wellman, M. P. Learning to play against any mixture of opponents. *arXiv preprint arXiv:2009.14180*, 2020b.
- Vezhnevets, A., Wu, Y., Eckstein, M., Leblond, R., and Leibo, J. Z. OPTIONS as RESPONSES: Grounding behavioural hierarchies in multi-agent reinforcement learning. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 9733–9742. PMLR, 13–18 Jul 2020. URL <http://proceedings.mlr.press/v119/vezhnevets20a.html>.
- Wellman, M. P. Methods for empirical game-theoretic analysis. In *AAAI*, pp. 1552–1556, 2006.
- Wright, J. R. and Leyton-Brown, K. Predicting human behavior in unrepeated, simultaneous-move games. *Games and Economic Behavior*, 106:16–37, 2017.
- Wu, Z., Li, K., Zhao, E., Xu, H., Zhang, M., Fu, H., An, B., and Xing, J. L2e: Learning to exploit your opponent. *arXiv preprint arXiv:2102.09381*, 2021.
- Zheng, Y., Meng, Z., Hao, J., Zhang, Z., Yang, T., and Fan, C. A deep bayesian policy reuse approach against non-stationary agents. *Advances in neural information processing systems*, 31, 2018.
- Zintgraf, L., Shiarlis, K., Igl, M., Schulze, S., Gal, Y., Hofmann, K., and Whiteson, S. Varibad: A very good method for bayes-adaptive deep rl via meta-learning. In *International Conference on Learning Representations*, 2019.

A. Results

A.1. Goofspiel

A.1.1. ENVIRONMENT SETTINGS

The specific implementation of the game is available as part of OpenSpiel (Lanctot et al., 2019), instantiated with the following game string:

```
goofspiel (i mp_i nfo=true,
           egocentric=True,
           num_cards=5,
           points_order=descending,
           returns_type=point_difference))
```

We consider the imperfect information, two-player zero-sum of *goofspiel* with 5 point cards revealed in descending order. At turn t , each player observes o_t consisting of the revealed point card p_t , its action history $a_{<t} = (a_0, a_1, \dots, a_{t-1})$ as well as the win, loss and draw history $w_{<t} (w_i \in \{1, 0, 1\})$ of previously revealed point cards. The player then plays one of the remaining bidding cards a_t from its hand. We denote all valid action history $A_{<t}$, corresponding to all permutations of t cards from the initial hand of 5 bidding cards for each player.

A.1.2. STRATEGIC EXPLORATION

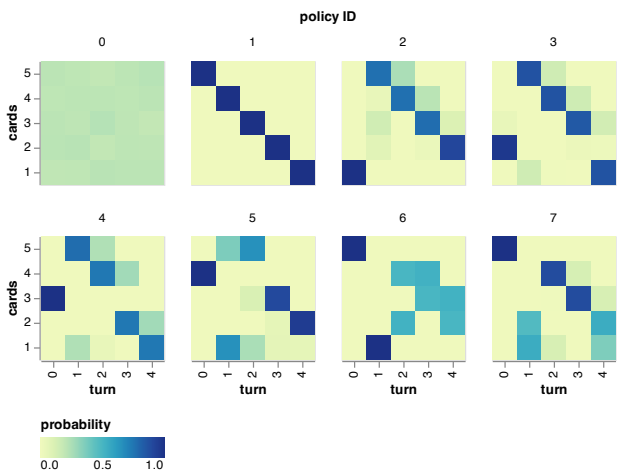


Figure 7. Policy profiles of a learned neural population $f_{\theta}(j_{o_t}, \sigma_i), \sigma_i \in \mathcal{G}$ in the game of *goofspiel*. The initial policy $\theta(j_{o_t}, \sigma_0)$ is fixed to play randomly. We simulate each pair of policies (with replacement) for 32 games and average the probability of playing each card at each turn, across all such pairs.

We visualise the set of policies optimised via simplex-NeuPL in the game of *goofspiel* described in Appendix A.1.1 in Figure 7, following PSRO-NASH. We note that the policies are conditioned on observation-history and the action profiles visualised are averaged across all pairwise matches.

A.1.3. ANY-MIXTURE OPTIMALITY

To establish the expected returns of different policies against arbitrarily opponent mixture policies, we instantiate with 256 σ sampled from DIRICHLET distributions across 7 levels of concentrations. This led to 7 sets of prior distributions with different levels of entropy $H(\sigma)$, ranging from 0.46 to 2.03. Note that a uniform distribution over a population of 8 policies corresponds to an entropy of 2.08. Given the set of sampled mixture policies, we evaluate each candidate policy against each sampled mixture policy for 32 episodes, yielding expected returns as reported in Figure 2.

For exact best-response solving in *goofspiel*, we resort to the open-source implementation of policy aggregator `open_spiel/python/algorithms/policy_aggregator.py` and `open_spiel/algorithms/best_response.h` from OpenSpiel (Lanctot et al., 2019) to compute a best-response policy against each of the sampled mixture policies.

A.1.4. IMPLICIT POSTERIOR INFERENCE IN ACTION

Figure 8 visualises 8 example episodes where the uninformed policy (j_{o_t}, σ) plays against each opponent in the policy population $f_{\theta}(j_{o_t}, \sigma_i)_{i=0}^7$. We note that across all opponents, the implicit posterior as inferred from the internal observation-history representation of the agent is predictive of the true opponent identity, with increasing accuracy as the game progresses. We recall that players do not observe the actual bidding cards played by the opponent, but rather, the win-loss history of past point cards. This contributes to the uncertainty in opponent inference, in addition to the stochasticity present in the policy itself. Finally, we note that at turn 0, the implicit posterior readout reproduces the uniform prior, as expected.

A.2. Running-with-Scissors

A.2.1. ENVIRONMENT

Figure 9 visualises an example scene of *running-with-scissors* during initialisation. The two players are randomly positioned and oriented in the game at the beginning of an episode with a 4x4 partial views of its surroundings. Each player maintains its own inventory, representing the proportion of resources in their possession, composed of *rock*, *paper* or *scissors* as visualised in coloured blocks. During a close encounter, each player may choose to tag its opponent, highlighting a small region in front of itself. If the opponent falls within the highlighted area, then the game resolves and the two players compare their inventories and receive rewards according to the *rock-paper-scissors* rules.

Similar to Liu et al. (2022), we observed a range of strategic plays as explored by the neural population. This includes a policy that *observe-and-exploit* (shown as (j_{o_t}, σ_3)

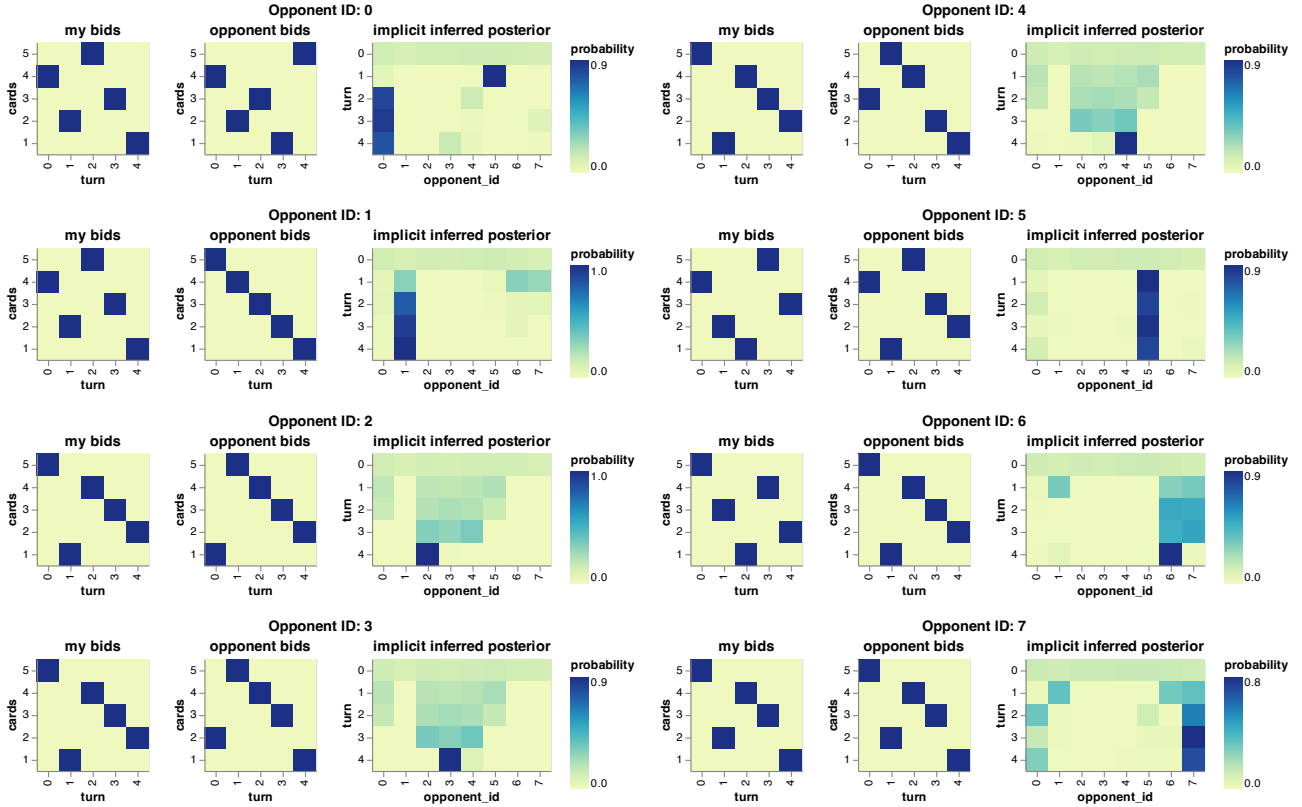


Figure 8. Example episodes where the uninformed policy $\theta(j_o, \tau, \sigma)$ plays against each one of the opponent policy. We visualise the bids from both players as well as the posterior inference from the first player’s perspective. Note that players do not directly observe their opponents’ previous bids but only observe the binary win-loss of previous point cards.




Figure 9. An example frame of the *running-with-scissors* environment. Figure from Liu et al. (2022).

in Figure 4) as well as a policy that plays deceptively to counter this policy.

A.2.2. RELATIVE POPULATION PERFORMANCE

Relative Population Performance (RPP, Balduzzi et al. (2019)) compares the performance of two populations in

relative terms. Specifically, it computes the expected returns of one population versus another, if both parties play their respective NE mixture policies. We recall the formal definition of RPP.

Definition A.1 (Relative Population Performance). Given two populations of policies $\mathfrak{B}, \mathfrak{D}$ and let (\mathbf{p}, \mathbf{q}) be a Nash equilibrium over the zero-sum game on $\mathcal{U}_{\mathfrak{B}, \mathfrak{D}} \subseteq \mathbb{R}^M \times \mathbb{R}^N$, Relative Population Performance measures their relative performance: $v(\mathfrak{B}, \mathfrak{D}) := \mathbf{p}^T \mathcal{U}_{\mathfrak{B}, \mathfrak{D}} \mathbf{q}$.

B. Experimental Setup

B.1. Agent Architecture

The high-level agent network architecture is visualised in Figure 3 which remains identical to that of Liu et al. (2022) besides the introduction of the posterior readout head. We recall that the readout head does not affect the representation learning of the main RL agent. Across both domains, we used the same MPO agent (Abdolmaleki et al., 2018) as in Liu et al. (2022), with 20 action samples drawn and evaluated by the learned Q-function per state

at each gradient update. The target (Q-value and policy) networks are updated every 100 gradient steps. The policy head and Q-value networks are parameterised by MLPs of (512, 256, 128, NUMACTIONS) and (512, 512, 128, 1) respectively with ELU activation.

We describe domain-specific configuration below.

B.1.1. *Goofspiel*

In the game of *goofspiel*, we use a feed-forward agent without a memory component. This is possible because at each step, the environment observation contains the entire observation history for each player with perfect recall, rendering a recurrent network unnecessary. Instead of a recurrent network, we used a simple MLP network with 512 neurons for the Memory module. The encoder consists of another MLP network of size (128, 64), encoding the observation-history into a fixed size vector at each step. As is common in OpenSpiel (Lanctot et al., 2019), the observation primarily consists of binary indicators of past events.

B.1.2. *running-with-scissors*

In *running-with-scissors*, the observation consists of a 4x4 pixel grid at each timestep as well as the player’s current inventory. For observation-history encoding, we used a small convolutional network with a kernel shapes of (1, 1) and 6 output channels followed by a MLP network of (64, 64) dimensions. The pixel embedding is then concatenated with the inventory information and fed into another encoder parameterised by a 2-layer MLP network of size (256, 256). This final per-timestep representation is then provided to a recurrent LSTM network, with a hidden size of 512. We note that the memory component is critically important in this game, as the observation at each timestep provides little information about the environment and the opponent.

B.2. Neural Population Configuration

We used a maximum population size of 8 across all our experiments. We used an off the shelf linear program solver for SOLVE-NE and the iterative NE solving required by the meta-graph solver requires around 50ms to complete on a desktop CPU. In *goofspiel* we invoke the meta-graph solver every 10,000 gradient updates to allow time for the underlying RL agent to optimise. In *running-with-scissors*, we update the meta-graph every 1,000 gradient updates. We did not extensively adjust these hyper-parameters during our experimentation.

B.3. Meta-Graph Solver $F_{\text{PSRO-N}}$

As in Liu et al. (2022), we iteratively apply SOLVE-NE (Shoham & Leyton-Brown, 2008) to the sub-payoff matrices as the subsequent opponent mixture policy to best-respond

Algorithm 2 MGS implementing PSRO-NASH.

```

1: function  $F_{\text{PSRO-N}}(U)$   $\triangleright U \in \mathbb{R}^{N \times N}$ .
2:   Initialize  $U$  with zeros.
3:   for  $i \in \{1, \dots, N\}$  do
4:      $U_{1:i;1:i} \leftarrow \text{SOLVE-NASH}(U_{1:i;1:i})$ 
5:   return

```

to. This procedure is described in Algorithm 2.

B.4. Training Setup

Across both domains, we used a single TPU-v2 both to perform gradient updates for neural population of policies and to serve their inference requests during simulation. The game simulation is then performed on 256 remote CPU actors for *running-with-scissors* and 128 for *goofspiel*. Across both domains the neural population seem to have converged after 2M gradient updates (shared across all population members). This corresponds to about 1-days wall clock time for *goofspiel* and 3-days for *running-with-scissors*. During training, actors generate and write experience data to a replay server, with a maximum buffer size of 100,000 trajectories across both domains. Data are sampled uniformly from the replay server, without further prioritisation.