No-Press Diplomacy from Scratch

Anonymous Author(s)
Affiliation
Address
email

Abstract

Prior AI successes in complex games have largely focused on settings with at most hundreds of actions at each decision point. In contrast, Diplomacy is a game with more than 10^{20} possible actions per turn. Previous attempts to address games with large branching factor, such as Diplomacy, StarCraft and Dota, used human data to bootstrap the policy or used handcrafted reward shaping. In this paper we describe an algorithm for action exploration and equilibrium approximation in games with combinatorial action spaces. This algorithm simultaneously performs value iteration RL while learning a policy proposal network. A double oracle step is used to explore additional actions to add to the policy proposals. At each state, the target state value and policy for RL training are computed via an equilibrium search procedure. Using this algorithm, we train an agent, *DORA*, completely from scratch for a popular two-player variant of Diplomacy and show that it achieves superhuman performance. Additionally, we extend our methods to full-scale no-press Diplomacy and for the first time train an agent from scratch with no human data. We show this agent differs radically from past agents that required human data.

1 Introduction

Classic multi-agent AI research domains such as chess, go, and poker feature action spaces with at most thousands of actions per state, and are therefore possible to explore exhaustively a few steps deep. In contrast, modern multi-agent AI benchmarks such as StarCraft, Dota, and Diplomacy feature combinatorial action spaces that are incredibly large. In StarCraft and Diplomacy, this challenge has so far been addressed by bootstrapping from human data [28, 21, 2, 9]. In Dota, this challenge was addressed through careful expert-designed reward shaping [3].

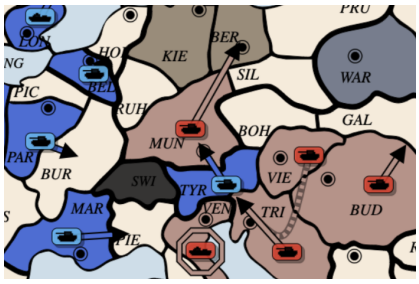

Figure 1: France (blue) would be in a strong position were it not for exactly two actions that Austria (red) can play to dislodge the unit in Tyr. The probability of sampling one of them randomly is on the order of 10^{-6}.

In contrast, in this paper we describe an algorithm that trains agents through self play without using human data and can accommodate the large action space of Diplomacy, a long-standing AI benchmark game in which the number of legal actions for a player on a turn is often more than 10^{20} [21].

We propose a form of deep value iteration RL where at each visited state the agent computes a Nash equilibrium for the stage game with 1-step lookahead search, plays the equilibrium policy and uses the equilibrium value as the training target. Since the full action space is too large to consider, this equilibrium is computed over a smaller set of candidate actions. These actions are sampled from an action proposal network that is trained in unison with the value model. The proposal actions are augmented with exploratory actions chosen through a double-oracle-like

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

process, and if these discovered actions end up as part of the equilibrium then they will be learned by the proposal network.

Using this technique, which we call *DORA* (Double Oracle Reinforcement learning for Action exploration), we develop a superhuman agent for a popular two-player variant of Diplomacy. Our agent is trained purely through self-play with no human data and no reward shaping. We also train a *DORA* agent for the 7-player game of no-press Diplomacy, the first competitive no-press Diplomacy agent to be trained entirely from scratch with no human data. When one *DORA* agent plays with 6 human-data-based agents, the bot performs at a reasonable human level. When 6 *DORA* agents play with 1 human-data-based agent, all human-data-based agents perform extremely poorly. This suggests that, unlike prior agents that had to be trained with human data, *DORA* is able to approximate equilibria that differ radically from how humans play. Our approach therefore opens up the ability to investigate novel ways of playing the game. Further highlighting the effectiveness of our methods, we show that our deep Nash learning algorithm, when initialized with a human-bootstrapped model, defeats one of the prior SOTA human-derived agents in no-press Diplomacy.

2 Background and Related Work

2.1 Description of Diplomacy

In this section we summarize the rules of Diplomacy. For a more detailed description, see [21]. No-press Diplomacy is a zero-sum board game where 7 players compete for control on a map of Europe. The map contains 75 locations, 34 of which are *supply centers* (SCs). Each player begins with 3-4 units and 3-4 SCs and a player wins by control a majority (18) of the SCs. If no player controls a majority and all remaining players agree to a draw, then the victory is divided between the uneliminated players. In case of a draw, we use the **Sum-of-Squares (SoS)** scoring system that defines the score of player i as $C_i^2 / \sum_{i'} C_{i'}^2$, where C_i is the SC count for player i.

Each turn during movement, all players simultaneously choose an action composed of one order for every unit they control. An order may direct a unit to hold stationary, to move to an adjacent location, to support a nearby unit's hold or move, to move across water via convoy, or to act as the convoy for another unit. Because a player may control up to 17 units with an average of 26 valid orders for each unit, the number of possible actions is typically too large to enumerate. Additionally, in Diplomacy, often supports or other critical orders may be indistinguishable from a no-op except in precise combinations with other orders or opponent actions. This makes the action space challenging to learn with black-box RL. Actions are simultaneous, so players typically must play a mixed strategy to avoid being predictable and exploitable.

In the popular *France vs. Austria (FvA)* variant of Diplomacy, the number of players is reduced to two, but the victory condition and all game rules remain the same. FvA serves as an effective way to isolate the challenge of an astronomically large imperfect-information action space from the challenge of Diplomacy's multi-agent nature.

2.2 Stochastic Games

Stochastic games [22] are a multi-agent generalization of Markov decision processes. At each state, each player chooses an action and the players' joint action determines the transition to the next state. Formally, at each state $s \in S$ agents simultaneously choose a joint action $\boldsymbol{a} = (a_1, \ldots, a_{\mathcal{N}}) \in A_1 \times \cdots \times A_{\mathcal{N}}$, where $\mathcal{N}$ is the number of players. Agents then receive a reward according to a vector function $\boldsymbol{r}(s, \boldsymbol{a}) \in \mathbb{R}^{\mathcal{N}}$ and transition to a new state $s' = f(s, \boldsymbol{a})$. The game starts at initial state s^0 and continues until a terminal state is reached, tracing a trajectory $\tau = ((s^0, \boldsymbol{a}^0), ..., (s^{t-1}, \boldsymbol{a}^{t-1}))$ where $f(s^{t-1}, \boldsymbol{a}^{t-1})$ is terminal. Each agent i plays according to a policy $\pi_i : S \to \Delta A_i$, where ΔA_i is a set of probability distributions over actions in A_i. Given a policy profile $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_{\mathcal{N}})$, the probability of a trajectory is $p(\tau|\boldsymbol{\pi}) = \prod_{(s,\boldsymbol{a})\in\tau} p(\boldsymbol{a}|\, \boldsymbol{a} \sim \boldsymbol{\pi}(s))$, and we define the vector of expected values of the game rooted at state s as $\boldsymbol{V}(\boldsymbol{\pi}, s) = \sum_{\tau \in T(s)} p(\tau|\boldsymbol{\pi}) \sum_{(s,\boldsymbol{a})\in\tau} \boldsymbol{r}(s, \boldsymbol{a})$ where $T(s)$ is the set of terminal trajectories beginning at s. We denote by $\boldsymbol{V}(\boldsymbol{\pi}) := \boldsymbol{V}(\boldsymbol{\pi}, s^0)$ the value of a game rooted in the initial state. We denote by $\boldsymbol{\pi}_{-i}$ the policies of all players other than i.

A **best response (BR)** for agent i against policy profile π is any policy that achieves maximal expected value. That is, π_i^* is a BR to π if $V_i(\pi_i^*, \pi_{-i}) = \max_{\pi_i'} V_i(\pi_i', \pi_{-i})$. A policy profile is a **Nash equilibrium (NE)** if all players' policies are BRs to each other. Letting V_i be the ith component of $\boldsymbol{V}$, $\boldsymbol{\pi}$ is a NE if:

$$\forall i \, \forall \pi_i' : V_i(\boldsymbol{\pi}) \geq V_i(\pi_i', \boldsymbol{\pi}_{-i})$$

An **ϵ-NE** is a policy profile in which an agent can only improve its expected reward by at most ϵ by switching to a different policy. If a policy profile is an ϵ-NE, we say its **exploitability** is ϵ.

We may also talk about Nash equilibria for one stage of the game rather than the whole game. For any state s, supposing we have values $Q(s, \boldsymbol{a})$ that represent the expected values of playing joint action $\boldsymbol{a}$ in state s assuming some policy profile is played thereafter, we may define a new 1-step game where instead the game ends immediately after this action and $Q(s, \boldsymbol{a})$ is itself the final reward. For such stage games, notationally we denote NE using the symbol σ rather than π.

2.3 Nash Q-Learning

Value iteration methods such as Nash-Q have been developed that provably converge to Nash equilibria in some classes of stochastic game, including 2p0s stochastic games [17, 12]. Nash-Q learns joint Q values $\boldsymbol{Q}(s, \boldsymbol{a})$ that aim to converge to the state-action value of $(s, \boldsymbol{a})$ assuming that some NE $\boldsymbol{\pi}$ is played thereafter. This is done by performing 1-step updates on a current estimated function $\boldsymbol{Q}$ as in standard Q-learning, but replacing the max operation with a stage game NE computation.

Formally, suppose that $\boldsymbol{a}$ is played in s leading to s'. A NE σ, $\sigma_i \in \Delta A_i$, is computed for the stage game at s' whose payoffs for each joint action $\boldsymbol{a}'$ are specified by the current values of $\boldsymbol{Q}(s', \boldsymbol{a}')$. The value of this game to each player assuming this NE is then used as the target values of $\boldsymbol{Q}(s, \boldsymbol{a})$. The update rule is thus:

$$\boldsymbol{Q}(s, \boldsymbol{a}) \leftarrow (1 - \alpha)\boldsymbol{Q}(s, \boldsymbol{a}) + \alpha\Big((\boldsymbol{r}(s, \boldsymbol{a}) + \gamma \sum_{\boldsymbol{a}'} \sigma(\boldsymbol{a}')\boldsymbol{Q}(s', \boldsymbol{a}')\Big), \tag{1}$$

where α is the learning rate, γ is the discount factor, and $\sigma(\boldsymbol{a}) := \prod_i \sigma_i(a_i)$ is the probability of joint action $\boldsymbol{a}$.

As presented later in Section 3.1, the foundation of our algorithm is an adaptation of Nash Q-Learning, with several modifications to make it suitable for deep RL and adjusted to use a value function rather than a Q function.

2.4 Double Oracle

Double oracle (DO) is an iterative method for finding a NE over a large set of actions in a matrix game [19]. It leverages the fact that computing a best response is computationally much cheaper than computing a NE. Starting from a small set of candidate actions $A_i^0 \subseteq A_i$ for each player i, on each iteration it computes a NE σ^t in the matrix game restricted to the candidate actions $\boldsymbol{A}^t$. Then, at each iteration t for each player i it finds the best response action $a_i^{t+1} \in A_i$ to the restricted NE σ^t and creates $A_i^{t+1} := A_i^t \bigcup \{a_i^{t+1}\}$. Formally, a_i^{t+1} is an action in the full action set A_i that satisfies

$$a_i^{t+1} = \arg\max_{a' \in A_i} \sum_{\boldsymbol{a} \in \boldsymbol{A}^t} Q_i(s, a_1, \ldots, a_{i-1}, a', a_{i+1}, \ldots, a_{\mathcal{N}}) \prod_{j \neq i} \sigma_j^t(a_j). \tag{2}$$

DO-inspired methods have been used in a variety of other work for handling large or continuous action spaces [13, 6, 16, 1, 18]. Although in the worst case the equilibrium policy may have support over a large fraction of the action space, in practice in many problems a small subset of actions may suffice, and in this case DO-based methods can be computationally efficient due to being able to restrict equilibrium computation to only a tiny fraction of the total number of actions.

As we will present in Section 3.2, we make use of a DO-like procedure both within the RL loop and at inference time, and show it to be an effective way to introduce novel good actions to the policy.

2.5 Regret Matching

Both Nash-Q-Learning and DO require the computation of NEs in matrix games, which we approximate using the iterative **regret matching (RM)** algorithm [4, 11], as has been done in past work [9]. RM is only guaranteed to converge to a NE in 2p0s games and other special classes of games [10]. Nevertheless, past work has found that it produces competitive policies that have low exploitability in important non-2p0s games including six-player poker and no-press Diplomacy [8, 9]. We therefore use RM to approximate a NE σ in restricted matrix games during DO.

In RM, each player i has a **regret** value for each action $a_i \in A_i$. The regret on iteration t is denoted $\text{Regret}_i^t(a_i)$. Initially, all regrets are zero. On each iteration t of RM, $\sigma_i^t(a_i)$ is set according to

$$\sigma_i^t(a_i) = \begin{cases} \frac{\max\{0, \text{Regret}_i^t(a_i)\}}{\sum_{a'_i \in \mathcal{A}_i} \max\{0, \text{Regret}_i^t(a'_i)\}} & \text{if } \sum_{a'_i \in \mathcal{A}_i}, \max\{0, \text{Regret}_i^t(a'_i)\} > 0 \\ \frac{1}{|\mathcal{A}_i|} & \text{otherwise} \end{cases} \tag{3}$$

Next, each player samples an action a_i^* from $\mathcal{A}_i$ according to σ_i^t and all regrets are updated such that

$$\text{Regret}_i^{t+1}(a_i) = \text{Regret}_i^t(a_i) + v_i(a_i, a_{-i}^*) - \sum_{a'_i \in \mathcal{A}_i} \sigma_i^t(a'_i) v_i(a'_i, a_{-i}^*) \tag{4}$$

This sampled form of RM guarantees that the *average* policy over all iterations converges to a ϵ-NE in 2p0s games at a rate of at worst $\mathcal{O}(\frac{1}{\epsilon^2})$ iterations with high probability [15]. In order to improve empirical performance, we use linear RM [7], which weighs updates on iteration t by t.[1] We also use optimism [26], in which the most recent iteration is counted twice when computing regret.

3 Algorithms

We now describe DORA in detail. DORA simultaneously learns a state-value function and an action proposal distribution via with neural networks trained by bootstrapping on an approximate Nash equilibrium for the stage game each turn. A DO-based process is used to discover actions.

3.1 Deep Nash Learning

The core of the learning procedure is based on Nash Q-Learning as presented in Section 2.3, but simplified to use only a value function, with adaptations for large action spaces and function approximation.

Approximating a Q-function in a game with an action space as large as Diplomacy is challenging. However, in Diplomacy as in many other domains, the reward $\boldsymbol{r}(s, a)$ depends only on the next state $s' = f(s, a)$ and not on how it was reached, and since the transition function is known and can be simulated exactly, we can redefine $\boldsymbol{r}$ to be a function of just the next state s'. In that case, we can rewrite Update Rule 1 in terms of a state value function:

$$\boldsymbol{V}(s) \leftarrow (1-\alpha)\boldsymbol{V}(s) + \alpha(\boldsymbol{r}(s) + \gamma \sum_{\boldsymbol{a'}} \sigma(\boldsymbol{a'})\boldsymbol{V}(f(s, \boldsymbol{a'}))) \tag{5}$$

Since the state space is large, we use a deep neural network $\boldsymbol{V}(s; \theta_v)$ with parameters θ_v to approximate $\boldsymbol{V}(s)$. Since the action space is large, exactly computing the Nash equilibrium σ for the 1-step matrix game at a state is also infeasible. Therefore, we train a policy proposal network $\pi(s; \theta_\pi)$ with parameters θ_π that approximates the distribution of actions under the equilibrium policy at state s. We generate a candidate action set at state s by sampling N_b actions from $\pi(s; \theta_\pi)$ for each player and selecting the $N_c \ll N_b$ actions with highest likelihood. We then approximate a NE σ via RM (described in Section 2.5) in the restricted matrix game in which each player is limited to their N_c actions assuming successor states' values are given by our current value network $\boldsymbol{V}(s; \theta_v)$.

We explore states via self-play to generate data, with both players following the computed NE policy with an ϵ probability of exploring a random action from among the N_c proposals. Also, in place of θ_v and θ_π, self-play uses slightly-lagged versions of the networks $\hat{\theta}_v$ and $\hat{\theta}_\pi$ to compute σ that are only updated periodically from the latest θ_v and θ_π in training, a normal practice that improves training stability [20]. We then regress the value network towards the computed stage game value under the NE σ using MSE loss (whose gradient step is simply the above update rule), and regress the policy proposal network towards σ using cross entropy loss:

$$\begin{aligned} \text{ValueLoss}(\theta_v) &= \frac{1}{2}\left(\boldsymbol{V}(s; \theta_v) - \boldsymbol{r}(s) - \gamma \sum_{\boldsymbol{a'}} \sigma(\boldsymbol{a'})\boldsymbol{V}\left(f(s, \boldsymbol{a'}); \hat{\theta}_v\right)\right)^2 \\ \text{PolicyLoss}(\theta_\pi) &= -\sum_i \sum_{a_i \in A_i} \sigma_i(a) \log \pi_i(s, a_i; \theta_\pi), \end{aligned} \tag{6}$$

[1] In practice, rather than weigh iteration t's updates by t we instead discount prior iterations by $\frac{t}{t+1}$ in order to reduce numerical instability. The two options are mathematically equivalent.

We show in Appendix B that in 2p0s games, the exact tabular form of the above algorithm with mild assumptions and without the various approximations for deep RL provably converges to a NE. Outside of 2p0s games and other special classes, both Nash Q-learning and RM lack theoretical guarantees; however, prior work on Diplomacy and 6-player poker have shown that similar approaches still often perform well in practice [9, 2, 8], and we also observe that DORA performs well in 7-player no-press Diplomacy.

At a high level our algorithm is similar to previous algorithms such as AlphaZero [24] and ReBeL [5] that update both a policy and value network based on self-play, leveraging a simulator of the environment to perform search. The biggest differences between our algorithm and AlphaZero are that 1-ply RM acts as the search algorithm instead of Monte Carlo tree search (just as in [9]), the value update is uses a 1-step bootstrap instead of the end-of-game-value (which allows the trajectory generation to be off-policy if desired), and the policy network acts as an action proposer for the search but does not regularize the search thereafter. Our algorithm is also similar to a recent method to handle large action spaces for Q-learning [27] and reduces to it in single-agent settings.

3.2 Double-Oracle-Based Action Discovery

The algorithm in Section 3.1 relies on the policy network to generate good candidate actions. However, good actions initially assigned too low of a probability by the network may never be sampled and reinforced. To address this, we use a mechanism based on double oracle, described in Section 2.4, to introduce novel actions. However, vanilla DO requires computing the expected value of all actions against the restricted matrix game NE. Since our action space is too large to enumerate, we generate only a pool of $N_p \gg N_c$ actions to consider.

A natural idea to generate this pool would be to sample actions uniformly at random from all actions. However, in our results in Section 5.1 we show this approach performs poorly. Intuitively, in a combinatorial action space as large as in Diplomacy often only a minuscule fraction of actions are reasonable, so random sampling is exceedingly unlikely to find useful actions (see Fig 1 for an example).

We instead generate the pool of N_p actions for a player by generating local modifications of the candidate actions already in the restricted matrix game equilibrium. We first sample uniformly among the player's N_d most probable actions in that equilibrium. Next, for each sampled action, we randomly sample a single location on the map and consider all units of the player adjacent to the location. For those units, we randomly choose legal valid new orders. We add the order in this pool that best-responds to the equilibrium as a new candidate action, and DO proceeds as normal.

As shown in Table 3, we found this method of action sampling to outperform simply picking legal and valid combinations of orders uniformly at random. Intuitively, this is likely because locality in Diplomacy, as in a wide variety of other problem domains and algorithms, is a powerful inductive bias and heuristic. Perturbations of good actions are vastly more likely than random actions to also be good, and spatially-nearby actions are more likely to need to covary than spatially-distant ones.

To further improve the scalability of our algorithm, we apply several additional minor approximations, including computing some of the above expected values using sampling and capping the number of DO iterations we perform. For details, see Appendix D.

4 Implementation and Methods

We describe our implementation of the above algorithms and training methods as applied to training in Diplomacy. For details on the exact parameters, see Appendix D.

4.1 Data Generation and Training

We implement deep Nash learning following the same considerations as deep Q-learning. Multiple data generation workers run self-play in parallel. The policy and value targets on each turn of each game along with the game state they correspond to are collected in a shared experience replay buffer and sampled by a set of training workers to train the networks. The policy proposal and value networks used by the data generation workers are only periodically updated from those in training.

During self play, each agent plays according to its part of the computed approximate-NE policy σ with probability $1 - \epsilon$, and with probability ϵ plays a uniformly random action from the N_c actions obtained from the policy proposal network. We also use larger values of ϵ on the first two turns

to quickly introduce diversity from the opening state, similar to the use of larger early exploration settings in other work, for example AlphaZero [25, 24]. Detailed parameters are in Appendix D.

Algorithm 1 Approximated Double Oracle

```
function FindEquilibriumWithDO(s, N_c, N_p, N_iters, ϵ)
    A ← SampleActionsFromProposalNet(s, N_p)
    σ ← ComputeEquilibrium(s, A)
    for t = (1..N_iters) do
        modified ← 0
        for i = (1..N_players) do
            Â_i ← GenerateActions(s, A_i, N_c)          ▷ Uniform sampling or local modification
            v_i ← E_{a~σ} [V_i(f(s, a); θ_v)]
            a*_i, v*_i ← arg max_{a_i ∈ Â_i} E_{a_{-i}~σ_{-i}} [V_i(f(s, a); θ_v)]
            if v*_i − v_i > ϵ then
                A_i ← A_i ⋃ {a*_i}
                σ ← ComputeEquilibrium(s, A)
                modified ← 1
        if !modified then
            Break
    return σ

function GenerateActionsLocalModification(s, A_base, N_c)
    A ← {}
    while |A| < N_c do
        a ← SampleAction(A_base)
        loc ← SampleMapLocation()
        clique ← GetNeighbours(loc)
        a' ← RandomCoordinatedModification(a, clique)
        A ← A ⋃ {a'}
    return A
```

4.2 Model Architecture

Our model closely follows the encoder-decoder architecture of [9]. We use the policy head to model the policy proposal distribution and the value head to predict the expected values of the players.

We found several adjustments to be effective for our use in deep Nash-based RL. Firstly, we removed all dropout layers, as we found this to improve training stability. Secondly, we replaced the graph convolution-based encoder with a simpler transformer encoder, improving model quality and removing the need to hardwire the adjacency matrix of the Diplomacy map into the architecture. Lastly, we also split the network into separate policy and value networks. Since our regret minimization procedure is essentially only a 1-ply search, the states on which we need to evaluate policy (the current state) and value (the states after one turn) are disjoint, so there is no gain to computing both policy and value in one network query, and we found splitting the networks improved the quality and effective capacity of the models. We measure the impact of these choices in Table 2 and describe the architecture in more detail in Appendix A.

4.3 Exploitability Testing

To obtain a metric of the exploitablity of our final agents, we also train additional exploiter agents that rather than attempting to play an equilibrium, instead hold the exploited agent fixed and attempt to converge to a best response. We compare two approaches.

In the first approach, we reimplement the same actor-critic method with sampled entropy regularization previously used in [9] for 7-player no-press Diplomacy and apply it to 2-player FvA. Unlike [9], we optimized versus our final agents directly and did not train proxy agents in their place.

In the second approach, we begin with the policy and value models of the exploited agent and resume deep Nash learning, except with a single simple modification, described in Appendix E, that results in the learning of a best response rather than an equilibrium.

5 Results

Diplomacy features two main challenges for AI research: its large combinatorial action space and its multi-agent (non-2p0s) nature. *DORA* is intended to address the former challenge. To measure its effectiveness on just the action-exploration aspect, we first present results in a 2p0s variant of Diplomacy called France vs. Austria (FvA). We show *DORA* decisively beats top humans in FvA.

Next, we present results for 7-player no-press Diplomacy, which has been a domain for much prior research [21, 2, 9]. We train a *DORA* agent from scratch with no human data and test it against pre-trained models from previous work [21, 9][2]. In the 1v6 setting (one *DORA* agent playing with six copies of another bot), *DORA* performs roughly on par with some past human-data-based agents. In the 6v1 setting, *DORA* wins overwhelmingly. Our results suggest that *DORA* converges to strong play within an equilibrium (i.e., a "metagame") that is very different from how humans play. Finally, we present results showing that adding our RL algorithm on top of a human-bootstrapped model clearly beats prior bots in both the 1v6 and 6v1 settings.

5.1 Results in 2-player France vs. Austria

	Austria	**France**	**Total**
Average score (+variance reduction)	$89\% \pm 5\%$	$66\% \pm 7\%$	$78\% \pm 4\%$
Raw average score	$100\% \pm 0\%$	$73\% \pm 12\%$	$87\% \pm 6\%$
# games	13	13	26

Table 1: Head-to-head score of *DORA* versus top FvA players on `webdiplomacy.net`. The $\pm$ shows one standard error. Variance reduction (Appendix C) was used to compute average scores from raw average scores.

In 2-player FvA, we tested our final *DORA* agent against multiple top players of this variant, inviting them to play a series of games against our agent. All players were allowed to play an arbitrary number of practice games against the agent before playing official counted games. Prior to running the games, we also implemented a simple variance reduction method, described in Appendix C.

In Table 1 we report both the raw and variance-reduced average scores. DORA defeats top human players by a large and highly significant margin, with an average score of 78% after variance reduction. It also significantly more than half of games as France, the side considered much harder to play[3]. These results show DORA achieves a level of play in FvA past that of current top humans.

We ran several ablations to loosely estimate the relative importance of the various components contributing to our final agent for France vs Austria.

Architecture	Final Elo	Params
GraphConv [9]	898 ± 13	46.8M
TransformerEnc 5x192	1041 ± 7	9.3M
TransformerEnc 10x224	1061 ± 12	12.4M
TransformerEnc 10x224 (Separate)	1123 ± 13	24.8M

Table 2: Effect in FvA of the architecture on final strength after RL, and number of parameters. GraphConv is the architecture achieving SOTA in supervised learning from [9], and similar to other past work [21, 2]. TransformerEnc b x f replaces the graph-conv encoder with a transformer encoder with b blocks and f feature channels. "Separate" uses separate value and policy networks instead of a shared network with two heads. Elos are based on 100 games against each of several internal baseline agents, including the agents in Table 4.

GraphConv vs Transformers See Table 2. Although not a primary focus of this paper, we found that at least for use in RL, our model architecture change to use a transformer encoder rather than a graph-convolution encoder in Section 4.2 resulted in significant gains over prior neural network architectures in Diplomacy [21, 2, 9], while using fewer parameters and still being similarly fast or faster in performance. Enlarging the model and splitting the network led to further gains.

[2] See Appendix D for the details.

[3] Between agents of similar strength we often observe France winning only about a third of games vs Austria.

DO Method During RL Training	Final Elo
No DO	927 ± 13
DO, uniform random sampling	956 ± 13
DO, local modification sampling	1023 ± 13

Table 3: Effect of DO methods during training on final FvA agent strength. Local modification sampling improves results greatly, whereas uniform random sampling is relatively ineffective. Elos $\pm$ one standard error were computed based on 100 games against each of several baseline agents, including the agents in Table 4.

Model used by both players → DO Method used by measured player ↓	Model1	Model2	Model3
None	0.5±0	0.5±0	0.5±0
DO, uniform random sampling	0.75±0.02	0.48±0.02	0.47±0.02
DO, local modification sampling	0.89±0.02	0.59±0.02	0.56±0.02

Table 4: Effect of using DO at inference time in FvA. Each entry is the average score when both players play the column agent, but one player uses a DO variant at inference time. Model1-Model3 are models of increasing quality developed over the course of our research, with the major differences being that Model1 uses a trained value network with a fixed proposal network, Model2 uses GraphConv encoder, and Model3 uses a small Transformer encoder. Stronger agents require better sampling methods to find candidate actions that help.

Double Oracle at Training Time See Table 3. We demonstrate the effectiveness of our DO method in finding improved actions during training. DO's ability to discover new strong actions for the policy proposal network to learn makes a big difference in the final agent. Local modification sampling is effective, whereas random sampling is relatively ineffective.

Double Oracle at Inference Time See Table 4. We also demonstrate the effectiveness of our DO method in finding improved candidate actions at inference time. Uniform random sampling among legal actions to generate the action pool for finding best responses only improves the weakest of our baseline agent at inference time, due to being ineffective at finding better actions within the huge action space, whereas local modification sampling improves all baselines at inference time.

5.1.1 Exploitability

	vs DORA	**vs DORA w/ inference-time DO**
Actor-Critic PG [9]	$77.0\% \pm 0.6\%$	-
Deep Nash-Derived Exploiter	$90.5\% \pm 0.4\%$	$85.8\% \pm 0.5\%$
Deep Nash-Derived Exploiter, values only	$65.7\% \pm 0.6\%$	$60.5\% \pm 0.7\%$

Table 5: Win rates for exploiter bots trained with different techniques against our best FvA agent with and without inference-time DO. The $\pm$ shows one standard error. Actor-Critic PG vs DORA with inference-time DO was not run due to computational cost.

We implemented and tested both algorithms from Section 4.3 to estimate a lower bound on the exploitability of our agents. We also test a variant where we adapt deep Nash-learning for exploitation just as in the second method, but disable policy learning - the value network learns to predict states of higher exploitability, but the policy proposal network is fixed and does not learn.

The results against our best agent are in Table 5. We find that despite being of superhuman strength, the agent is still relatively exploitable. In a 2p0s large-action-space setting, this is not necessarily surprising. For example in Go, another game with a moderately large action space, there exist gaps of more than 1000 Elo (nearly 100% winrate) between different agents, all of which are superhuman, even with no explicit training for exploitation [23]. Moreover, in imperfect-information games such as Diplomacy, exploiting a fixed agent is much easier than training an agent with low exploitability, since the exploiting agent does not need to precisely balance the probabilities of its actions.

Nonetheless, we also observe that our adaptation of the deep Nash learning method greatly outperforms the simpler actor-critic method from [9], and that both value learning and DO-based policy learning are important for its performance, with value learning improving over the baseline from 50% to 60.5%, and policy learning improving further to 85.8% winrate over *DORA*.

We also observe that applying DO at inference time reduces the exploitability, from 90.5% using our strongest exploiter agent to 85.8%. Even at inference time, DO continues to find and suggest new effective actions that would otherwise be overlooked by the policy proposal network.

5.2 Results in 7-player No-Press Diplomacy

We also applied our methods to 7-player no-press Diplomacy, training a *DORA* agent entirely from scratch with no human data in this more-challenging domain. For comparison, we also trained an agent via deep Nash learning beginning from policy and value networks first trained on human data, which we label as *HumanDNL*[4]. The results are reported in Table 6. Note that as these are 7-player games, the expected score if the agents perform equally well is $1/7 \sim 14.3\%$.

1x Agent vs	6x DipNet [21]	6x SearchBot [9]	6x *DORA*	6x *HumanDNL*
DipNet [21]	-	$0.2\% \pm 0.1\%$	$0.1\% \pm 0.1\%$	$0\% \pm 0\%$
SearchBot [9]	$48.8\% \pm 2.6\%$	-	$0.1\% \pm 0.1\%$	$3\% \pm 0.6\%$
DORA	$14.6\% \pm 1.9\%$	$1.2\% \pm 0.4\%$	-	$0.5\% \pm 0.2\%$
HumanDNL	$44.6\% \pm 2.6\%$	$28\% \pm 1.5\%$	$8.4\% \pm 1.4\%$	-

Table 6: SoS scores of various agents playing against 6 copies of another agent. The $\pm$ shows one standard error. Note that equal performance is 1/7 $\approx$ 14.3%. The non-human DORA performs very well as a 6x opponent, while being on-par with DipNet when playing on the 1x side with 6 human-like opponents.

Unlike a 2p0s setting, in a competitive multi-agent game, particularly one such as Diplomacy where agents must temporarily ally and coordinate to make progress, one expects that there may be a large number of different equilibria that are mutually incompatible. An agent playing optimally under one equilibrium may perform very poorly in a pool of agents playing a different equilibrium.

Our results are consistent with *DORA* converging roughly to an equilibrium that is very different than the one played by humans, but playing well within that equilibrium. *DORA* achieves only about a 1% winrate in a 1-vs-6 setting against one of the strongest prior human-data-based agents, SearchBot [9], but SearchBot fares even worse with only a 0.1% winrate in a 1-vs-6 setting against *DORA*. In the 1-vs-6 setting *DORA* still manages to perform on par with DipNet, a human-based model that performs about as well as a moderately-experienced human player.

Our results with *HumanDNL* underscore this point as well. *HumanDNL* soundly defeats SearchBot, the prior SOTA, winning 28% of games 1-vs-6, while SearchBot wins only 3% in the reverse matchup. This result provides additional evidence that deep Nash learning is effective, but limited by the compatibility of different equilibria. Indeed, SearchBot's *higher* 3% winrate versus *HumanDNL* compared to its 0.1% winrate versus DORA likely reflects how *HumanDNL* may be playing an approximate equilibrium that by being more human-like is *easier* for SearchBot to survive in.

6 Conclusions and Future Work

Diplomacy has been a benchmark for AI research for two main reasons: its massive combinatorial action space and its multi-agent nature. We show that *DORA* achieves superhuman performance in 2-player Diplomacy without any human data or reward shaping, despite the combinatorial action space and imperfect information. *DORA* therefore effectively addresses the first challenge of Diplomacy.

Additionally, we show that *DORA* in full-scale 7-player no-press Diplomacy approximates an equilibrium that differs radically from human-data-based agents. When one human-data-based agent plays against a population of *DORA* agents, the one agent loses by a wide margin. Similarly, one *DORA* agent against a population of human-data-based agents does not exceed an average human level of performance. These 7-player results provide evidence that self-play from scratch may be insufficient to achieve superhuman performance in Diplomacy, unlike other multiplayer games like 6-player poker [8]. Furthermore, the results suggest that there may be a wide space of equilibria, with some differing dramatically from human conventions. *DORA* provides the first competitive technique for exploring this wider space of equilibria without being constrained to human priors.

[4]Our results with human-initialization are preliminary, DO was not yet tested, only deep Nash learning.

References

[1] Lukáš Adam, Rostislav Horčík, Tomáš Kasl, and Tomáš Kroupa. Double oracle algorithm for computing equilibria in continuous games. 09 2020.

[2] Thomas Anthony, Tom Eccles, Andrea Tacchetti, János Kramár, Ian Gemp, Thomas C Hudson, Nicolas Porcel, Marc Lanctot, Julien Pérolat, Richard Everett, et al. Learning to play no-press diplomacy with best response policy iteration. *arXiv preprint arXiv:2006.04635*, 2020.

[3] Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dębiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.

[4] David Blackwell et al. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1):1–8, 1956.

[5] Noam Brown, Anton Bakhtin, Adam Lerer, and Qucheng Gong. Combining deep reinforcement learning and search for imperfect-information games. *arXiv preprint arXiv:2007.13544*, 2020.

[6] Noam Brown and Tuomas Sandholm. Simultaneous abstraction and equilibrium finding in games. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.

[7] Noam Brown and Tuomas Sandholm. Solving imperfect-information games via discounted regret minimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1829–1836, 2019.

[8] Noam Brown and Tuomas Sandholm. Superhuman AI for multiplayer poker. *Science*, page eaay2400, 2019.

[9] Jonathan Gray, Adam Lerer, Anton Bakhtin, and Noam Brown. Human-level performance in no-press diplomacy via equilibrium search. *arXiv preprint arXiv:2010.02923*, 2020.

[10] James Hannan. Approximation to bayes risk in repeated play. *Contributions to the Theory of Games*, 3:97–139, 1957.

[11] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.

[12] Junling Hu and Michael P Wellman. Nash q-learning for general-sum stochastic games. *Journal of machine learning research*, 4(Nov):1039–1069, 2003.

[13] Manish Jain, Dmytro Korzhyk, Ondřej Vaněk, Vincent Conitzer, Michal Pěchouček, and Milind Tambe. A double oracle algorithm for zero-sum security games on graphs. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 327–334, 2011.

[14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[15] Marc Lanctot, Kevin Waugh, Martin Zinkevich, and Michael Bowling. Monte carlo sampling for regret minimization in extensive games. In *Advances in neural information processing systems*, pages 1078–1086, 2009.

[16] Marc Lanctot, Vinicius Zambaldi, Audrunas Gruslys, Angeliki Lazaridou, Karl Tuyls, Julien Pérolat, David Silver, and Thore Graepel. A unified game-theoretic approach to multiagent reinforcement learning. In *nternational Conference on Neural Information Processing Systems*, pages 4193—4206, 2017.

[17] Michael L Littman. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*, pages 157–163. Elsevier, 1994.

[18] Stephen Mcaleer, JB Lanier, Roy Fox, and Pierre Baldi. Pipeline psro: A scalable approach for finding approximate nash equilibria in large games. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 20238–20248. Curran Associates, Inc., 2020.

[19] Brendan McMahan, Geoffrey Gordon, and Avrim Blum. Planning in the presence of cost functions controlled by an adversary. In *International conference on machine learning*, pages 536—-543, 2003.

[20] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

[21] Philip Paquette, Yuchen Lu, Seton Steven Bocco, Max Smith, O-G Satya, Jonathan K Kummerfeld, Joelle Pineau, Satinder Singh, and Aaron C Courville. No-press diplomacy: Modeling multi-agent gameplay. In *Advances in Neural Information Processing Systems*, pages 4474–4485, 2019.

[22] Lloyd S Shapley. Stochastic games. *Proceedings of the national academy of sciences*, 39(10):1095–1100, 1953.

[23] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484, 2016.

[24] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, et al. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018.

[25] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354, 2017.

[26] Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E Schapire. Fast convergence of regularized learning in games. In *Advances in Neural Information Processing Systems*, pages 2989–2997, 2015.

[27] Tom Van de Wiele, David Warde-Farley, Andriy Mnih, and Volodymyr Mnih. Q-learning in enormous action spaces via amortized approximate maximization. *arXiv preprint arXiv:2001.08116*, 2020.

[28] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.

7 Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?
   [Yes]
   (b) Did you describe the limitations of your work?
   [Yes] Particularly with regards to non-2p0s settings, we discuss limitations interlaced through the discussion of theoretical grounding and through discussion of our results in 7-player and the conclusion.
   (c) Did you discuss any potential negative societal impacts of your work?
   [N/A] We do not expect near-term societal impacts from this work.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them?
   [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results?
   [Yes]
   (b) Did you include complete proofs of all theoretical results?
   [Yes] In Appendix B that also refers to prior work that proves the main result we rely on.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)?
   [No] The experiments in the paper require an extensive code base that is challenging to check for any violation of the Double Blind principles. We plan to publish the repository with camera-ready version.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)?
   [Yes] Detailed parameters are provided in Appendix D.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)?
   [Yes] Error bars are given on relevant measurements and statistics.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)?
   [Yes] We report the requirements in Appendix D.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators?
   [Yes]
   (b) Did you mention the license of the assets?
   [Yes] We specify the license of the model and code we use in appendix D. We plan to release the code and the models under MIT license.
   (c) Did you include any new assets either in the supplemental material or as a URL?
   [No] As explained in 3a we plan to release the model and the code after the Double Blind review ends.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating?
   [Yes] We discuss data used in appendix F.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content?
   [Yes] We do not use any data with personally identifiable information or offensive content (see appendix F).
5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable?
[Yes] We included the instructions given to participants in appendix F.

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable?
[N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation?
[N/A]