

Human-Level Performance in No-Press Diplomacy via Equilibrium Search

Jonathan Gray*, Adam Lerer*, Anton Bakhtin, Noam Brown

Facebook AI Research

{jsgray, alerer, yolo, noambrown}@fb.com

Abstract

Prior AI breakthroughs in complex games have focused on either the purely adversarial or purely cooperative settings. In contrast, Diplomacy is a game of shifting alliances that involves both cooperation and competition. For this reason, Diplomacy has proven to be a formidable research challenge. In this paper we describe an agent for the no-press variant of Diplomacy that combines supervised learning on human data with one-step lookahead search via external regret minimization. External regret minimization techniques have been behind previous AI successes in adversarial games, most notably poker, but have not previously been shown to be successful in large-scale games involving cooperation. We show that our agent greatly exceeds the performance of past no-press Diplomacy bots, is unexploitable by expert humans, and achieves a rank of 23 out of 1,128 human players when playing anonymous games on a popular Diplomacy website.

Introduction

A primary goal for AI research is to develop agents that can act optimally in real-world multi-agent interactions (i.e., **games**). In recent years, AI agents have achieved expert-level or even superhuman performance in benchmark games such as backgammon (Tesauro 1994), chess (Campbell, Hoane Jr, and Hsu 2002), Go (Silver et al. 2016, 2017, 2018), poker (Moravčík et al. 2017; Brown and Sandholm 2017, 2019b), and real-time strategy games (Berner et al. 2019; Vinyals et al. 2019). However, previous large-scale game AI results have focused on either purely competitive or purely cooperative settings. In contrast, real-world games, such as business negotiations, politics, and traffic navigation, involve a far more complex mixture of cooperation and competition. In such settings, the theoretical grounding for the techniques used in previous AI breakthroughs falls apart.

In this paper we augment neural policies trained through imitation learning with regret minimization search techniques, and evaluate on the benchmark game of no-press Diplomacy. Diplomacy is a longstanding benchmark for research that features a rich mixture of cooperation and competition. Like previous researchers, we evaluate on the

widely played no-press variant of Diplomacy, in which communication can only occur through the actions in the game (i.e., no cheap talk is allowed).

Specifically, we begin with a **blueprint** policy that approximates human play in a dataset of Diplomacy games. We then improve upon the blueprint during play by approximating an equilibrium for the current phase of the game, assuming all players (including our agent) play the blueprint for the remainder of the game. Our agent then plays its part of the computed equilibrium. The equilibrium is computed via regret matching (RM) (Blackwell et al. 1956; Hart and Mas-Colell 2000).

Search via RM has led to remarkable success in poker. However, RM only converges to a Nash equilibrium in two-player zero-sum games and other special cases, and RM was never previously shown to produce strong policies in a mixed cooperative/competitive game as complex as no-press Diplomacy. Nevertheless, we show that our agent exceeds the performance of prior agents and for the first time convincingly achieves human-level performance in no-press Diplomacy. Specifically, we show that our agent soundly defeats previous agents, that our agent is far less exploitable than previous agents, that an expert human cannot exploit our agent even in repeated play, and, most importantly, that our agent achieves a score of 25.6% when playing anonymously with humans on a popular Diplomacy website, compared to an average human score of 14.3%.

Background and Related Work

Search has previously been used in almost every major game AI breakthrough, including backgammon (Tesauro 1994), chess (Campbell, Hoane Jr, and Hsu 2002), Go (Silver et al. 2016, 2017, 2018), poker (Moravčík et al. 2017; Brown and Sandholm 2017, 2019b), and Hanabi (Lerer et al. 2020). A major exception is real-time strategy games (Vinyals et al. 2019; Berner et al. 2019). Similar to SPARTA as used in Hanabi (Lerer et al. 2020), our agent conducts one-ply lookahead search (i.e., changes the policy just for the current game turn) and thereafter assumes all players play according to the blueprint. Similar to the Pluribus poker agent (Brown and Sandholm 2019b), our search technique uses regret matching to compute an approximate equilibrium. In a manner similar to the sampled best response algorithm of (Anthony et al. 2020), we sample a limited number of actions

*Equal contribution

from the blueprint policy rather than search over all possible actions, which would be intractable.

Learning effective policies in games involving cooperation and competition has been studied extensively in the field of multi-agent reinforcement learning (MARL) (Shoham, Powers, and Grenager 2003). Nash-Q and CE-Q applied Q learning for general sum games by using Q values derived by computing Nash (or correlated) equilibrium values at the target states (Hu and Wellman 2003; Greenwald, Hall, and Serrano 2003). Friend-or-foe Q learning treats other agents as either cooperative or adversarial, where the Nash Q values are well defined (Littman 2001). The recent focus on "Deep" MARL has led to learning rules from game theory such as fictitious play and regret minimization being adapted to Deep reinforcement learning (Heinrich and Silver 2016; Brown et al. 2019), as well as work on game-theoretic challenges of mixed cooperative/competitive settings such as social dilemmas and multiple equilibria in the MARL setting (Leibo et al. 2017; Lerer and Peysakhovich 2017, 2019).

Diplomacy in particular has served for decades as a benchmark for multi-agent AI research (Kraus and Lehmann 1988; Kraus, Ephrati, and Lehmann 1994; Kraus and Lehmann 1995; Johansson and Hård 2005; Ferreira, Cardoso, and Reis 2015). Recently, Paquette et al. (2019) applied imitation learning (IL) via deep neural networks on a dataset of more than 150,000 Diplomacy games. This work greatly improved the state of the art for no-press Diplomacy, which was previously a handcrafted agent (van Hal 2013). Paquette et al. (2019) also tested reinforcement learning (RL) in no-press Diplomacy via Advantage Actor-Critic (A2C) (Mnih et al. 2016). (Anthony et al. 2020) introduced sampled best response policy iteration, a self-play technique, which further improved upon the performance of (Paquette et al. 2019).

Description of Diplomacy

The rules of no-press Diplomacy are complex; a full description is provided by Paquette et al. (2019). No-press Diplomacy is a seven-player zero-sum board game in which a map of Europe is divided into 75 provinces. 34 of these provinces contain supply centers (SCs), and the goal of the game is for a player to control a majority (18) of the SCs. Each player begins the game controlling three or four SCs and an equal number of units.

The game consists of three types of phases: movement phases in which each player assigns an order to each unit they control, retreat phases in which defeated units retreat to a neighboring province, and adjustment phases in which new units are built or existing units are destroyed.

During a movement phase, a player assigns an order to each unit they control. A unit's order may be to hold (defend its province), move to a neighboring province, convoy a unit over water, or support a neighboring unit's hold or move order. Support may be provided to units of any player. We refer to a tuple of orders, one order for each of a player's units, as an **action**. That is, each player chooses one action each turn. There are an average of 26 valid orders for each unit (Paquette et al. 2019), so the game's branching factor

is massive and on some turns enumerating all actions is intractable.

Importantly, all actions occur simultaneously. In live games, players write down their orders and then reveal them at the same time. This makes the game an imperfect-information game in which an optimal policy may need to be stochastic in order to prevent predictability.

Diplomacy is designed in such a way that cooperation with other players is almost essential in order to achieve victory, even though only one player can ultimately win.

A game may end in a draw on any turn if all remaining players agree. Draws are a common outcome among experienced players because players will often coordinate to prevent any individual from reaching 18 centers. The two most common scoring systems for draws are **draw-size scoring (DSS)**, in which all surviving players equally split a win, and **sum-of-squares scoring (SoS)**, in which player i receives a score of $\frac{C_i^2}{\sum_{j \in \mathcal{N}} C_j^2}$, where C_i is the number of SCs that player i controls (Fogel 2020). Throughout this paper we use SoS scoring except in anonymous games against humans where the human host chooses a scoring system.

Regret Matching

Regret Matching (RM) (Blackwell et al. 1956; Hart and Mas-Colell 2000) is an iterative algorithm that converges to a Nash equilibrium (NE) (Nash 1951) in two-player zero-sum games and other special cases, and converges to a coarse correlated equilibrium (CCE) (Hannan 1957) in general.

We consider a game with \mathcal{N} players where each player i chooses an action a_i from a set of actions \mathcal{A}_i . We denote the joint action as $a = (a_1, a_2, \dots, a_N)$, the actions of all players other than i as a_{-i} , and the set of joint actions as \mathcal{A} . After all players simultaneously choose an action, player i receives a reward of $v_i(a)$ (which can also be represented as $v_i(a_i, a_{-i})$). Players may also choose a probability distribution over actions, where the probability of action a_i is denoted $\pi_i(a_i)$ and the vector of probabilities is denoted π_i .

Normally, each iteration of RM has a computational complexity of $\prod_{i \in \mathcal{N}} |\mathcal{A}_i|$. In a seven-player game, this is typically intractable. We therefore use a sampled form of RM in which each iteration has a computational complexity of $\sum_{i \in \mathcal{N}} |\mathcal{A}_i|$. We now describe this sampled form of RM.

Each agent i maintains an **external regret** value for each action $a_i \in \mathcal{A}_i$, which we refer to simply as **regret**. The regret on iteration t is denoted $R_i^t(a_i)$. Initially, all regrets are zero. On each iteration t of RM, $\pi_i^t(a_i)$ is set according to

$$\pi_i^t(a_i) = \begin{cases} \frac{R_i^t(a_i)_+}{\sum_{a'_i \in \mathcal{A}_i} R_i^t(a'_i)_+} & \text{if } \sum_{a'_i} R_i^t(a'_i)_+ > 0 \\ \frac{1}{|\mathcal{A}_i|} & \text{otherwise} \end{cases} \quad (1)$$

where $R_i^t(a_i)_+ = \max\{0, R_i^t(a_i)\}$. Next, each player samples an action a_i^* from \mathcal{A}_i according to π_i^t and all regrets are updated such that

$$R_i^{t+1}(a_i) = R_i^t(a_i) + v_i(a_i, a_{-i}^*) - \sum_{a'_i \in \mathcal{A}_i} \pi_i^t(a'_i) v_i(a'_i, a_{-i}^*) \quad (2)$$

This sampled form of RM guarantees that $R_i^t(a_i) \in \mathcal{O}(\sqrt{t})$ with high probability (Lanctot et al. 2009). If $R_i^t(a_i)$ grows sublinearly for all players’ actions, as in RM, then the *average* policy over all iterations converges to a NE in two-player zero-sum games and in general the empirical distribution of players’ joint policies converges to a CCE as $t \rightarrow \infty$.

In order to improve empirical performance, we use linear RM (Brown and Sandholm 2019a), which weighs updates on iteration t by t .¹ We also use optimism (Syrgekani et al. 2015), in which the most recent iteration is counted twice when computing regret. Additionally, the action our agent ultimately plays is sampled from the *final* iteration’s policy, rather than the average policy over all iterations. This reduces the risk of sampling a non-equilibrium action due to insufficient convergence. In theory sampling from the final iteration may increase exploitability, but this technique has been used successfully in past poker agents (Brown and Sandholm 2019b).

Agent Description

Our agent is composed of two major components. The first is a **blueprint** policy and state-value function trained via imitation learning on human data. The second is a search algorithm that utilizes the blueprint. This algorithm is executed on every turn, and approximates an equilibrium policy (for all players, not just the agent) for the current turn via RM, assuming that the blueprint is played by all players for the remaining game beyond the current turn.

Supervised Learning

We construct a blueprint policy via imitation learning on a corpus of 46,148 Diplomacy games collected from online play, building on the methodology and model architecture described by Paquette et al. (2019) and Anthony et al. (2020). A blueprint policy and value function estimated from human play is ideal for performing search in a general-sum game, because it is likely to realistically approximate state values and other players’ actions when playing with humans. Our blueprint supervised model is based on the DipNet agent from (Paquette et al. 2019), but we make a number of modifications to the architecture and training.

We trained the blueprint policy using only a subset of the data used by Paquette et al. (2019), specifically those games obtained from webdiplomacy.net. For this subset of the data, we obtained metadata about the press variant (full-press vs. no-press) which we add as a feature to the model, and anonymized player IDs for the participants in each game. Using the IDs, we computed ratings s_i for each player i and only trained the policy on actions from players with above-average ratings.

To compute player ratings, we used a regularized logistic outcome model. Specifically, we optimized the loss

$$L(s|\mathcal{D}) = \mathbb{E}_{(i,j) \in \mathcal{D}} [\sigma(s_i - s_j)] + \lambda |s|_2$$

¹In 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.

across all pairs of players $(i, j) \in \mathcal{D}$ where player i achieved a "better" outcome than player j in a game. We found this approach led to more plausible scores than Elo (Elo 1978) or TrueSkill (Herbrich, Minka, and Graepel 2007) ratings.

Paquette et al. (2019) took an orthogonal approach to filter poor players from the training data: they only trained on "winning" powers, i.e. those who ended the game with at least 7 SCs. This filtering is sensible for training a policy for play, but is problematic for training policies for search. In a general-sum game, it is crucial for the agent to be able to predict the empirical distribution of actions even for other agents who are destined to lose.

Our model closely follows the architecture of Paquette et al. (2019), with additional dropout of 0.4 between GNN encoder layers. We model sets of build orders as single tokens because there are a small number of build order combinations and it is tricky to predict *sets* auto-regressively with teacher forcing. We adopt the encoder changes of Anthony et al. (2020), but do not adopt their relational order decoder because it is more expensive to compute and leads to only marginal accuracy improvements after tuning dropout.

We make a small modification to the encoder GNN architecture that improves modeling. In addition to the standard residual that skips the entire GNN layer, we replace the graph convolution² with the sum of a graph convolution and a linear layer. This allows the model to learn a hierarchy of features for each graph node (through the linear layer) without requiring a concomitant increase in graph smoothing (the GraphConv (GrCv)). The resulting GNN layer computes (modification in bold)

$$x_{i+1} = Dropout(ReLU(BN(GrCv(x_i) + \mathbf{A}x_i))) + x_i \quad (3)$$

where A is a learned linear transformation.

Finally, we achieve a substantial improvement in order prediction accuracy using a *featurized order decoder*. Diplomacy has over 13,000 possible orders, many of which will be observed infrequently in the training data. Therefore, by featurizing the orders by the order type, and encodings of the source, destination, and support locations, we observe improved prediction accuracy.

Specifically, in a standard decoder each order o has a learned representation e_o , and for some board encoding x and learned order embedding e_o , $P(o) = softmax(x \cdot e_o)$. With order featurization, we use $\tilde{e}_o = e_o + A f_o$, where f_o are static order features and A is a learned linear transformation. The order featurization we use is the concatenation of the one-hot order type with the board encodings for the source, destination, and support locations. We found that representing order location features by their location encodings works better than one-hot locations, presumably because the model can learn more state-contextual features.³

²As noted by Anthony et al. (2020), DipNet uses a variant of a GNN that learns a separate weight matrix at each graph location.

³We note the likelihood that a transformer with token-based decoding should capture a similar featurization, although Paquette et al. (2019) report worse performance for both a transformer and token-based decoding.

We add an additional value head to the model immediately after the dipnet encoder, that is trained to estimate the final SoS scores given a board situation. We use this value estimate during equilibrium search (Sec.) to estimate the value of a Monte Carlo rollout after a fixed number of steps. The value head is an MLP with one hidden layer that takes as input the concatenated vector of all board position encodings. A softmax over powers’ SoS scores is applied at the end to enforce that all players’ SoS scores sum to 1.

Equilibrium Search

The policy that is actually played results from a search algorithm which utilizes the blueprint policy. Let s be the current state of the game. On each turn, the search algorithm computes an equilibrium for a **subgame** and our agent plays according to its part of the equilibrium solution for its next action.

Conceptually, the subgame is a well-defined game that begins at state s . The set of actions available to each player is a subset of the possible actions in state s in the full game, and are referred to as the **subgame actions**. Each player i chooses a subgame action a_i , resulting in joint subgame action a . After a is taken, the players make no further decisions in the subgame. Instead, the players receive a reward corresponding to the players sampling actions according to the blueprint policy π^b for the remaining game.

The subgame actions for player i are the M_i highest-probability actions according to the blueprint model. M_i is a hyperparameter that is proportional to the number of units controlled by player i . The effect of different choices for M_i is plotted in Figure 1 (left).

Rolling out π^b to the end of the game is very expensive, so in practice we instead roll out π^b for a small number of turns (usually 2 or 3 movement phases in our experiments) until state s' is reached, and then use the value for s' from the blueprint’s value network as the reward vector. Figure 1 (right) shows the performance of our search agent using different rollout lengths. We do not observe improved performance for rolling out farther than 3 or 4 movement phases.

We compute a policy for each agent by running the sampled regret matching algorithm described in Equation (1) and (2). The search algorithm used 256 – 4,096 iterations of RM and typically required between 2 minutes and 20 minutes per turn using a single Volta GPU and 8 CPU cores, depending on the hyperparameters used for the game..

Results

Using the techniques described in the previous section, we developed an agent we call SearchBot. Our experiments focus on two formats. The first evaluates SearchBot playing head-to-head against other bots and against the population of human players on a popular Diplomacy website. The second measures the exploitability of SearchBot.

⁴This is slightly different than reported in (Paquette et al. 2019) because we compute token-accuracy treating a full set of build orders as a single token.

⁵Policy accuracy for these model are not comparable to above because training data was modified.

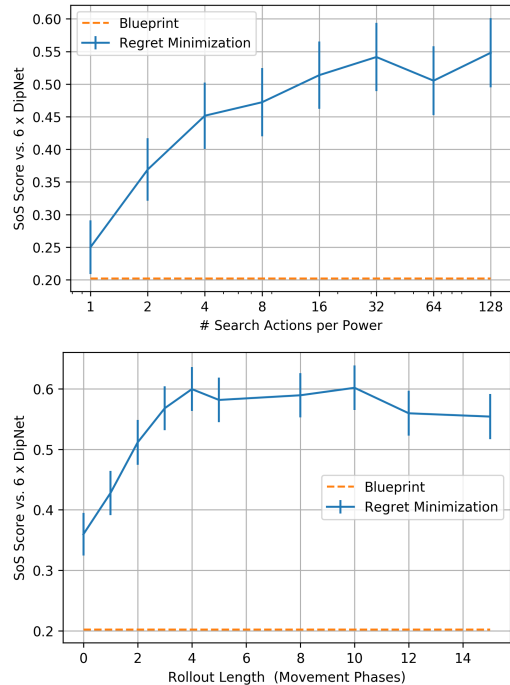


Figure 1: **Left:** Score of SearchBot using different numbers of sampled subgame actions M_i , against 6 DipNet agents ((Paquette et al. 2019) at temperature 0.1). A score of 14.3% would be a tie. Even when sampling only two actions, SearchBot dramatically outperforms our blueprint, which achieves a score of 20.2%. **Right:** The effect of different rollout lengths on SearchBot performance.

Table 2 compares our search agent (SearchBot) with the supervised and RL DipNet agents. Following prior work, we compute average scores in ‘1v6’ games containing a single agent of type A and six agents of type B. The average SoS score of an identical agent should therefore be $1/7 \approx 14.3\%$. The SearchBot agent achieves its highest 1v6 SoS score when matched against its own blueprint, since it is most accurately able to approximate the behavior of that agent. It outperforms all three agents by a large margin, and none of the three baselines is able to achieve more than 1% SoS score against our search agent.

Performance against a population of human players

The ultimate test of an AI system is how well it performs in the real world with humans. To measure this, we had SearchBot anonymously play no-press Diplomacy games on the popular Diplomacy website webdiplomacy.net. Since there are 7 players in each game, average human performance is a score of 14.3%. In contrast, SearchBot scored $25.6\% \pm 4.8\%$.⁶ If the bot’s performance for each of the 7 powers is weighed equally, this score increases to $27.0\% \pm 5.3\%$. The

⁶Most games used draw-size scoring in the case of draws, while our agent was trained based on sum-of-squares scoring. If all games are scored using sum-of-squares, our agent achieves a score of $30.2\% \pm 5.3\%$.

Model	Policy Accuracy	SoS v. DipNet	
		temp=0.5	temp=0.1
DipNet ((Paquette et al. 2019))	60.5% ⁴	0.143	
+ combined build orders & encoder dropout	62.0%		
+ encoder changes from (Anthony et al. 2020)	62.4%	0.150	0.198
<hr/>			
<i>switch to webdiplomacy training data only</i>	61.3% ⁵	0.175	0.206
+ output featurization	62.0%	0.184	0.188
+ improved GNN layer	62.4%	0.183	0.205
+ merged GNN trunk	62.9%	0.199	0.202

Table 1: Effect of model and training data changes on supervised model quality. We measure policy accuracy as well as average SoS score achieved by each agent against 6 of the original DipNet model. We measure the SoS scores in two settings: with all 7 agents sampling orders at a temperature of either 0.5 or 0.1.

Agent A (1x)	Agent B (6x)	Average Score	Games
SearchBot	SL DipNet	54.5% \pm 2%	670
SearchBot	RL DipNet	35.6% \pm 2%	699
SearchBot	Blueprint	60.8% \pm 2%	699
SL DipNet	SearchBot	0.2% \pm 0.2%	138
RL DipNet	SearchBot	0.7% \pm 0.3%	140
Blueprint	SearchBot	0.6% \pm 0.3%	140

Table 2: Comparison of average sum-of-squares scores for our agent (SearchBot) in 1v6 games with DipNet agents from Paquette et al. (2019), as well as our own blueprint imitation learning agent. All agents other than SearchBot use a temperature of 0.1.

agent’s performance is shown in Table 3.

In addition to raw score, we measured SearchBot’s performance using the Ghost-Rating system (Anthony 2020), which is a Diplomacy rating system inspired by the Elo system that accounts for the relative strength of opponents and that is used to semi-officially rank players on webdiplomacy.net. Among no-press Diplomacy players on the site, our agent ranked 23 out of 1,128 players with a Ghost-Rating of 176.0 as of September 30th, 2020.⁷

Most games on webdiplomacy.net were played with 24-hour turns, though the agent also played in some “live” games with 5-minute turns. Different hyperparameters were used for live games versus non-live games. In non-live games, we typically ran RM for 2,048 iterations with a rollout length of 3 movement phases, and set M_i equal to 5 times the number of units a player controls. This typically required about 20 minutes to compute. In live games, including games in which one human played against six bots, we typically ran RM for 256 iterations with a rollout length of 2 movement phases, and set M_i equal to 3.5 times the number of units a player controls. This typically required about 2 minutes to compute. In all cases, the temperature for the blueprint in rollouts was set to 0.75.

The experiments on webdiplomacy.net occurred over a three-month timespan, with games commonly taking one to

⁷Our agent played games under different accounts; we report the Ghost-Rating for these accounts merged.

two months to complete (players are typically given 24 hours to act). Freezing research and development over such a period would have been impractical, so our agent was not fixed for the entire time period. Instead, serious bugs were fixed, improvements to the algorithm were made, and the model was updated.

Exploitability

While performance of an agent within a population of human players is the most important metric, that metric alone does not capture how the population of players might adapt to the agent’s presence. For example, if our agent is extremely strong then over time other players might adopt the bot’s playstyle. As the percentage of players playing like the bot increases, other players might adopt a policy that seeks to exploit this playstyle. Thus, if the bot’s policy is highly exploitable then it might eventually do poorly even if it initially performs well against the population of human players.

This can partly be interpreted through an evolutionary lens using the notion of an evolutionarily stable strategy (ESS) (Taylor and Jonker 1978; Smith 1982), which is a refinement of Nash equilibrium (Nash 1951). If our agent’s policy is an ESS, then a population of players all playing the agent’s policy could not be “invaded” by a different policy. That is, no other policy could do better than tie against the population’s policy.

Motivated by this, we measure the **exploitability** of our agent. Exploitability of a policy profile π (denoted $e(\pi)$) measures worst-case performance when all but one agents follows π . Formally, the exploitability of π is defined as $e(\pi) = \sum_{i \in \mathcal{N}} \max_{\pi_i} v_i(\pi_i, \pi_{-i}) / N$, where π_{-i} denotes the policies of all players other than i . Agent i ’s **best response** to π_{-i} is defined as $BR(\pi_{-i}) = \arg \max_{\pi_i} v_i(\pi_i, \pi_{-i})$.

We estimate our agent’s full-game exploitability in two ways: by training an RL agent to best respond to the bot, and by having expert humans repeatedly play against six copies of the bot. We also measure the ‘local’ exploitability in the search subgame and show that it converges to an approximate Nash equilibrium.

Performance against a best-responding agent When the policies of all players but one are fixed, the game becomes

Power	Bot Score	Human Mean	Games	Wins	Draws	Losses
All Games	25.6% \pm 4.8%	14.3%	50	7	16	27
Normalized By Power	27.0% \pm 5.3%	14.3%	50	7	16	27

Table 3: Performance of our agent in anonymous games against humans on webdiplomacy.net. Average human performance is 14.3%. Score in the case of draws was determined by the rules of the joined game. The \pm shows one standard error.

a Markov Decision Process (MDP) (Howard 1960) for the non-fixed player because the actions of the fixed players can be viewed as stochastic transitions in the “environment”. Thus, we can estimate the exploitability of π by first training a best response policy $BR(\pi_{-i})$ for each agent i using any single-agent RL algorithm, and then computing $\sum_{i \in \mathcal{N}} v_i(BR(\pi_{-i}), \pi_{-i})/N$. Since the best response RL policy will not be an exact best response (which is intractable to compute in a game as complex as no-press Diplomacy) this only gives us a lower-bound estimate of the exploitability.

Following other work on environments with huge action spaces (Vinyals et al. 2019; Berner et al. 2019), we use a distributed asynchronous actor-critic RL approach to optimize the exploiter policy (Espeholt et al. 2018). We use the same architecture for the exploiter agent as for the fixed model. Moreover, to simplify the training we initialize the exploiter agent from the fixed model.

We found that training becomes unstable when the policy entropy gets too low. The standard remedy is to use an entropy regularization term. However, due to the immense action space, an exact computation of the entropy term, $E_a \log p_\theta(a)$, is infeasible. Instead, we optimize a surrogate loss that gives an unbiased estimate of the gradient of the entropy loss (see the extended version of this paper for details). We found this to be critical for the stability of the training.

Training an RL agent to exploit SearchBot is prohibitively expensive. Even when choosing hyperparameters that would result in the agent playing as fast as possible, SearchBot typically requires at least a full minute in order to act each turn. Instead, we collect a dataset of self-play games of SearchBot and train a supervised agent on this dataset. The resulting agent, which we refer to as SearchBot-clone, is weaker than SearchBot but requires only a single pass through the neural network in order to act on a turn. By training an agent to exploit SearchBot-clone, we can obtain a (likely) upper bound on what the performance would be if a similar RL agent were trained against SearchBot. We report the reward of the exploiter agents against the blueprint and SearchBot-clone agents in Figure 2.

Performance Against Expert Human Exploiters In addition to training a best-responding agent, we also invited the 1st and 2nd place finishers in the 2017 World Diplomacy Convention (widely considered the world championship for full-press Diplomacy) to play games against six copies of our agent. The purpose was to determine whether the human experts could discover exploitable weaknesses in the bot.

The humans played games against three types of bots: DipNet (Paquette et al. 2019) (with temperature set to 0.5),




Figure 2: Score of the exploiting agent against the blueprint and SearchBot-clone as a function of training time. We report the average of six runs. The shaded area corresponds to three standard errors. We use temperature 0.5 for both agents as it minimizes exploitability for the blueprint. Since SearchBot-clone is trained through imitation learning of SearchBot, the exploitability of SearchBot is almost certainly lower than SearchBot-clone.

our blueprint agent (with temperature set to 0.5), and SearchBot. In total, the participants played 35 games against each bot; each of the seven powers was controlled by a human player five times, while the other six powers were controlled by identical copies of the bot. The performance of the humans is shown in Table 4. While the sample size is relatively small, the results suggest that our agent is less exploitable than prior bots.

Exploitability in Local Subgame We first investigate the exploitability of our agent in the local subgame defined by a given board state, sampled actions, and assumed blueprint policy for the rest of the game. We simulate 7 games between a search agent and 6 DipNet agents, and plot the total exploitability of the average strategy of the search procedure as a function of the number of RM iterations, as well as the exploitability of the blueprint policies. Utilities u_i are computed using Monte Carlo rollouts with the same (blueprint) rollout policy used during RM, and total exploitability for a joint policy π is computed as $e(\pi) = \sum_i \max_{a_i \in \mathcal{A}_i} u_i(a_i, \pi_{-i}) - u_i(\pi)$. The exploitability curves aggregated over all phases are shown in Figure 3 (left).

In Figure 3 (right), we verify that the average of policies from multiple independent executions of RM also converges to an approximate Nash. For example, it is possible that if each agent independently running RM converged to a different incompatible equilibrium and played their part of it, then the joint policy of all the agents would not be an ESS. However we observe that the exploitability of the average of

Power	1 Human vs. 6 DipNet	1 Human vs. 6 Blueprint	1 Human vs. 6 SearchBot
All Games	39.1%	22.5%	5.7%

Table 4: Performance of one expert human playing against six bots under repeated play. A score less than 14.3% means the human is unable to exploit the bot. Five games were played for each power for each agent, for a total of 35 games per agent. For each power, the human first played all games against DipNet, then the blueprint model, and then finally SearchBot.

policies closely matches the exploitability of the individual policies.

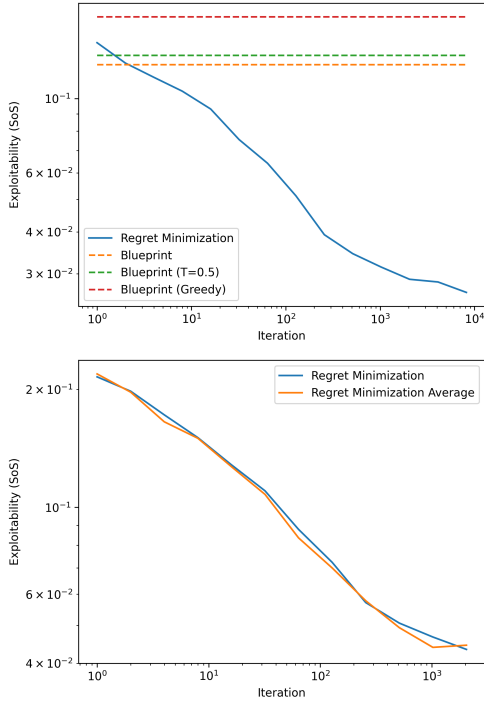


Figure 3: **Left:** Distance of the RM average strategy from equilibrium as a function of the RM iteration, computed as the sum of all agents’ exploitability in the matrix game in which RM is employed. RM reduces exploitability, while the blueprint policy has only slightly lower exploitability than the uniform distribution over the 50 sampled actions used in RM (i.e. RM iteration 1). **Right:** Comparison of convergence of individual strategies to the average of two independently computed strategies. The similarity of these curves suggests that independent RM computations lead to compatible equilibria. Note: In both figures, exploitability is averaged over all phases in 7 simulated games.

Conclusions

No-press Diplomacy is a complex game involving both cooperation and competition that poses major theoretical and practical challenges for past AI techniques. Nevertheless, our AI agent achieves human-level performance in this game with a combination of supervised learning on human data and one-ply search using external regret minimization. The massive improvement in performance from conducting search just one action deep matches a larger trend seen in

other games, such as chess, Go, poker, and Hanabi, in which search dramatically improves performance. While external regret minimization has been behind previous AI breakthroughs in purely competitive games, it was never previously shown to be successful in a complex game involving cooperation. The success of RM in no-press Diplomacy suggests that its use is not limited to purely adversarial games.

Our work points to several avenues for future research. SearchBot conducts search only for the current turn. In principle, this search could extend deeper into the game tree using counterfactual regret minimization (CFR) (Zinkevich et al. 2008). However, the size of the subgame grows exponentially with the depth of the subgame. Developing search techniques that scale more effectively with the depth of the game tree may lead to substantial improvements in performance. Another direction is combining our search technique with reinforcement learning. Combining search with reinforcement learning has led to tremendous success in perfect-information games (Silver et al. 2018) and more recently in two-player zero-sum imperfect-information games as well (Brown et al. 2020). Finally, it remains to be seen whether similar search techniques can be developed for variants of Diplomacy that allow for coordination between agents.

References

- Anthony, T. 2020. Ghost-Ratings. URL <https://sites.google.com/view/webdipinfo/ghost-ratings>.
- Anthony, T.; Eccles, T.; Tacchetti, A.; Kramár, J.; Gemp, I.; Hudson, T. C.; Porcel, N.; Lanctot, M.; Pérolat, J.; Ev-erett, R.; et al. 2020. Learning to Play No-Press Diplomacy with Best Response Policy Iteration. *arXiv preprint arXiv:2006.04635*.
- Berner, C.; Brockman, G.; Chan, B.; Cheung, V.; Debiak, P.; Dennison, C.; Farhi, D.; Fischer, Q.; Hashme, S.; Hesse, C.; et al. 2019. Dota 2 with Large Scale Deep Reinforcement Learning. *arXiv preprint arXiv:1912.06680*.
- Blackwell, D.; et al. 1956. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics* 6(1): 1–8.
- Brown, N.; Bakhtin, A.; Lerer, A.; and Gong, Q. 2020. Combining Deep Reinforcement Learning and Search for Imperfect-Information Games. *arXiv preprint arXiv:2007.13544*.
- Brown, N.; Lerer, A.; Gross, S.; and Sandholm, T. 2019. Deep Counterfactual Regret Minimization. In *International Conference on Machine Learning*, 793–802.

- Brown, N.; and Sandholm, T. 2017. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science* eaa01733.
- Brown, N.; and Sandholm, T. 2019a. Solving imperfect-information games via discounted regret minimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 1829–1836.
- Brown, N.; and Sandholm, T. 2019b. Superhuman AI for multiplayer poker. *Science* eaay2400.
- Campbell, M.; Hoane Jr, A. J.; and Hsu, F.-h. 2002. Deep Blue. *Artificial intelligence* 134(1-2): 57–83.
- Elo, A. E. 1978. *The rating of chessplayers, past and present*. Arco Pub.
- Espeholt, L.; Soyer, H.; Munos, R.; Simonyan, K.; Mnih, V.; Ward, T.; Doron, Y.; Firoiu, V.; Harley, T.; Dunning, I.; et al. 2018. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. *arXiv preprint arXiv:1802.01561*.
- Ferreira, A.; Cardoso, H. L.; and Reis, L. P. 2015. Dipblue: A diplomacy agent with strategic and trust reasoning. In *ICAART International Conference on Agents and Artificial Intelligence, Proceedings*.
- Fogel, B. 2020. To Whom Tribute is Due: The Next Step in Scoring Systems. URL <http://windycityweasels.org/wp-content/uploads/2020/04/2020-03-To-Whom-Tribute-Is-Due-The-Next-Step-in-Scoring-Systems.pdf>.
- Greenwald, A.; Hall, K.; and Serrano, R. 2003. Correlated Q-learning. In *ICML*, volume 20, 242.
- Hannan, J. 1957. Approximation to Bayes risk in repeated play. *Contributions to the Theory of Games* 3: 97–139.
- Hart, S.; and Mas-Colell, A. 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68(5): 1127–1150.
- Heinrich, J.; and Silver, D. 2016. Deep reinforcement learning from self-play in imperfect-information games. *arXiv preprint arXiv:1603.01121*.
- Herbrich, R.; Minka, T.; and Graepel, T. 2007. TrueSkill™: a Bayesian skill rating system. In *Advances in neural information processing systems*, 569–576.
- Howard, R. A. 1960. Dynamic programming and markov processes. .
- Hu, J.; and Wellman, M. P. 2003. Nash Q-learning for general-sum stochastic games. *Journal of machine learning research* 4(Nov): 1039–1069.
- Johansson, S. J.; and Håård, F. 2005. Tactical coordination in no-press diplomacy. In *International Joint Conference on Autonomous Agents and Multiagent Systems*, 423–430.
- Kraus, S.; Ephrati, E.; and Lehmann, D. 1994. Negotiation in a non-cooperative environment. *Journal of Experimental & Theoretical Artificial Intelligence* 3(4): 255–281.
- Kraus, S.; and Lehmann, D. 1988. Diplomat, an agent in a multi agent environment: An overview. In *IEEE International Performance Computing and Communications Conference*, 434–435. IEEE Computer Society.
- Kraus, S.; and Lehmann, D. 1995. Designing and building a negotiating automated agent. *Computational Intelligence* 11(1): 132–171.
- Lanctot, M.; Waugh, K.; Zinkevich, M.; and Bowling, M. 2009. Monte Carlo sampling for regret minimization in extensive games. In *Advances in neural information processing systems*, 1078–1086.
- Leibo, J. Z.; Zambaldi, V.; Lanctot, M.; Marecki, J.; and Graepel, T. 2017. Multi-agent reinforcement learning in sequential social dilemmas. *arXiv preprint arXiv:1702.03037*.
- Lerer, A.; Hu, H.; Foerster, J.; and Brown, N. 2020. Improving Policies via Search in Cooperative Partially Observable Games. In *AAAI Conference on Artificial Intelligence*.
- Lerer, A.; and Peysakhovich, A. 2017. Maintaining cooperation in complex social dilemmas using deep reinforcement learning. *arXiv preprint arXiv:1707.01068*.
- Lerer, A.; and Peysakhovich, A. 2019. Learning Existing Social Conventions via Observationally Augmented Self-Play. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 107–114. ACM.
- Littman, M. L. 2001. Friend-or-foe Q-learning in general-sum games. In *ICML*, volume 1, 322–328.
- Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, 1928–1937.
- Moravčík, M.; Schmid, M.; Burch, N.; Lisý, V.; Morrill, D.; Bard, N.; Davis, T.; Waugh, K.; Johanson, M.; and Bowling, M. 2017. Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science* 356(6337): 508–513.
- Nash, J. 1951. Non-cooperative games. *Annals of mathematics* 286–295.
- Paquette, P.; Lu, Y.; Bocco, S. S.; Smith, M.; Satya, O.-G.; Kummerfeld, J. K.; Pineau, J.; Singh, S.; and Courville, A. C. 2019. No-press diplomacy: Modeling multi-agent gameplay. In *Advances in Neural Information Processing Systems*, 4474–4485.
- Shoham, Y.; Powers, R.; and Grenager, T. 2003. Multi-agent reinforcement learning: a critical survey. *Web manuscript 2*.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature* 529(7587): 484.
- Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; et al. 2018. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* 362(6419): 1140–1144.
- Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; et al. 2017. Mastering the game of go without human knowledge. *Nature* 550(7676): 354.

Smith, J. M. 1982. *Evolution and the Theory of Games*. Cambridge university press.

Syrskanis, V.; Agarwal, A.; Luo, H.; and Schapire, R. E. 2015. Fast convergence of regularized learning in games. In *Advances in Neural Information Processing Systems*, 2989–2997.

Taylor, P. D.; and Jonker, L. B. 1978. Evolutionary stable strategies and game dynamics. *Mathematical biosciences* 40(1-2): 145–156.

Tesauro, G. 1994. TD-Gammon, a self-teaching backgammon program, achieves master-level play. *Neural computation* 6(2): 215–219.

van Hal, J. 2013. Diplomacy AI - Albert. URL <https://sites.google.com/site/diplomacyai/>.

Vinyals, O.; Babuschkin, I.; Czarnecki, W. M.; Mathieu, M.; Dudzik, A.; Chung, J.; Choi, D. H.; Powell, R.; Ewalds, T.; Georgiev, P.; et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* 575(7782): 350–354.

Zinkevich, M.; Johanson, M.; Bowling, M.; and Piccione, C. 2008. Regret minimization in games with incomplete information. In *Advances in neural information processing systems*, 1729–1736.