
Deep Reinforcement Learning from Self-Play in Imperfect-Information Games

Johannes Heinrich

David Silver

University College London, UK

J.HEINRICH@CS.UCL.AC.UK

D.SILVER@CS.UCL.AC.UK

Abstract

Many real-world applications can be described as large-scale games of imperfect information. To deal with these challenging domains, prior work has focused on computing Nash equilibria in a handcrafted abstraction of the domain. In this paper we introduce the first scalable end-to-end approach to learning approximate Nash equilibria without any prior knowledge. Our method combines fictitious self-play with deep reinforcement learning. When applied to Leduc poker, Neural Fictitious Self-Play (NFSP) approached a Nash equilibrium, whereas common reinforcement learning methods diverged. In Limit Texas Hold'em, a poker game of real-world scale, NFSP learnt a competitive strategy that approached the performance of human experts and state-of-the-art methods.

1. Introduction

Games have a tradition of encouraging advances in artificial intelligence and machine learning (Samuel, 1959; Tesauro, 1995; Campbell et al., 2002; Riedmiller et al., 2009; Gelly et al., 2012; Bowling et al., 2015). Game theory defines a game as a domain of conflict or cooperation between several entities (Myerson, 1991). One motivation of studying the simpler recreational games is to develop algorithms that will scale to more complex, real-world games such as airport and network security, financial and energy trading, traffic control and routing (Lambert III et al., 2005; Nevmyvaka et al., 2006; Bazzan, 2009; Tambe, 2011; Urieli & Stone, 2014; Durkota et al., 2015). Most of these real-world games involve decision making with imperfect information and high-dimensional information state spaces. Unfortunately, many machine learning methods, that have been applied to classical games, lack convergence guarantees for learning in imperfect-information games. On the other hand, many game-theoretic approaches lack the ability to extract relevant patterns and generalise from data. This results in limited scalability to large games, unless the domain is abstracted to a manageable size using

human expert knowledge, heuristics or modelling. However, acquiring human expertise often requires expensive resources and time. In addition, humans can be easily fooled into irrational decisions or assumptions (Selten, 1990; Ariely & Jones, 2008). This motivates algorithms that learn useful strategies end-to-end.

In this paper we introduce NFSP, a deep reinforcement learning method for learning approximate Nash equilibria of imperfect-information games. NFSP agents learn by playing against themselves without explicit prior knowledge. Technically, NFSP extends and instantiates Fictitious Self-Play (FSP) (Heinrich et al., 2015) with neural network function approximation. An NFSP agent consists of two neural networks and two kinds of memory. Memorized experience of play against fellow agents is used by reinforcement learning to train a network that predicts the expected values of actions. Experience of the agent's own behaviour is stored in a separate memory, which is used by supervised learning to train a network that predicts the agent's own average behaviour. An NFSP agent acts cautiously by sampling its actions from a mixture of its average, routine strategy and its greedy strategy that maximizes its predicted expected value. NFSP approximates fictitious play, which is a popular game-theoretic model of learning in games that converges to Nash equilibria in some classes of games, e.g. two-player zero-sum and many-player potential games.

We empirically evaluate our method in two-player zero-sum computer poker games. In this domain, current game-theoretic approaches use heuristics of card strength to abstract the game to a tractable size (Zinkevich et al., 2007; Gilpin et al., 2007; Johanson et al., 2013). While Limit Texas Hold'em (LHE), a poker game of real-world scale, has got within reach of being solved with current computational resources (Bowling et al., 2015), most other poker and real-world games remain far out of scope without abstraction. Our approach does not rely on engineering such abstractions or any other prior knowledge. NFSP agents leverage deep reinforcement learning to learn directly from their experience of interacting in the game. When applied to Leduc poker, NFSP approached a Nash equilibrium, whereas common reinforcement learning methods di-

verged. We also applied NFSP to LHE, learning directly from the raw inputs. NFSP learnt a competitive strategy, approaching the performance of state-of-the-art methods based on handcrafted abstractions.

2. Background

In this section we provide a brief overview of reinforcement learning, extensive-form games and fictitious self-play. For a more detailed exposition we refer the reader to (Sutton & Barto, 1998), (Myerson, 1991), (Fudenberg, 1998) and (Heinrich et al., 2015).

2.1. Reinforcement Learning

Reinforcement learning (Sutton & Barto, 1998) agents typically learn to maximize their expected future rewards from interaction with an environment. The environment is usually modelled as a **Markov decision process (MDP)**. An agent behaves according to a **policy** that specifies a distribution over available actions at each state of the MDP. The agent’s goal is to improve its policy in order to maximize its **gain**, $G_t = \sum_{i=t}^T R_{i+1}$, which is a random variable of the agent’s cumulative future rewards starting from time t .

Many reinforcement learning algorithms learn from sequential **experience** in the form of transition tuples, $(s_t, a_t, r_{t+1}, s_{t+1})$, where s_t is the state at time t , a_t is the action chosen in that state, r_{t+1} the reward received thereafter and s_{t+1} the next state that the agent transitioned to. A common objective is to learn the **action-value function**, $Q(s, a) = \mathbb{E}^\pi [G_t | S_t = s, A_t = a]$, defined as the expected gain of taking action a in state s and following policy π thereafter. An agent is learning **on-policy** if it learns about the policy that it is currently following. In the **off-policy** setting an agent learns from experience of another agent or another policy, e.g. a previous policy.

Q-learning (Watkins & Dayan, 1992) is a popular off-policy reinforcement learning method. It learns about the greedy policy, which at each state takes the action of the highest estimated value. Storing and replaying past experience by applying off-policy reinforcement learning to the respective transition tuples is known as experience replay (Lin, 1992). Fitted Q Iteration (FQI) (Ernst et al., 2005) is a batch reinforcement learning method that replays experience with Q-learning. Neural Fitted Q Iteration (NFQ) (Riedmiller, 2005) and Deep Q Network (DQN) (Mnih et al., 2015) are extensions of FQI that use neural network function approximation with batch and online updates respectively.

2.2. Extensive-Form Games

Extensive-form games are a model of sequential interaction involving multiple players. Assuming rationality, each

player’s goal is to maximize his payoff in the game. In imperfect-information games, each player only observes his respective **information states**, e.g. in a poker game a player only knows his own private cards but not those of other players. Each player chooses a **behavioural strategy** that maps information states to probability distributions over available actions. We assume games with **perfect recall**, i.e. each player’s current information state s_t^i implies knowledge of the sequence of his information states and actions, $s_1^i, a_1^i, s_2^i, a_2^i, \dots, s_t^i$, that led to this information state. The **realization-probability** (Von Stengel, 1996), $x_{\pi^i}(s_t^i) = \prod_{k=1}^{t-1} \pi^i(s_k^i, a_k^i)$, determines the probability that player i ’s behavioural strategy, π^i , contributes to realizing his information state s_t^i . A **strategy profile** $\pi = (\pi^1, \dots, \pi^n)$ is a collection of strategies for all players. π^{-i} refers to all strategies in π except π^i . Given a fixed strategy profile π^{-i} , any strategy of player i that achieves optimal payoff performance against π^{-i} is a **best response**. An approximate or ϵ -best response is suboptimal by no more than ϵ . A **Nash equilibrium** is a strategy profile such that each player’s strategy in this profile is a best response to the other strategies. Similarly, an approximate or ϵ -Nash equilibrium is a profile of ϵ -best responses. In a Nash equilibrium no player can gain by deviating from his strategy. Therefore, a Nash equilibrium can be regarded as a fixed point of rational self-play learning. In fact, Nash equilibria are the only strategy profiles that rational agents can hope to converge on in self-play (Bowling & Veloso, 2001).

2.3. Fictitious Self-Play

Fictitious play (Brown, 1951) is a game-theoretic model of learning from self-play. Fictitious players choose best responses to their opponents’ average behaviour. The average strategies of fictitious players converge to Nash equilibria in certain classes of games, e.g. two-player zero-sum and many-player potential games (Robinson, 1951; Monderer & Shapley, 1996). Leslie & Collins (2006) introduced generalised weakened fictitious play. It has similar convergence guarantees as common fictitious play, but allows for approximate best responses and perturbed average strategy updates, making it particularly suitable for machine learning.

Fictitious play is commonly defined in normal form, which is exponentially less efficient for extensive-form games. Heinrich et al. (2015) introduce **Full-Width Extensive-Form Fictitious Play (XFP)** that enables fictitious players to update their strategies in behavioural, extensive form, resulting in linear time and space complexity. A key insight is that for a convex combination of normal-form strategies, $\hat{\sigma} = \lambda_1 \hat{\pi}_1 + \lambda_2 \hat{\pi}_2$, we can achieve a realization-equivalent behavioural strategy σ , by setting it to be proportional to the respective convex combination of realization-

probabilities,

$$\sigma(s, a) \propto \lambda_1 x_{\pi_1}(s) \pi_1(s, a) + \lambda_2 x_{\pi_2}(s) \pi_2(s, a) \forall s, a, \quad (1)$$

where $\lambda_1 x_{\pi_1}(s) + \lambda_2 x_{\pi_2}(s)$ is the normalizing constant for the strategy at information state s . In addition to defining a full-width average strategy update of fictitious players in behavioural strategies, equation (1) prescribes a way to sample data sets of such convex combinations of strategies. Heinrich et al. (2015) introduce **Fictitious Self-Play (FSP)**, a sample- and machine learning-based class of algorithms that approximate XFP. FSP replaces the best response computation and the average strategy updates with reinforcement and supervised learning respectively. In particular, FSP agents generate datasets of their experience in self-play. Each agent stores its experienced transition tuples, $(s_t, a_t, r_{t+1}, s_{t+1})$, in a memory, \mathcal{M}_{RL} , designated for reinforcement learning. Experience of the agent’s own behaviour, (s_t, a_t) , is stored in a separate memory, \mathcal{M}_{SL} , designated for supervised learning. Self-play sampling is set up in a way that an agent’s reinforcement learning memory approximates data of an MDP defined by the other players’ average strategy profile. Thus, an approximate solution of the MDP by reinforcement learning yields an approximate best response. Similarly, an agent’s supervised learning memory approximates data of the agent’s own average strategy, which can be learned by supervised classification.

3. Neural Fictitious Self-Play

NFSP is an evolution of FSP, introducing multiple extensions such as neural network function approximation, reservoir sampling, anticipatory dynamics and a fully agent-based approach. An NFSP agent interacts with the other players in a game and memorizes its experience of game transitions and its own behaviour. NFSP treats these memories as two datasets suitable for deep reinforcement learning and supervised classification. In particular, the agent trains a neural network, F_Q , to predict action values, $Q(s, a)$, from the data in \mathcal{M}_{RL} using off-policy reinforcement learning. The resulting network defines the agent’s approximate best response strategy, $\beta = \epsilon\text{-greedy}(F_Q)$, which selects a random action with probability ϵ and otherwise chooses the action that maximizes the predicted action values. The NFSP agent trains a separate neural network, F_S , to imitate its own past behaviour using supervised classification on the data in \mathcal{M}_{SL} . This network maps states to action probabilities and defines the agent’s average strategy, $\pi = F_S$. During play, the agent chooses its actions from a mixture of its two strategies, β and π .

While fictitious players usually best respond to the average strategy of their opponents, in continuous-time dynamic fictitious play (Shamma & Arslan, 2005) players choose

Algorithm 1 Neural Fictitious Self-Play (NFSP) with DQN

Require:

Γ {Game}
 $\mathcal{M}_{RL}, \mathcal{M}_{SL}$ {RL and SL memories}
 F_Q, F_S {Action value and policy networks}
 $\beta = \epsilon\text{-GREEDY}(F_Q)$ {Best response policy}
 $\pi = F_S$ {Average policy}
 σ {Current policy}

Ensure: π an approximate Nash equilibrium in self-play

function STEP()

$s_t, r_t, c_t \leftarrow \text{OBSERVE}(\Gamma)$

$a_t \leftarrow \text{THINK}(s_t, r_t, c_t)$

ACT(Γ, a_t)

end function

function THINK(s_t, r_t, c_t)

if $c_t = 0$ {episode terminated} **then**

$\sigma \leftarrow \text{SAMPLEPOLICY}(\beta, \pi)$

end if

if $s_{t-1} \neq \text{nil}$ **then**

$\tau_t \leftarrow (s_{t-1}, a_{t-1}, r_t, s_t, c_t)$

UPDATERLMEMORY(\mathcal{M}_{RL}, τ_t)

end if

$a_t \leftarrow \text{SAMPLEACTION}(\sigma)$

if $\sigma = \beta$ **then**

UPDATESLMEMORY($\mathcal{M}_{SL}, (s_t, a_t)$)

end if

$s_{t-1} \leftarrow s_t$

$a_{t-1} \leftarrow a_t$

$\beta \leftarrow \text{REINFORCEMENTLEARNING}(\mathcal{M}_{RL})$

$\pi \leftarrow \text{SUPERVISEDLEARNING}(\mathcal{M}_{SL})$

end function

function REINFORCEMENTLEARNING(\mathcal{M}_{RL})

$F_Q \leftarrow \text{DQN}(\mathcal{M}_{RL})$

return $\epsilon\text{-GREEDY}(F_Q)$

end function

function SUPERVISEDLEARNING(\mathcal{M}_{SL})

$F_S \leftarrow$ Apply stochastic gradient descent to loss

$\mathbb{E}_{(s,a) \sim \mathcal{M}_{SL}} [-\log \pi(s, a)]$

return F_S

end function

best responses to a short-term prediction of their opponents’ average normal-form strategies, $\hat{\pi}_t^{-i} + \eta \frac{d}{dt} \hat{\pi}_t^{-i}$. The authors show that for appropriate, game-dependent choice of η stability of fictitious play at equilibrium points can be improved. NFSP uses $\hat{\beta}_{t+1}^i - \hat{\pi}_t^i \approx \frac{d}{dt} \hat{\pi}_t^i$ as a discrete-time approximation of the derivative that is used in these anticipatory dynamics. Note that $\Delta \hat{\pi}_t^i \propto \hat{\beta}_{t+1}^i - \hat{\pi}_t^i$ is the normal-form update direction of common discrete-time fictitious play. In order for an NFSP agent to compute an approximate best response, β^i , to its opponents’ anticipated average strategy profile, $\sigma^{-i} \equiv \hat{\pi}_t^{-i} + \eta(\hat{\beta}^{-i} - \hat{\pi}_t^{-i})$, the

agent iteratively evaluates and maximizes its action values, $Q^i(s, a) \approx \mathbb{E}_{\beta^i, \sigma^{-i}} [G_t^i | S_t = s, A_t = a]$. This can be achieved by off-policy reinforcement learning, e.g. Q-learning or DQN, on experience of play against the opponents' anticipated strategy, σ^{-i} . To ensure that the agents' reinforcement learning memories, \mathcal{M}_{RL} , contain this kind of experience, NFSP requires all agents to choose their actions from $\sigma \equiv (1 - \eta)\hat{\pi} + \eta\hat{\beta}$, where $\eta \in \mathbb{R}$ is called the **anticipatory parameter**.

Fictitious play usually keeps track of the average of normal-form best response strategies that players have chosen in the game, $\hat{\pi}_T^i = \frac{1}{T} \sum_{t=1}^T \hat{\beta}_t^i$. Heinrich et al. (2015) propose to use sampling and machine learning to generate data on and learn convex combinations of normal-form strategies in extensive form. E.g. we can generate a set of extensive-form data of π_T^i by sampling whole episodes of the game, using $\beta_t^i, t = 1, \dots, T$, in proportion to their weight, $\frac{1}{T}$, in the convex combination. NFSP uses reservoir sampling (Vitter, 1985; Osborne et al., 2014) to memorize experience of its average best responses. The agent's supervised learning memory, \mathcal{M}_{SL} , is a reservoir to which it only adds experience when following its approximate best response policy β . An NFSP agent regularly trains its average policy network, $\pi = F_S$, to match its average behaviour stored in its supervised learning memory, e.g. by optimizing the log-probability of past actions taken. Algorithm 1 presents NFSP, using DQN for reinforcement learning.

4. Experiments

We evaluate NFSP and related algorithms in Leduc (Southey et al., 2005) and Limit Texas Hold'em poker games. Most of our experiments measure the exploitability of learned strategy profiles. In a two-player zero-sum game, the exploitability of a strategy profile is defined as the expected average payoff that a best response profile achieves against it. An exploitability of 2δ yields at least a δ -Nash equilibrium.

4.1. Robustness of XFP

To understand how function approximation interacts with FSP, we begin with some simple experiments that emulate approximation and sampling errors in the full-width algorithm XFP. Firstly, we explore what happens when the perfect averaging used in XFP is replaced by an incremental averaging process closer to gradient descent. Secondly, we explore what happens when the exact table lookup used in XFP is replaced by an approximation with epsilon error.

Figure 1 shows the performance of XFP with default, $1/T$, and constant stepsizes for its strategy updates. We see improved asymptotic but lower initial performance for smaller

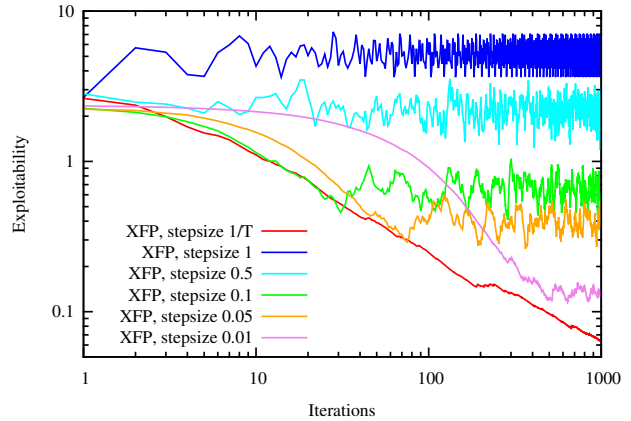


Figure 1. The impact of constant stepsizes on the performance of full-width fictitious play in Leduc Hold'em.

stepsizes. For constant stepsizes the performance seems to plateau rather than diverge. With reservoir sampling we can achieve an effective stepsize of $1/T$. However, the results suggest that exponentially-averaged reservoir sampling can be a viable choice too, as exponential averaging of past memories would approximately correspond to using a constant stepsize.

XFP with stepsize 1 is equivalent to a full-width iterated best response algorithm. While this algorithm converges to a Nash equilibrium in finite perfect-information two-player zero-sum games, the results suggest that with imperfect information this is not generally the case. The Poker-CNN algorithm introduced by Yakovenko et al. (2016) stores a small number of past strategies which it iteratively computes new strategies against. Replacing strategies in that set is similar to updating an average strategy with a large stepsize. This might lead to similar problems as shown in Figure 1.

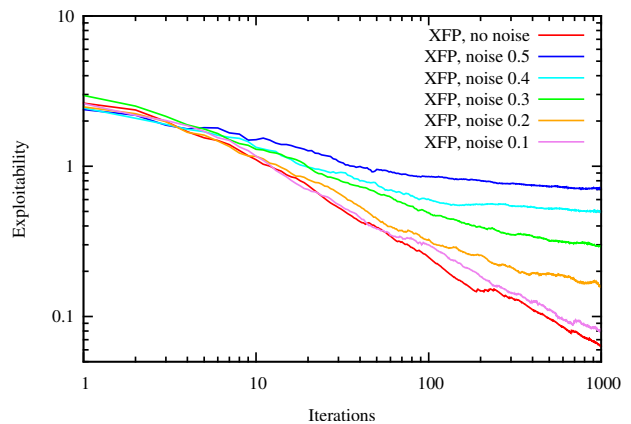


Figure 2. The performance of XFP in Leduc Hold'em with uniform-random noise added to the best response computation.

Our NFSP agents add random exploration to their policies and use noisy stochastic gradient updates to learn action values, which determine their approximate best responses. Therefore, we investigated the impact of random noise added to the best response computation, which XFP performs by dynamic programming. At each backward induction step, we pass back a uniform-random action’s value with probability ϵ and the best action’s value otherwise. Figure 2 shows monotonically decreasing performance with added noise. However, performance remains stable and keeps improving for all noise levels.

4.2. Convergence of NFSP

We empirically investigate the convergence of NFSP to Nash equilibria in Leduc Hold’em. We also study whether removing or altering some of NFSP’s components breaks convergence.

One of our goals is to minimize reliance on prior knowledge. Therefore, we attempt to define an objective encoding of information states in poker games. Contrary to other work on computer poker (Zinkevich et al., 2007; Gilpin et al., 2007; Johanson et al., 2013), we do not engineer any higher-level features. Poker games usually consist of multiple rounds. At each round new cards are revealed to the players. We represent each rounds’ cards by a k-of-n encoding. E.g. LHE has a card deck of 52 cards and on the second round three cards are revealed. Thus, this round is encoded with a vector of length 52 and three elements set to 1 and the rest to 0. In Limit Hold’em poker games, players usually have three actions to choose from, namely {fold, call, raise}. Note that depending on context, calls and raises can be referred to as checks and bets respectively. Betting is capped at a fixed number of raises per round. Thus, we can represent the betting history as a tensor with 4 dimensions, namely {player, round, number of raises, action taken}. E.g. heads-up LHE contains 2 players, 4 rounds, 0 to 4 raises per round and 3 actions. Thus we can represent a LHE betting history as a $2 \times 4 \times 5 \times 3$ tensor. In a heads-up game we do not need to encode the fold action, as a two-player game always ends if one players gives up. Thus, we can flatten the 4-dimensional tensor to a vector of length 80. Concatenating with the card inputs of 4 rounds, we encode an information state of LHE as a vector of length 288. Similarly, an information state of Leduc Hold’em can be encoded as a vector of length 30, as it contains 6 cards with 3 duplicates, 2 rounds, 0 to 2 raises per round and 3 actions.

For learning in Leduc Hold’em, we manually calibrated NFSP for a fully connected neural network with 1 hidden layer of 64 neurons and rectified linear activations. We then repeated the experiment for various network architectures with the same parameters. In particular, we set the sizes

of memories to 200k and 2m for \mathcal{M}_{RL} and \mathcal{M}_{SL} respectively. \mathcal{M}_{RL} functioned as a circular buffer containing a recent window of experience. \mathcal{M}_{SL} was updated with reservoir sampling (Vitter, 1985). The reinforcement and supervised learning rates were set to 0.1 and 0.005, and both used vanilla Stochastic Gradient Descent (SGD) without momentum for stochastic optimization of the neural networks. Each agent performed 2 stochastic gradient updates of mini-batch size 128 per network for every 128 steps in the game. The target network of the DQN algorithm was refitted every 300 updates. NFSP’s anticipatory parameter was set to $\eta = 0.1$. The ϵ -greedy policies’ exploration started at 0.06 and decayed to 0, proportionally to the inverse square root of the number of iterations.

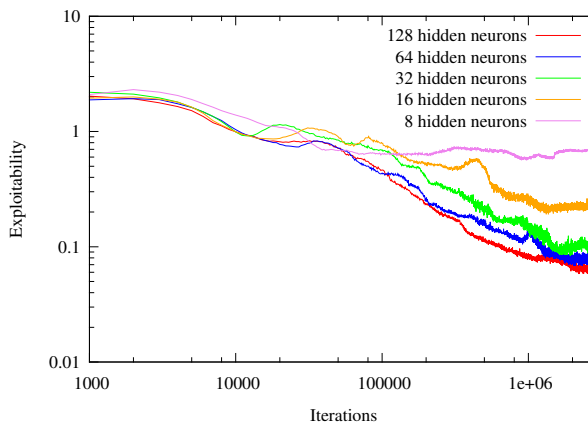


Figure 3. Learning performance of NFSP in Leduc Hold’em.

Figure 3 shows NFSP approaching Nash equilibria for various network architectures. We observe a monotonic performance increase with size of the networks. NFSP achieved an exploitability of 0.06, which full-width XFP typically achieves after around 1000 full-width iterations.

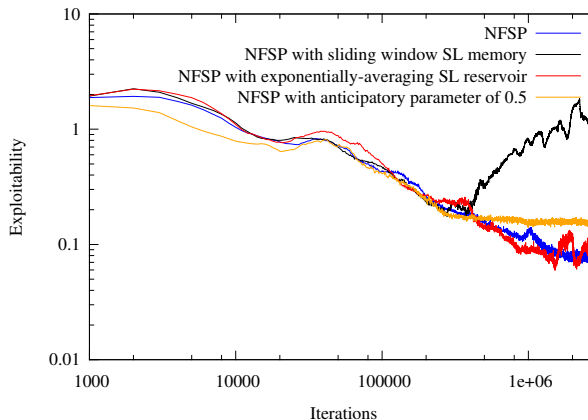


Figure 4. Breaking learning performance in Leduc Hold’em by removing essential components of NFSP.

In order to investigate the relevance of various components of NFSP, e.g. reservoir sampling and anticipatory dynamics, we conducted an experiment that isolated their effects. Figure 4 shows that these modifications led to decremental performance. In particular, using a fixed-size sliding window to store experience of the agents’ own behaviour led to divergence. For a high anticipatory parameter of 0.5 NFSP’s performance plateaued. Finally, using exponentially-averaged reservoir sampling for supervised learning memory updates led to noisy performance.

4.3. Comparison to DQN

Several stable algorithms have previously been proposed for deep reinforcement learning, notably the DQN algorithm (Mnih et al., 2015). However, the empirical stability of these algorithms was only previously established in single-agent, perfect (or near-perfect) information MDPs. Here, we investigate the stability of DQN in multi-agent, imperfect-information games, in comparison to NFSP.




Figure 5. Comparing performance to DQN in Leduc Hold’em.

DQN learns a deterministic, greedy strategy. This is sufficient to behave optimally in MDPs, which the algorithm is designed for. Imperfect-information games, on the other hand, generally require stochastic strategies for optimal behaviour. Thus, in addition to DQN’s ϵ -greedy strategy, we store its actions in a supervised learning memory, \mathcal{M}_{SL} , and learn its average behaviour. This average policy does not affect DQN’s runtime behaviour at all, as it is never played. We implement this variant of DQN by using NFSP with an anticipatory parameter of $\eta = 1$. We set most of DQN’s parameters to be equal to those found for NFSP in the previous section’s experiments. This is motivated by the supervised learning parameters not directly affecting DQN’s performance. We trained DQN with all combinations of the following parameters: Learning rate $\{0.2, 0.1, 0.05\}$, decaying exploration starting at $\{0.06, 0.12\}$ and reinforcement learning mem-

ory $\{2m \text{ reservoir}, 2m \text{ sliding window}\}$. We then chose the best-performing result of DQN and compared to the performance of NFSP that was achieved in the previous section’s experiment. DQN achieved its best performing result with a learning rate of 0.1, exploration starting at 0.12 and a sliding window memory of size 2m.

Figure 5 shows that DQN’s deterministic strategy is highly exploitable, which is expected as imperfect-information games usually require stochastic policies. DQN’s average behaviour does not approach a Nash equilibrium either. This is notable because DQN stores its experience in a replay memory and thus would effectively learn against the opponents’ average behaviour as long as its replay memory is big enough to keep track of it. This is quite similar to a fictitious play. However, because DQN agents use their ϵ -greedy strategies in self-play their experience is highly correlated over time and focussed on only a subset of states. We believe this is the main reason for NFSP’s superior performance in our experiments. NFSP agents use an ever more slowly changing average policy in self-play. Thus, their experience varies more slowly, resulting in a more stable data distribution contained in their memories. This might help their training of neural networks and adaptation to each other. Other common reinforcement learning methods have been shown to exhibit similarly stagnating performance in poker games (Ponsen et al., 2011; Heinrich & Silver, 2015).

4.4. Limit Texas Hold’em

We applied NFSP to LHE, a game that is popular with humans. Since in 2008 a computer program beat expert human LHE players for the first time in a public competition, modern computer agents are widely considered to have achieved super-human performance (Newall, 2013). The game was *essentially solved* by Bowling et al. (2015). We evaluated our agents against SmooCT, a Smooth UCT (Heinrich & Silver, 2015) agent which achieved 3 silver medals in the Annual Computer Poker Competition (ACPC) in 2014. Learning performance was measured in milli-big-blinds won per hand, mbb/h, i.e. one thousandth of a big blind that players post at the beginning of a hand.

We manually calibrated NFSP by trying 9 configurations. We achieved the best performance with the following parameters. The neural networks were fully connected with four hidden layers of 1024, 512, 1024 and 512 neurons with rectified linear activations. The memory sizes were set to 600k and 30m for \mathcal{M}_{RL} and \mathcal{M}_{SL} respectively. \mathcal{M}_{RL} functioned as a circular buffer containing a recent window of experience. \mathcal{M}_{SL} was updated with exponentially-averaged reservoir sampling (Osborne et al., 2014), replacing entries in \mathcal{M}_{SL} with minimum probability 0.25. We used vanilla SGD without momentum for both reinforce-

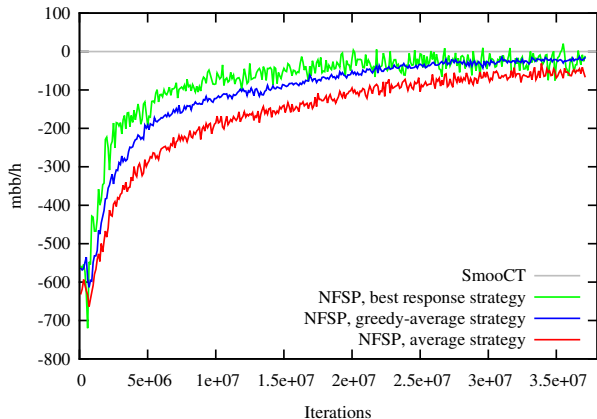


Figure 6. Performance of playing against SmooCT. The estimated standard error of each evaluation is less than 10 mbb/h.

ment and supervised learning, with learning rates set to 0.1 and 0.01 respectively. Each agent performed 2 stochastic gradient updates of mini-batch size 256 per network for every 256 steps in the game. The target network of the DQN algorithm was refitted every 1000 updates. NFSP’s anticipatory parameter was set to $\eta = 0.1$. The ϵ -greedy policies’ exploration started at 0.08 and decayed to 0, more slowly than in Leduc Hold’em. In addition to NFSP’s main, average strategy profile we also evaluated the best response and greedy-average strategies, which deterministically choose actions that maximize the predicted action values or probabilities respectively.

To provide some intuition for win rates in heads-up LHE, a player that always folds will lose 750 mbb/h, and expert human players typically achieve expected win rates of 40-60 mbb/h at online high-stakes games. Similarly, the top half of computer agents in the ACPC 2014 achieved up to 50 mbb/h between themselves. While training, we periodically evaluated NFSP’s performance against SmooCT from symmetric play for 25000 hands each. Figure 6 presents the learning performance of NFSP. NFSP’s average and greedy-average strategy profiles exhibit a stable and relatively monotonic performance improvement, and achieve win rates of around -50 and -20 mbb/h respectively. The best response strategy profile exhibited more noisy performance, mostly ranging between -50 and 0 mbb/h. We also evaluated the final greedy-average strategy against the other top 3 competitors of the ACPC 2014. Table 1 presents the results.

5. Related work

Reliance on human expert knowledge can be expensive, prone to human biases and limiting if such knowledge is suboptimal. Yet many methods that have been ap-

Match-up	NFSP
escabeche	-52.1 ± 8.5
SmooCT	-17.4 ± 9.0
Hyperborean	-13.6 ± 9.2

Table 1. Performance of NFSP’s greedy-average strategy against the top 3 agents of the ACPC 2014.

plied to games have relied on human expert knowledge. Deep Blue used a human-engineered evaluation function for chess (Campbell et al., 2002). In computer Go, Maddison et al. (2015) and Clark & Storkey (2015) trained deep neural networks from data of expert human play. In computer poker, current game-theoretic approaches use heuristics of card strength to abstract the game to a tractable size (Zinkevich et al., 2007; Gilpin et al., 2007; Johanson et al., 2013). Waugh et al. (2015) recently combined one of these methods with function approximation. However, their full-width algorithm has to implicitly reason about all information states at each iteration, which is prohibitively expensive in large domains. In contrast, NFSP focuses on the sample-based reinforcement learning setting where the game’s states need not be exhaustively enumerated and the learner may not even have a model of the game’s dynamics.

Many successful applications in games have relied on local search (Campbell et al., 2002; Browne et al., 2012). Local search algorithms efficiently plan decisions in a game at runtime, e.g. via Monte Carlo simulation or limited-depth backward induction. However, common simulation-based local search algorithms have been shown to diverge when applied to imperfect-information poker games (Ponsen et al., 2011; Heinrich & Silver, 2015). Furthermore, even game-theoretic methods do not generally achieve unexploitable behaviour when planning locally in imperfect-information games (Burch et al., 2014; Ganzfried & Sandholm, 2015; Lisý et al., 2015). Another problem of local search is the potentially prohibitive cost at runtime if no prior knowledge is injected to guide the search. This poses the question of how to obtain this prior knowledge. Silver et al. (2016) trained convolutional neural networks on human expert data and then used a self-play reinforcement learning procedure to optimize these networks further. By using these neural networks to guide a high-performance local search, they beat a Go grandmaster 5 to 0. In this work, we evaluate our agents without any local search at runtime. If local search methods for imperfect-information games were developed, strategies trained by NFSP could be a promising choice for guiding the search.

Nash equilibria are the only strategy profiles that rational agents can hope to converge on in self-play (Bowling & Veloso, 2001). TD-Gammon (Tesauro, 1995) is a world-class backgammon agent, whose main component is a neu-

ral network trained from self-play reinforcement learning. While its algorithm, based on temporal-difference learning, is sound in two-player zero-sum perfect-information games, it does not generally converge in games with imperfect information. DQN (Mnih et al., 2015) combines temporal-difference learning with experience replay and deep neural network function approximation. It achieved human-level performance in a majority of Atari games, learning from raw sensory inputs. However, these Atari games were set up as single-agent domains with potential opponents fixed and controlled by the Atari emulator. Our experiments showed that DQN agents were unable to approach a Nash equilibrium in Leduc Hold'em, where players were allowed to adapt dynamically. Yakovenko et al. (2016) trained deep neural networks in self-play in computer poker, including two poker games that are popular with humans. Their networks performed strongly against heuristic-based and simple computer programs. Expert human players were able to outperform their agent, albeit over a statistically insignificant sample size. It remains to be seen whether their approach converges in practice or theory. In contrast, we have empirically shown NFSP's convergence to approximate Nash equilibria in Leduc Hold'em. Furthermore, the approach is principled and builds on the theory of fictitious play in extensive-form games.

6. Conclusion

We have introduced NFSP, the first end-to-end deep reinforcement learning approach to learning approximate Nash equilibria of imperfect-information games from self-play. NFSP addresses three problems. Firstly, NFSP agents learn without prior knowledge. Secondly, they do not rely on local search at runtime. Thirdly, they converge to approximate Nash equilibria in self-play. Our empirical results provide the following insights. The performance of fictitious play degrades gracefully with various approximation errors. NFSP converges reliably to approximate Nash equilibria in a small poker game, whereas DQN's greedy and average strategies do not. NFSP learned a competitive strategy in a real-world scale imperfect-information game from scratch without using explicit prior knowledge.

In this work, we focussed on imperfect-information two-player zero-sum games. Fictitious play, however, is also guaranteed to converge to Nash equilibria in cooperative, potential games. It is therefore conceivable that NFSP can be successfully applied to these games as well. Furthermore, recent developments in continuous-action reinforcement learning (Lillicrap et al., 2015) could enable NFSP to be applied to continuous-action games, which current game-theoretic methods cannot deal with directly.

Acknowledgements

We thank Peter Dayan, Marc Lanctot and Marc Bellemare for helpful discussions and feedback. This research was supported by the UK Centre for Doctoral Training in Financial Computing and by the NVIDIA Corporation.

References

- Ariely, Dan and Jones, Simon. *Predictably irrational*. HarperCollins New York, 2008.
- Bazzan, Ana LC. Opportunities for multiagent systems and multiagent reinforcement learning in traffic control. *Autonomous Agents and Multi-Agent Systems*, 18(3):342–375, 2009.
- Bowling, Michael and Veloso, Manuela. Rational and convergent learning in stochastic games. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, volume 17, pp. 1021–1026, 2001.
- Bowling, Michael, Burch, Neil, Johanson, Michael, and Tammelin, Oskari. Heads-up limit holdem poker is solved. *Science*, 347(6218):145–149, 2015.
- Brown, George W. Iterative solution of games by fictitious play. *Activity analysis of production and allocation*, 13(1):374–376, 1951.
- Browne, Cameron B, Powley, Edward, Whitehouse, Daniel, Lucas, Simon M, Cowling, Peter I, Rohlfshagen, Philipp, Tavener, Stephen, Perez, Diego, Samothrakis, Spyridon, and Colton, Simon. A survey of Monte Carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1):1–43, 2012.
- Burch, Neil, Johanson, Michael, and Bowling, Michael. Solving imperfect information games using decomposition. In *28th AAAI Conference on Artificial Intelligence*, 2014.
- Campbell, Murray, Hoane, A Joseph, and Hsu, Feng-hsiung. Deep blue. *Artificial intelligence*, 134(1):57–83, 2002.
- Clark, Christopher and Storkey, Amos. Training deep convolutional neural networks to play go. In *Proceedings of the 32nd International Conference on Machine Learning*, pp. 1766–1774, 2015.
- Durkota, Karel, Lisý, Viliam, Bošanský, Branislav, and Kiekintveld, Christopher. Optimal network security hardening using attack graph games. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, 2015.

- Ernst, Damien, Geurts, Pierre, and Wehenkel, Louis. Tree-based batch mode reinforcement learning. In *Journal of Machine Learning Research*, pp. 503–556, 2005.
- Fudenberg, Drew. *The theory of learning in games*, volume 2. MIT press, 1998.
- Ganzfried, Sam and Sandholm, Tuomas. Endgame solving in large imperfect-information games. In *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems*, 2015.
- Gelly, Sylvain, Kocsis, Levente, Schoenauer, Marc, Sebag, Michèle, Silver, David, Szepesvári, Csaba, and Teytaud, Olivier. The grand challenge of computer go: Monte Carlo tree search and extensions. *Communications of the ACM*, 55(3):106–113, 2012.
- Gilpin, Andrew, Hoda, Samid, Pena, Javier, and Sandholm, Tuomas. Gradient-based algorithms for finding Nash equilibria in extensive form games. In *Internet and Network Economics*, pp. 57–69. Springer, 2007.
- Heinrich, Johannes and Silver, David. Smooth UCT search in computer poker. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, 2015.
- Heinrich, Johannes, Lanctot, Marc, and Silver, David. Fictitious self-play in extensive-form games. In *Proceedings of the 32nd International Conference on Machine Learning*, 2015.
- Johanson, Michael, Burch, Neil, Valenzano, Richard, and Bowling, Michael. Evaluating state-space abstractions in extensive-form games. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems*, pp. 271–278, 2013.
- Lambert III, Theodore J, Epelman, Marina A, and Smith, Robert L. A fictitious play approach to large-scale optimization. *Operations Research*, 53(3):477–489, 2005.
- Leslie, David S and Collins, Edmund J. Generalised weakened fictitious play. *Games and Economic Behavior*, 56(2):285–298, 2006.
- Lillicrap, Timothy P, Hunt, Jonathan J, Pritzel, Alexander, Heess, Nicolas, Erez, Tom, Tassa, Yuval, Silver, David, and Wierstra, Daan. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Lin, Long-Ji. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine learning*, 8(3-4):293–321, 1992.
- Lisý, Viliam, Lanctot, Marc, and Bowling, Michael. Online monte carlo counterfactual regret minimization for search in imperfect information games. In *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems*, 2015.
- Maddison, Chris J, Huang, Aja, Sutskever, Ilya, and Silver, David. Move evaluation in go using deep convolutional neural networks. *The International Conference on Learning Representations*, 2015.
- Mnih, Volodymyr, Kavukcuoglu, Koray, Silver, David, Rusu, Andrei A, Veness, Joel, Bellemare, Marc G, Graves, Alex, Riedmiller, Martin, Fidjeland, Andreas K, Ostrovski, Georg, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- Monderer, Dov and Shapley, Lloyd S. Fictitious play property for games with identical interests. *Journal of economic theory*, 68(1):258–265, 1996.
- Myerson, Roger B. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.
- Nevmyvaka, Yuriy, Feng, Yi, and Kearns, Michael. Reinforcement learning for optimized trade execution. In *Proceedings of the 23rd International Conference on Machine Learning*, pp. 673–680. ACM, 2006.
- Newall, P. *Further Limit Hold 'em: Exploring the Model Poker Game*. Two Plus Two Publishing, LLC, 2013.
- Osborne, Miles, Lall, Ashwin, and Van Durme, Benjamin. Exponential reservoir sampling for streaming language models. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, pp. 687–692, 2014.
- Ponsen, Marc, de Jong, Steven, and Lanctot, Marc. Computing approximate Nash equilibria and robust best-responses using sampling. *Journal of Artificial Intelligence Research*, 42(1):575–605, 2011.
- Riedmiller, Martin. Neural fitted q iteration—first experiences with a data efficient neural reinforcement learning method. In *Machine Learning: ECML 2005*, pp. 317–328. Springer, 2005.
- Riedmiller, Martin, Gabel, Thomas, Hafner, Roland, and Lange, Sascha. Reinforcement learning for robot soccer. *Autonomous Robots*, 27(1):55–73, 2009.
- Robinson, Julia. An iterative method of solving a game. *Annals of Mathematics*, pp. 296–301, 1951.
- Samuel, Arthur L. Some studies in machine learning using the game of checkers. *IBM Journal of research and development*, 3(3):210–229, 1959.

- Selten, Reinhard. Bounded rationality. *Journal of Institutional and Theoretical Economics*, pp. 649–658, 1990.
- Shamma, Jeff S and Arslan, Gürdal. Dynamic fictitious play, dynamic gradient play, and distributed convergence to Nash equilibria. *IEEE Transactions on Automatic Control*, 50(3):312–327, 2005.
- Silver, David, Huang, Aja, Maddison, Chris J., Guez, Arthur, Sifre, Laurent, van den Driessche, George, Schrittwieser, Julian, Antonoglou, Ioannis, Panneershelvam, Veda, Lanctot, Marc, Dieleman, Sander, Grewe, Dominik, Nham, John, Kalchbrenner, Nal, Sutskever, Ilya, Lillicrap, Timothy, Leach, Madeleine, Kavukcuoglu, Koray, Graepel, Thore, and Hassabis, Demis. Mastering the game of go with deep neural networks and tree search. *Nature*, 529:484–489, 2016.
- Southey, Finnegan, Bowling, Michael, Larson, Bryce, Piccione, Carmelo, Burch, Neil, Billings, Darse, and Rayner, Chris. Bayes bluff: Opponent modelling in poker. In *In Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence*, pp. 550–558, 2005.
- Sutton, Richard S and Barto, Andrew G. *Reinforcement learning: An introduction*, volume 1. Cambridge Univ Press, 1998.
- Tambe, Milind. *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge University Press, 2011.
- Tesauro, Gerald. Temporal difference learning and td-gammon. *Communications of the ACM*, 38(3):58–68, 1995.
- Urieli, Daniel and Stone, Peter. Tactex’13: a champion adaptive power trading agent. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems*, pp. 1447–1448, 2014.
- Vitter, Jeffrey S. Random sampling with a reservoir. *ACM Transactions on Mathematical Software (TOMS)*, 11(1): 37–57, 1985.
- Von Stengel, Bernhard. Efficient computation of behavior strategies. *Games and Economic Behavior*, 14(2):220–246, 1996.
- Watkins, Christopher JCH and Dayan, Peter. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- Waugh, Kevin, Morrill, Dustin, Bagnell, J. Andrew, and Bowling, Michael. Solving games with functional regret estimation. In *29th AAAI Conference on Artificial Intelligence*, 2015.
- Yakovenko, Nikolai, Cao, Liangliang, Raffel, Colin, and Fan, James. Poker-cnn: A pattern learning strategy for making draws and bets in poker games using convolutional networks. In *30th AAAI Conference on Artificial Intelligence*, 2016.
- Zinkevich, Martin, Johanson, Michael, Bowling, Michael, and Piccione, Carmelo. Regret minimization in games with incomplete information. In *Advances in Neural Information Processing Systems*, pp. 1729–1736, 2007.