
Regret Minimization for Partially Observable Deep Reinforcement Learning

Peter Jin¹ Kurt Keutzer¹ Sergey Levine¹

Abstract

Deep reinforcement learning algorithms that estimate state and state-action value functions have been shown to be effective in a variety of challenging domains, including learning control strategies from raw image pixels. However, algorithms that estimate state and state-action value functions typically assume a fully observed state and must compensate for partial observations by using finite length observation histories or recurrent networks. In this work, we propose a new deep reinforcement learning algorithm based on counterfactual regret minimization that iteratively updates an approximation to an advantage-like function and is robust to partially observed state. We demonstrate that this new algorithm can substantially outperform strong baseline methods on several partially observed reinforcement learning tasks: learning first-person 3D navigation in Doom and Minecraft, and acting in the presence of partially observed objects in Doom and Pong.

1. Introduction

Many reinforcement learning problems of practical interest have the property of partial observability, where observations of state are generally non-Markovian. Practical deep reinforcement learning algorithms fall into two broad classes, neither of which satisfactorily deals with partial observability despite the prevalence of partial observations in the real world. One class of algorithms consists of value function-based methods such as deep Q-learning (Mnih et al., 2013; 2015), which are known to be highly sample efficient but generally assume a Markovian observation space. The other class of algorithms consists of Monte Carlo policy gradient methods such as trust region policy optimization (Schulman et al., 2015), which do not need to assume Markovian observations but are less sample efficient than value

function-based methods. Some policy gradient methods such as advantage actor-critic (Mnih et al., 2016) introduce the Markov assumption through a critic or state-dependent baseline to improve sample efficiency.

There are two common workarounds for the problem of partial observation: (a) learning policies and value functions on finite length observation histories, and (b) learning recurrent policies and recurrent value functions. Finite length observation histories concatenate the most recent raw observations into a stack of observations, and are a simple but effective technique to approximate the appearance of full observability. Recurrent functions can potentially incorporate an infinite observation history, but they can be harder to optimize. When using either finite length observation histories or recurrent functions, the same value function-based methods and policy gradient methods are employed with their respective tradeoffs in partially observable domains.

We are interested in developing methods that combine the best of both value function-based methods and policy gradient methods for partial observable domains. That is, can we develop methods that are sample efficient but are also robust to partial observation spaces?

Our contribution is a new model-free deep reinforcement learning algorithm based on the principle of regret minimization which does not require access to a Markovian state. Our method learns a policy by estimating an advantage-like function which approximates a quantity called the counterfactual regret. Counterfactual regret is central to the family of counterfactual regret minimization (CFR) (Zinkevich et al., 2007) algorithms for solving incomplete information games. Hence we call our algorithm “advantage-based regret minimization” (ARM).

We evaluate our approach on three partially observable visual reinforcement learning tasks: first-person 3D navigation in Doom and Minecraft (Kempka et al., 2016; Johnson et al., 2016), avoiding partially observed objects in Doom, and playing Pong with partial observation perturbations (Bellemare et al., 2013). In our experiments, we find that our method is more robust and offers substantial improvement over prior methods on partially observable tasks.

¹Department of Electrical Engineering and Computer Sciences, University of California, Berkeley. Correspondence to: Peter Jin <phj@eecs.berkeley.edu>.

2. Related Work

Deep reinforcement learning algorithms have been demonstrated to achieve excellent results on a range of complex tasks, including playing games (Mnih et al., 2015; Oh et al., 2016) and continuous control (Schulman et al., 2015; Lillicrap et al., 2016; Levine et al., 2016). Prior deep reinforcement learning algorithms either learn state or state-action value functions (Mnih et al., 2013), learn policies using policy gradients (Schulman et al., 2015), or perform a combination of the two using actor-critic architectures (Mnih et al., 2016). Policy gradient methods typically do not need to assume a Markovian state, but tend to suffer from poor sample complexity, due to their inability to use off-policy data. Methods based on learning Q-functions can use replay buffers to include off-policy data, accelerating learning (Lillicrap et al., 2016). However, learning Q-functions with Bellman error minimization typically requires a Markovian state space. When learning from partial observations such as images, the inputs might not be Markovian. Mnih et al. (2013) introduced the concatenation of short observation sequences. Prior methods (Hausknecht & Stone, 2017; Oh et al., 2016; Mnih et al., 2016; Heess et al., 2015) have proposed to mitigate this issue by learning recurrent critics and Q-functions, that depend on entire histories of observations. The recurrent deterministic policy gradient (Heess et al., 2015) for partially observable control uses a similar DDPG-style estimator as used in ARM (see Section 3.3). However, all of these changes increase the size of the input space, increase variance, or make the optimization problem more complex. Our method instead learns cumulative advantage functions that depend only on the current observation, but can still handle non-Markovian problems.

The form of our advantage function update resembles positive temporal difference methods (Peng et al., 2016; van Hasselt & Wiering, 2007). Additionally, our update rule for a modified cumulative Q-function resembles the average Q-function (Anschel et al., 2017) used for variance reduction in Q-learning. In both cases, the theoretical foundations of our method are based on cumulative regret minimization, and the motivation is substantively different. Previous work by Ross et al. (2011); Ross & Bagnell (2014) has connected regret minimization to reinforcement learning, imitation learning, and structured prediction, although not with counterfactual regret minimization. Regression regret matching (Waugh et al., 2015) is based on a closely related idea, which is to directly approximate the regret with a linear regression model, however the use of a linear model is limited in representation compared to deep function approximation.

Finally, regret and value functions both typically take the form of expectations in reinforcement learning. An alternative view of value functions in RL is through the lens of value distributions (Bellemare et al., 2017). A connec-

tion between regret and distributional RL could lead to very interesting future work.

3. Advantage-based Regret Minimization

In this section, we provide some background on counterfactual regret minimization, describe ARM in detail, and give some intuition for why ARM works.

3.1. Counterfactual Regret Minimization (CFR)

First, we give a reinforcement learning-centric exposition of counterfactual regret minimization (CFR) (Zinkevich et al., 2007; Bowling et al., 2015).

The CFR model for partial observation is as follows. Let \mathcal{S} be the state space. Let \mathcal{I} be the space of *information sets*: an information set $I \in \mathcal{I}$ is a set of states $s \in I$, where $s \in \mathcal{S}$, such that only the information set I is directly observable, and the individual states contained in I are hidden. An information set I is therefore a kind of *aggregate state*.

In CFR, an extensive-form game is repeatedly played between N players, indexed by i . We consider an iterative learning setting, where at the t -th learning iteration, the i -th player follows a fixed policy π_t^i , choosing actions $a \in \mathcal{A}$ conditioned on information sets $I \in \mathcal{I}$ according to their policy $\pi_t^i(a|I)$. Let σ_t denote the players' joint policy (their strategy profile). For any strategy profile σ , the i -th player has an expected value for playing the game, $J^i(\sigma)$. Let σ_t^{-i} represent the joint policy of all players except the i -th player; similarly, let the tuple (π_t^i, σ_t^{-i}) be the product of the i -th player's policy π_t^i and the other players' joint policy σ_t^{-i} .

The i -th player's *overall regret* after T learning iterations is defined:

$$(R_T^i)^{(\text{overall})} = \max_{\pi_*^i} \sum_{t=1}^T J^i(\pi_*^i, \sigma_t^{-i}) - J^i(\pi_t^i, \sigma_t^{-i}). \quad (1)$$

What is the interpretation of the overall regret? It is essentially how much better the i -th player could have done, had it always followed an optimal policy in hindsight instead of the actual sequence of policies π_t^i it executed. Intuitively, the difference $J^i(\pi_*^i, \sigma_t^{-i}) - J^i(\pi_t^i, \sigma_t^{-i})$ is the suboptimality of π_t^i compared to π_*^i , and the sum of the suboptimalities over learning iterations yields the area inside the learning curve. In other words, a smaller overall regret implies better sample efficiency.

Let $Q_{\sigma_t}^i(I, a)$ be the *counterfactual value* of the i -th player, where the i -th player is assumed to reach I and always chooses the action a in the aggregate state I , and otherwise follows the policy π_t^i , while all other players follow the strategy profile σ_t^{-i} . Similarly, let the expected counterfactual value be calculated as $V_{\sigma_t}^i(I) = \sum_{a \in \mathcal{A}} \pi_t^i(a|I) Q_{\sigma_t}^i(I, a)$.

Let $(R_T^i)^{\text{CF}}(I, a)$ be the *counterfactual regret* of the i -th player, which is the sum of the advantage-like quantities $Q_{\sigma_t}^i(I, a) - V_{\sigma_t}^i(I)$ after T learning iterations:

$$(R_T^i)^{\text{CF}}(I, a) = \sum_{t=1}^T Q_{\sigma_t}^i(I, a) - V_{\sigma_t}^i(I). \quad (2)$$

Similarly, the *immediate counterfactual regret* can be obtained from the counterfactual regret by maximization over the player’s action space:

$$(R_T^i)^{\text{immCF}}(I) = \max_{a \in \mathcal{A}} \sum_{t=1}^T Q_{\sigma_t}^i(I, a) - V_{\sigma_t}^i(I). \quad (3)$$

The immediate counterfactual regret (Equation (3)) possesses a similar interpretation to the overall regret (Equation (1)), except that the immediate counterfactual regret and its constituent terms are additionally functions of the aggregate state I .

Suppose that one were to briefly consider each aggregate state I as a separate, independent subproblem. By naively treating the counterfactual regret $(R_T^i)^{\text{CF}}(I, a)$ for each I as analogous to regret in an online learning setting, then at each learning iteration one may simply plug the counterfactual regret into a regret matching policy update (Hart & Mas-Colell, 2000):

$$(\pi_{t+1}^i)^{\text{RM}}(a|I) = \frac{\max(0, (R_t^i)^{\text{CF}}(I, a))}{\sum_{a' \in \mathcal{A}} \max(0, (R_t^i)^{\text{CF}}(I, a'))}. \quad (4)$$

In the online learning setting, the regret matching policy achieves regret that increases with rate $O(\sqrt{T})$ in the number of iterations T (Cesa-Bianchi & Lugosi, 2003; Gordon, 2007); we can then say that the regret matching policy is *regret minimizing*.

It turns out that updating players’ policies in the extensive game setting by iteratively minimizing the immediate counterfactual regrets according to Equation (4) will also minimize the overall regret $(R_T^i)^{\text{(overall)}}$ with upper bound $O(|\mathcal{I}| \sqrt{|\mathcal{A}| T})$. The overall regret’s $O(\sqrt{T})$ dependence on the number of iterations T is not impacted by the structure of the information set space \mathcal{I} , which is why CFR can be said to be robust to a certain kind of partial observability. This result forms the basis of the counterfactual regret minimization algorithm and was proved by Zinkevich et al. (2007).

Since we are interested in the application of CFR to reinforcement learning, we can write down “1-player” versions of the components above: the counterfactual value, reinterpreted as a *stationary* state-action value function $Q_{\pi|I \rightarrow a}(I, a)$, where the action a is always chosen in the aggregate state I , and the policy π is otherwise followed

(Bellemare et al., 2016); the counterfactual regret, including its recursive definition:

$$\begin{aligned} R_T^{\text{CF}}(I, a) &= \sum_{t=1}^T Q_{\pi_t|I \rightarrow a}(I, a) - V_{\pi_t|I}(I) \\ &= R_{T-1}^{\text{CF}}(I, a) + Q_{\pi_T|I \rightarrow a}(I, a) - V_{\pi_T|I}(I) \end{aligned} \quad (5)$$

where $V_{\pi_t|I}(I) = \sum_{a \in \mathcal{A}} \pi_t(a|I) Q_{\pi_t|I \rightarrow a}(I, a)$; and the regret matching policy update:

$$\pi_{t+1}^{\text{RM}}(a|I) = \frac{\max(0, R_t^{\text{CF}}(I, a))}{\sum_{a' \in \mathcal{A}} \max(0, R_t^{\text{CF}}(I, a'))}. \quad (7)$$

3.2. CFR+

CFR+ (Tammelin, 2014) consists of a modification to CFR, in which instead of calculating the full counterfactual regret as in Equation (6), instead the counterfactual regret is recursively positively clipped:

$$\begin{aligned} R_T^{\text{CF+}}(I, a) & \\ &= [R_{T-1}^{\text{CF+}}(I, a)]_+ + Q_{\pi_T|I \rightarrow a}(I, a) - V_{\pi_T|I}(I) \end{aligned} \quad (8)$$

where $[x]_+ = \max(0, x)$ is the positive clipping operator. Comparing Equation (6) with Equation (8), the only difference in CFR+ is that the previous iteration’s quantity is positively clipped in the recursion. This simple change turns out to yield a large practical improvement in the performance of the algorithm (Bowling et al., 2015). One intuition for why the positive clipping of CFR+ can improve upon CFR is that because $[R_{T-1}^{\text{CF+}}]_+$ is nonnegative, it provides a kind of “optimism under uncertainty,” adding a bonus to some transitions while ignoring others. The CFR+ update has also been shown to do better than CFR when the best action in hindsight changes frequently (Tammelin et al., 2015). In the rest of this work we will solely build upon the CFR+ update.

3.3. From CFR and CFR+ to ARM

Recall that CFR and CFR+ model partial observation in extensive games with a space \mathcal{I} of information sets or aggregate states. Because we are interested in general observation spaces \mathcal{O} in partially observed Markov decision processes (MDPs), we naively map CFR and CFR+ onto MDPs and replace \mathcal{I} with \mathcal{O} . It is known that Markov or stochastic games can be converted to extensive form (Littman, 1994; Sandholm & Singh, 2012; Kroer & Sandholm, 2014), and an MDP is a 1-player Markov game. Because we are also interested in high dimensional observations such as images, we estimate a counterfactual regret function approximation (Waugh et al., 2015) parameterized as a neural network.

The exact form of the counterfactual regret (either Equation (6) or Equation (8)), after translating aggregate states I to

observations o , utilizes a stationary action-value function $Q_{\pi|o \rightarrow a}(o, a)$. It is unclear how to learn a truly stationary action-value function, although Bellemare et al. (2016) propose a family of consistent Bellman operators for learning locally stationary action-value functions. We posit that the approximation $Q_{\pi|o \rightarrow a}(o, a) \approx Q_{\pi}(o, a)$, where $Q_{\pi}(o, a)$ is the usual action-value function, is acceptable when observations are rarely seen more than once in a typical trajectory.

The above series of approximations finally yields a learnable function using deep reinforcement learning; this function is a ‘‘clipped cumulative advantage function’’:

$$A_T^+(o, a) = \max(0, A_{T-1}^+(o, a)) + A_{\pi_T}(o, a) \quad (9)$$

where $A_{\pi}(o, a)$ is the usual advantage function. Advantage-based regret minimization (ARM) is then the resulting batch-mode deep RL algorithm that updates the policy to the regret matching distribution on the cumulative clipped advantage function:

$$\pi_{t+1}(a|o) = \frac{\max(0, A_t^+(o, a))}{\sum_{a' \in \mathcal{A}} \max(0, A_t^+(o, a'))}. \quad (10)$$

At the t -th batch iteration of ARM, a batch of data is collected by sampling trajectories using the current policy π_t , followed by two processing steps: (a) fit A_t^+ using Equation (9), then (b) set the next iteration’s policy π_{t+1} using Equation (10).

Below, we will use the subscript k to refer to a timestep within a trajectory, while the subscript t refers to the batch iteration. To implement Equation (9) with deep function approximation, we define two value function approximations, $V_{\pi_t}(o_k; \theta_t)$ and $Q_t^+(o_k, a_k; \omega_t)$, as well as a target value function $V'(o_k; \varphi)$, where θ_t , ω_t , and φ are the learnable parameters. The cumulative clipped advantage function is represented as $A_t^+(o_k, a_k) = Q_t^+(o_k, a_k; \omega_t) - V_{\pi_t}(o_k; \theta_t)$. Within each sampling iteration, the value functions are fitted using stochastic gradient descent by sampling minibatches and performing gradient steps. In practice, we use Adam to perform the optimization (Kingma & Ba, 2015). The state-value function $V_{\pi_t}(o_k; \theta_t)$ is fit using n -step returns with a moving target value function $V'(o_{k+n}; \varphi)$, essentially using the estimator of the deep deterministic policy gradient (DDPG) (Lillicrap et al., 2016). In the same minibatch, $Q_t^+(o_k, a_k; \theta_t)$ is fit to a similar loss, but with an additional target reward bonus that incorporates the previous iteration’s cumulative clipped advantage, $\max(0, A_{t-1}^+(o_k, a_k))$. The regression targets v_k and q_k^+ are defined in terms of the n -step returns $g_k^n = \sum_{k'=k}^{k+n-1} \gamma^{k'-k} r_{k'} + \gamma^n V'(o_{k+n}; \varphi)$:

$$v_k \triangleq g_k^n \quad (11)$$

$$q_k \triangleq r_k + \gamma g_{k+1}^{n-1} \quad (12)$$

$$\phi_k \triangleq Q_{t-1}^+(o_k, a_k; \omega_{t-1}) - V_{\pi_{t-1}}(o_k; \theta_{t-1}) \quad (13)$$

$$q_k^+ \triangleq \max(0, \phi_k) + q_k. \quad (14)$$

Altogether, each minibatch step of the optimization subproblem consists of the following three parameter updates:

$$\theta_t^{(\ell+1)} \leftarrow \theta_t^{(\ell)} - \frac{\alpha}{2} \nabla_{\theta_t^{(\ell)}} (V_{\pi_t}(o_k; \theta_t^{(\ell)}) - v_k)^2 \quad (15)$$

$$\omega_t^{(\ell+1)} \leftarrow \omega_t^{(\ell)} - \frac{\alpha}{2} \nabla_{\omega_t^{(\ell)}} (Q_t^+(o_k, a_k; \omega_t^{(\ell)}) - q_k^+)^2 \quad (16)$$

$$\varphi^{(\ell+1)} \leftarrow \varphi^{(\ell)} + \tau (\theta_t^{(\ell+1)} - \varphi^{(\ell)}). \quad (17)$$

The advantage-based regret minimization algorithm is summarized in Algorithm 1.

We note again that we use a biased value function estimator, whereas only the full returns are guaranteed to be unbiased in non-Markovian settings. In the high-dimensional domains we evaluated on, the most practical choice of deep advantage or value function approximation is based on biased but lower variance estimation, such as the n -step returns we use in practice (Schulman et al., 2016; Gu et al., 2017). We also note that ARM is essentially CFR with different design choices to facilitate function approximation. CFR is proven to converge in domains with non-Markovian information set spaces when using tabular representations, suggesting that our use of biased advantage estimation does not fundamentally preclude the applicability or effectiveness of ARM on non-Markovian domains.

Algorithm 1 Advantage-based regret minimization (ARM).

initialize $\pi_1 \leftarrow$ uniform, $\theta_0, \omega_0 \leftarrow$ arbitrary

for t in $1, \dots$ **do**

collect batch of trajectory data $\mathcal{D}_t \sim \pi_t$

initialize $\theta_t \leftarrow \theta_{t-1}, \omega_t \leftarrow \omega_{t-1}, \varphi \leftarrow \theta_{t-1}$

for ℓ in $0, \dots$ **do**

sample transitions:

$(o_k, a_k, r_k, o_{k+1}) \sim \mathcal{D}_t$

$\delta_k \leftarrow 1 - \mathbb{I}[o_k \text{ is terminal}]$

calculate n -step returns:

$g_k^n \leftarrow \sum_{k'=k}^{k+n-1} \gamma^{k'-k} r_{k'} + \gamma^n \delta_{k+n} V'(o_{k+n}; \varphi)$

calculate target values:

if $t = 0$ **then**

$\phi_k \leftarrow 0$

else

$\phi_k \leftarrow Q_{t-1}^+(o_k, a_k; \omega_{t-1}) - V_{\pi_{t-1}}(o_k; \theta_{t-1})$

end if

$v(o_k) \leftarrow g_k^n$

$q^+(o_k, a_k) \leftarrow \max(0, \phi_k) + r_k + \gamma g_{k+1}^{n-1}$

update θ_t (Equation (15))

update ω_t (Equation (16))

update φ (Equation (17))

end for

set $\pi_{t+1}(a|o) \propto \max(0, Q_t^+(o, a; \omega_t) - V_{\pi_t}(o; \theta_t))$

end for

3.4. ARM vs. Existing Policy Gradient Methods

We note that the fundamental operation in CFR-like algorithms, including ARM, is exemplified by the counterfactual regret update in Equation (6) which superficially looks quite similar to a policy gradient-style update: the update is in the direction of a Q-value minus a baseline. However, we can take the analogy between ARM and policy gradient methods further and show that ARM represents an inherently different update compared to existing policy gradient methods.

Recent work has shown that policy gradient methods and Q-learning methods are connected via maximum entropy RL (O’Donoghue et al., 2017; Haarnoja et al., 2017; Nachum et al., 2017; Schulman et al., 2017). One such perspective is from the soft policy iteration framework for batch-mode reinforcement learning (Haarnoja et al., 2018), where at the t -th batch iteration the updated policy is obtained by minimizing the average KL-divergence between the parametric policy class Π and a target policy f_t . Below is the soft policy iteration update, where the subscript t refers to the batch iteration:

$$\begin{aligned} \pi_{t+1} &= \arg \min_{\pi \in \Pi} \mathbb{E}_{o \sim \rho_t} [D_{\text{KL}}(\pi(\cdot|o) \| f_t(\cdot|o))] \\ &= \arg \min_{\pi \in \Pi} \mathbb{E}_{o \sim \rho_t, a \sim \pi(\cdot|o)} [\log(\pi(a|o)) - \log(f_t(a|o))]. \end{aligned} \quad (18)$$

Using the connection between policy gradient methods and Q-learning, we define the policy gradient target policy as a Boltzmann or softmax distribution on the entropy regularized advantage function $A^{\beta\text{-soft}}$:

$$f_t^{\text{PG}}(a|o) \triangleq \frac{\exp(\beta A_t^{\beta\text{-soft}}(o, a))}{\sum_{a' \in \mathcal{A}} \exp(\beta A_t^{\beta\text{-soft}}(o, a'))}. \quad (19)$$

Now, parameterizing the policy π in terms of an explicit parameter θ , we obtain the expression for the existing policy gradient update, where $b(o)$ is a baseline function:

$$\begin{aligned} \Delta \theta^{\text{PG}} &\propto \mathbb{E}_{o \sim \rho_t, a \sim \pi(\cdot|o; \theta)} \left[\nabla_{\theta} \log(\pi(o|a; \theta)) \cdot \right. \\ &\quad \left(- (1/\beta) \log(\pi(o|a; \theta)) \right. \\ &\quad \left. \left. + A_t^{\beta\text{-soft}}(o, a) - b(o) \right) \right]. \end{aligned} \quad (20)$$

The non-entropy-regularized policy gradient arises in the limit $\beta \rightarrow \infty$, at which point $(1/\beta)$ vanishes.

Note that an alternative choice of target policy f_t will lead to a different kind of policy gradient update. A policy gradient algorithm based on ARM instead proposes a target policy based on the regret matching distribution:

$$f_t^{\text{ARM}}(a|o) \triangleq \frac{\max(0, A_t^+(o, a))}{\sum_{a' \in \mathcal{A}} \max(0, A_t^+(o, a'))}. \quad (21)$$

Similarly, we can express the ARM-like policy gradient update, where again $b(o)$ is a baseline:

$$\begin{aligned} \Delta \theta^{\text{ARM}} &= \mathbb{E}_{o \sim \rho_t, a \sim \pi(\cdot|o; \theta)} \left[\nabla_{\theta} \log(\pi(o|a; \theta)) \cdot \right. \\ &\quad \left(- \log(\pi(o|a; \theta)) \right. \\ &\quad \left. \left. + \log(\max(0, A_t^+(o, a))) - b(o) \right) \right]. \end{aligned} \quad (22)$$

Comparing Equations (20) and (22), we see that the ARM-like policy gradient (Equation (22)) has a logarithmic dependence on the clipped advantage-like function $\max(0, A^+)$, whereas the existing policy gradient (Equation (20)) is only linearly dependent on the advantage function $A^{\beta\text{-soft}}$. This difference in logarithmic vs. linear dependence is responsible for a large part of the inherent distinction of ARM from existing policy gradient methods. In particular, the logarithmic dependence in an ARM-like update may be less sensitive to advantage overestimation compared to existing policy gradient methods, perhaps serving a similar purpose to the double Q-learning estimator (van Hasselt et al., 2016) or consistent Bellman operators (Bellemare et al., 2016).

We also see that for the existing policy gradient (Equation (20)), the $-(1/\beta) \log(\pi(o|a; \theta))$ term, which arises from the policy entropy, is vanishing for large β . On the other hand, for the ARM-like policy gradient (Equation (22)), there is no similar vanishing effect on the equivalent policy entropy term, suggesting that ARM may perform a kind of entropy regularization by default.

In practice we cannot implement an ARM-like policy gradient exactly as in Equation (22), because the positive clipping $\max(0, A^+)$ can yield $\log(0)$. However we believe this is not an intrinsic obstacle and leave the question of how to implement an ARM-like policy gradient to future work.

3.5. Comparing ARM with Other Methods in Partially Observable Domains

The regret matching policy which is fundamental to ARM can be interpreted as a more nuanced form of exploration compared to the epsilon-greedy policy used with Q-learning. In partially observable domains, the optimal policy may generally be stochastic (Jaakkola et al., 1994). So ϵ -greedy policies which put a substantial probability mass on $\arg \max_a Q(o, a)$, e.g. by setting $\epsilon = 0.01$, can be suboptimal, especially compared to more general distributions on discrete actions, such as ARM’s regret matching policy. The softmax policy typically learned by policy gradient methods is also quite general, but can still put too much probability mass on one action without compensation by an explicit entropy bonus as done in maximum entropy reinforcement learning.

Policy gradient methods have an overall regret bound of $R_T \leq B^2/\eta + \eta G^2 T$ derived from stochastic gradient de-

scent, where B is an upper bound on the policy parameter ℓ^2 -norm, G^2 is an upper bound on the second moments of the stochastic gradients, and η is the learning rate (Dick, 2015). Assuming an optimal learning rate $\eta = B/(G\sqrt{T})$, the policy gradient regret bound becomes $R_T \leq 2BG\sqrt{T}$. CFR has an overall regret bound of $R_T \leq \Delta|\mathcal{I}|\sqrt{|\mathcal{A}|T}$, where Δ is the positive range of returns, and $|\mathcal{I}|$ and $|\mathcal{A}|$ are cardinalities of the information set space and action space, respectively (Zinkevich et al., 2007). Both regret bounds have $O(\sqrt{T})$ dependence on T . As mentioned earlier, policy gradient methods can also be robust to partial observation, but unlike in the case of CFR, the dependence of the policy gradient regret bound on an optimal learning rate and on gradient variance may affect policy gradient convergence in the presence of estimation error. On the other hand, the constants in the CFR regret bound are constant properties of the environment.

For Q-learning per se, we are not aware of any known regret bound. Szepesvári proved that an upper bound on the convergence rate of Q-learning (specifically, the convergence rate of $\|Q_T(s, a) - Q^*(s, a)\|_\infty$), assuming a fixed exploration strategy, depends on an exploration condition number, which is the ratio of minimum to maximum state-action occupation frequencies (Szepesvári, 1998), and which describes how “balanced” the exploration strategy is. If partial observability leads to imbalanced exploration due to confounding of states from perceptual aliasing (McCallum, 1997), then Q-learning should be negatively affected in convergence and possibly in absolute performance.

4. Experiments

Because we hypothesize that ARM should perform well in partially observable domains, we conduct our experiments on visual tasks that naturally provide partial observations of state. Our evaluations use feedforward convnets with frame history inputs; our hyperparameters are listed in Section A1 of the Supplementary Material. We are interested in comparing ARM with other advantage-structured methods, primarily: (a) double deep Q-learning with dueling network streams (van Hasselt et al., 2016; Wang et al., 2016), which possesses an advantage-like parameterization of its Q-function and assumes Markovian observations; and (b) TRPO (Schulman et al., 2015; 2016), which estimates an empirical advantage using a baseline state-value function and can handle non-Markovian observations.

4.1. Learning First-person 3D Navigation

We first evaluate ARM on the task of learning first-person navigation in 3D maze-like tasks from two domains: ViZ-Doom (Kempka et al., 2016) based on the game Doom, and Malmö (Johnson et al., 2016) based on the game Minecraft. Doom and Minecraft both feature an egocentric viewpoint,

3D perspective, and rich visual textures. We expect that both domains exhibit a substantial degree of partial observability since only the immediate field-of-view of the environment is observable by the agent due to first-person perspective.

In Doom MyWayHome, the agent is randomly placed in one of several rooms connected in a maze-like arrangement, and the agent must reach an item that has a fixed visual appearance and is in a fixed location before time runs out. For Minecraft, we adopt the teacher-student curriculum learning task of Matiisen et al. (2017), consisting of 5 consecutive “levels” that successively increase the difficulty of completing the simple task of reaching a gold block: the first level (“L1”) consists of a single room; the intermediate levels (“L2”–“L4”) consist of a corridor with lava-bridge and wall-gap obstacles; and the final level (“L5”) consists of a 2×2 arrangement of rooms randomly separated by lava-bridge or wall-gap obstacles. Examples of the MyWayHome and the Minecraft levels are shown in Figure 1.

Our results on Doom and Minecraft are in Figures 2 and 3. Unlike previous evaluations which augmented raw pixel observations with extra information about the game state, e.g. elapsed time ticks or remaining health (Kempka et al., 2016; Dosovitskiy & Koltun, 2017), in our evaluation we forced all networks to learn using only visual input. Despite this restriction, ARM is still able to quickly learn policies with minimal tuning of hyperparameters and to reach close to the maximum achievable score in under 1 million simulator steps, which is quite sample efficient. On MyWayHome, we observed that ARM generally learned a well-performing policy more quickly than other methods. Additionally, we found that ARM is able to take advantage of an off-policy replay memory when learning on MyWayHome by storing the trajectories of previous sampling batches and applying an importance sampling correction to the n -step return estimator; please see Section A2 in the Supplementary Material.

We performed our Minecraft experiments using fixed curriculum learning schedules to evaluate the sample efficiency of different algorithms: the agent is initially placed in the first level (“L1”), and the agent is advanced to the next level whenever a preselected number of simulator steps have elapsed, until the agent reaches the last level (“L5”). We found that ARM and dueling double DQN both were able to learn on an aggressive “fast” schedule of only 62500 simulator steps between levels. TRPO required a “slow” schedule of 93750 simulator steps between levels to reliably learn. ARM was able to consistently learn a well performing policy on all of the levels, whereas double DQN learned more slowly on some of the intermediate levels. ARM also more consistently reached a high score on the final, most difficult level (“L5”).



Figure 1. Screenshots from (left) Doom MyWayHome, (middle) Minecraft level 1, and (right) Minecraft level 4.




Figure 2. Evaluation on the Doom MyWayHome task.

4.2. Learning with Partially Observed Objects

Previously in Section 3.5, we argued that the convergence results for CFR suggest that reducing the size of the observation space could improve the convergence of CFR compared to methods based on policy gradients or Q-learning. We would expect that by controlling the degree of partial observability in a fixed task, one could expect the relative sample efficiency or performance of ARM to improve compared to other methods.

Whereas the navigation tasks in Section 4.1 only dealt with first-person motion in static 3D environments, tasks that add other objects which may move and are not always visible to the agent intuitively ought to be less observable. To test the effect of the degree of partial observability on fixed tasks with objects, we conduct experiments with at least two variants of each task: one experiment is based on the unmodified task, but the other experiments occlude or mask essential pixels in observed frames to raise the difficulty of identifying relevant objects.

We evaluate ARM and other methods on two tasks with objects: Doom Corridor via ViZDoom, and Atari Pong via the Arcade Learning Environment (Bellemare et al., 2013). In Corridor, the agent is placed in a narrow 3D hallway and must avoid getting killed by any of several shooting monsters scattered along the route while trying to reach the end of the hallway. In Pong, the agent controls a paddle and must hit a moving ball past the computer opponent’s paddle on the other side of the screen.

In our implementation of Corridor, which we call “Corridor+,” we restrict the action space of the agent to only be able to turn left or right and to move to one side. Because




Figure 3. Evaluation on a Minecraft curriculum learning task. The simulator step counts at which each level begins are labeled and demarcated with dashed vertical lines.

the agent spawns facing the end of the hallway, the agent has to learn to first turn to its side to orient itself sideways before moving, while also trying to avoid getting hit from behind by monsters in close proximity. To induce partial observability in Corridor, we mask out a large square region in the center of the frame buffer. An example of this occlusion is shown in Figure 4. Because the only way for the agent to progress toward the end of the hallway is by moving sideways, the monsters that are closest to the agent, and hence the most dangerous, will appear near the center of the agent’s field-of-view. So the center pixels are the most natural ones to occlude in Corridor+.

For Pong, we choose a rectangular area in the middle of the frame between the two players’ paddles, then we set all pixels in the rectangular area to the background color. An illustration of the occlusion is shown in Figure 4. When the ball enters the occluded region, it completely disappears from view so its state cannot be reconstructed using a limited frame history. Intuitively, the agent needs to learn to anticipate the trajectory of the ball in the absence of visual cues of the ball’s position and velocity.

It is also known that for Atari games in general, with only 4 observed frames as input, it is possible to predict hundreds of frames into the future on some games using only a feed-forward dynamics model (Oh et al., 2015). By default, all of our networks receive as input a frame history of length 4. This suggests that limiting the frame history length in Pong is another effective perturbation for reducing observability. When the frame history length is limited to one, then reconstructing the velocity of moving objects in Pong becomes more difficult (Hausknecht & Stone, 2017).

Our results on Corridor+ are shown in Figure 5. Among all our experiments, it is on Corridor+ that TRPO performs the best compared to all other methods. One distinguishing feature of TRPO is its empirical full return estimator, which is unbiased on non-Markovian observations, but which can have greater variance than the n -step return estimator used in deep Q-learning and ARM. On Corridor+, there appears to be a benefit to using the unbiased full returns over the

biased n -step returns, which speaks to the non-Markovian character of the Corridor+ task. It is also evident on Corridor+ that the performance of deep Q-learning suffers when the observability is reduced by occlusion, and that by comparison ARM is relatively unaffected, despite both methods using biased n -step returns with the same value of n ($n = 5$). This suggests that even when ARM is handicapped by the bias of its return estimator, ARM is intrinsically more robust to the non-Markovian observations that arise from partial observability. One possible direction to improve ARM is through alternative return estimators (Schulman et al., 2016; Wang et al., 2017).

Our results on Pong are shown in Figure 6. The convergence of ARM on all three variants of Pong suggests that ARM is not affected much by partial observation in this domain. As expected, when observability is reduced in the limited frame history and occlusion experiments, deep Q-learning performs worse and converges more slowly than ARM. TRPO is generally less sample efficient than either ARM or DQN, although like ARM it seems to be resilient to the partial observation perturbations.



Figure 4. Screenshots from Atari Pong (left two) and Doom Corridor+ (right two), unmodified and with occlusion.




Figure 5. Results on Doom Corridor+: (left) unmodified, (right) with occlusion.




Figure 6. Results on Atari Pong: (left) unmodified, (middle) with limited frame history, (right) with occlusion.

5. Discussion

In this paper, we presented a novel deep reinforcement learning algorithm based on counterfactual regret minimization (CFR). We call our method advantage-based regret minimization (ARM). Similarly to prior methods that learn state or state-action value functions, our method learns a cumulative clipped advantage function of observation and action. However, in contrast to these prior methods, ARM is well suited to partially observed or non-Markovian environments, making it an appealing choice in a number of difficult domains. When compared to baseline methods, including deep Q-learning and TRPO, on partially observable tasks such as first-person navigation in Doom and Minecraft and interacting with partially observed objects in Doom and Pong, ARM is robust to the degree of partial observability and can achieve substantially better sample efficiency and performance. This illustrates the value of ARM for partially observable problems. In future work, we plan to further explore applications of ARM to more complex tasks, including continuous action spaces.

Acknowledgements

We thank Tabet Matiisen for providing the Malmö curriculum learning tasks, Deepak Pathak for help with ViZDoom experiments, and Tuomas Haarnoja, Richard Li, and Haoran Tang for helpful discussions. This work was partially supported by Berkeley DeepDrive, ADEPT Lab industrial sponsors and affiliates Intel, Google, Siemens, and SK Hynix, DARPA Award Number HR0011-12-2-0016, and ASPIRE Lab industrial sponsors and affiliates Intel, Google, HPE, Huawei, LGE, Nokia, NVIDIA, Oracle, and Samsung.

References

- Ansel, O., Baram, N., and Shimkin, N. Averaged-DQN: Variance Reduction and Stabilization for Deep Reinforcement Learning. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 176–185, 2017.
- Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. The Arcade Learning Environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- Bellemare, M. G., Ostrovski, G., Guez, A., Thomas, P. S., and Munos, R. Increasing the Action Gap: New Operators for Reinforcement Learning. In *AAAI Conference on Artificial Intelligence*, 2016.
- Bellemare, M. G., Dabney, W., and Munos, R. A Distributional Perspective on Reinforcement Learning. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.

- Bowling, M., Burch, N., Johanson, M., and Tammelin, O. Heads-up limit hold'em poker is solved. *Science*, 347(6218):145–149, 2015.
- Cesa-Bianchi, N. and Lugosi, G. Potential-Based Algorithms in On-Line Prediction and Game Theory. *Machine Learning*, 51(3):239–261, 2003.
- Dick, T. Policy Gradient Reinforcement Learning Without Regret. Master's thesis, University of Alberta, 2015.
- Dosovitskiy, A. and Koltun, V. Learning to Act by Predicting the Future. In *International Conference on Learning Representations*, 2017.
- Gordon, G. J. No-regret Algorithms for Online Convex Programs. In *Advances in Neural Information Processing Systems*, 2007.
- Gu, S., Lillicrap, T., Ghahramani, Z., Turner, R. E., Schölkopf, B., and Levine, S. Interpolated Policy Gradient: Merging On-Policy and Off-Policy Gradient Estimation for Deep Reinforcement Learning. *arXiv preprint arXiv:1706.00387*, 2017.
- Haarnoja, T., Tang, H., Abbeel, P., and Levine, S. Reinforcement Learning with Deep Energy-Based Policies. In *International Conference on Machine Learning*, 2017.
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- Hart, S. and Mas-Colell, A. A Simple Adaptive Procedure Leading to Correlated Equilibrium. *Econometrica*, 68(5): 1127–1150, 2000.
- Hausknecht, M. and Stone, P. Deep Recurrent Q-Learning for Partially Observable MDPs. *arXiv preprint arXiv:1507.06527*, 2017.
- Heess, N., Hunt, J. J., Lillicrap, T. P., and Silver, D. Memory-based control with recurrent neural networks. *arXiv preprint arXiv:1512.04455*, 2015.
- Jaakkola, T., Singh, S. P., and Jordan, M. I. Reinforcement Learning Algorithm for Partially Observable Markov Decision Problems. In *Advances in Neural Information Processing Systems*, 1994.
- Johnson, M., Hofmann, K., Hutton, T., and Bignell, D. The Malmo Platform for Artificial Intelligence Experimentation. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, 2016.
- Kempka, M., Wydmuch, M., Runc, G., Toczek, J., and Jaśkowski, W. ViZDoom: A Doom-based AI Research Platform for Visual Reinforcement Learning. *arXiv preprint arXiv:1605.02097*, 2016.
- Kingma, D. P. and Ba, J. L. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*, 2015.
- Kroer, C. and Sandholm, T. Extensive-Form Game Abstraction With Bounds. In *Proceedings of the 15th ACM Conference on Economics and Computation*, 2014.
- Levine, S., Finn, C., Darrell, T., and Abbeel, P. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1):1334–1373, 2016.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. In *International Conference on Learning Representations*, 2016.
- Littman, M. L. Markov games as a framework for multi-agent reinforcement learning. In *Proceedings of the 11th International Conference on Machine Learning*, pp. 157–163, 1994.
- Matiisen, T., Oliver, A., Cohen, T., and Schulman, J. Teacher-Student Curriculum Learning. *arXiv preprint arXiv:1707.00183*, 2017.
- McCallum, A. K. Efficient Exploration in Reinforcement Learning with Hidden State. In *AAAI Fall Symposium on Model-directed Autonomous Systems*, 1997.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing Atari with Deep Reinforcement Learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529–533, 2015.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pp. 1928–1937, 2016.
- Nachum, O., Norouzi, M., Xu, K., and Schuurmans, D. Bridging the Gap Between Value and Policy Based Reinforcement Learning. In *Advances in Neural Information Processing Systems*, 2017.

- O'Donoghue, B., Munos, R., Kavukcuoglu, K., and Mnih, V. Combining policy gradient and Q-learning. *arXiv preprint arXiv:1611.01626*, 2017.
- Oh, J., Guo, X., Lee, H., Lewis, R. L., and Singh, S. Action-Conditional Video Prediction using Deep Networks in Atari Games. In *Advances in Neural Information Processing Systems*, 2015.
- Oh, J., Chockalingam, V., Singh, S., and Lee, H. Control of memory, active perception, and action in minecraft. In *Proceedings of The 33rd International Conference on Machine Learning*, pp. 2790–2799, 2016.
- Peng, X. B., Berseth, G., and van de Penne, M. Terrain-Adaptive Locomotion Skills Using Deep Reinforcement Learning. *ACM Transactions on Graphics*, 35(4):81, 2016.
- Ross, S. and Bagnell, J. A. Reinforcement and Imitation Learning via Interactive No-Regret Learning. *arXiv preprint arXiv:1406.5979*, 2014.
- Ross, S., Gordon, G. J., and Bagnell, J. A. A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pp. 627–635, 2011.
- Sandholm, T. and Singh, S. Lossy Stochastic Game Abstraction with Bounds. In *Proceedings of the 13th ACM Conference on Economics and Computation*, 2012.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pp. 1889–1897, 2015.
- Schulman, J., Moritz, P., Levine, S., Jordan, M. I., and Abbeel, P. High-Dimensional Continuous Control Using Generalized Advantage Estimation. *arXiv preprint arXiv:1506.02438*, 2016.
- Schulman, J., Chen, X., and Abbeel, P. Equivalence Between Policy Gradients and Soft Q-Learning. *arXiv preprint arXiv:1704.06440*, 2017.
- Szepesvári, C. The Asymptotic Convergence-Rate of Q-learning. In *Advances in Neural Information Processing Systems*, 1998.
- Tammelin, O. Solving Large Imperfect Information Games Using CFR+. *arXiv preprint arXiv:1407.5042*, 2014.
- Tammelin, O., Burch, N., Johanson, M., and Bowling, M. Solving Heads-up Limit Texas Hold'em. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, 2015.
- van Hasselt, H. and Wiering, M. A. Reinforcement Learning in Continuous Action Spaces. In *Proceedings of the 2007 IEEE Symposium on Approximate Dynamic Programming and Reinforcement Learning*, pp. 272–279, 2007.
- van Hasselt, H., Guez, A., and Silver, D. Deep Reinforcement Learning and Double Q-Learning. In *AAAI Conference on Artificial Intelligence*, 2016.
- Wang, Z., Schaul, T., Hessel, M., van Hasselt, H., Lanctot, M., and de Freitas, N. Dueling Network Architectures for Deep Reinforcement Learning. In *Proceedings of the 33rd International Conference on Machine Learning*, pp. 1995–2003, 2016.
- Wang, Z., Bapst, V., Heess, N., Mnih, V., Munos, R., Kavukcuoglu, K., and de Freitas, N. Sample Efficient Actor-Critic with Experience Replay. *arXiv preprint arXiv:1611.01224*, 2017.
- Waugh, K., Morrill, D., Bagnell, J. A., and Bowling, M. Solving Games with Functional Regret Estimation. In *Workshops at the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015. Supplementary material in *arXiv preprint arXiv:1411.7974*.
- Zinkevich, M., Johanson, M., Bowling, M. H., and Piccione, C. Regret Minimization in Games with Incomplete Information. In *Advances in Neural Information Processing Systems*, 2007.