

A Unified View of Large-Scale Zero-Sum Equilibrium Computation

Kevin Waugh

waugh@cs.cmu.edu
School of Computer Science
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA 15213 USA

J. Andrew Bagnell

dbagnell@ri.cmu.edu
The Robotics Institute
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA 15213 USA

Abstract

The task of computing approximate Nash equilibria in large zero-sum extensive-form games has received a tremendous amount of attention due mainly to the Annual Computer Poker Competition. Immediately after its inception, two competing and seemingly different approaches emerged—one an application of no-regret online learning, the other a sophisticated gradient method applied to a convex-concave saddle-point formulation. Since then, both approaches have grown in relative isolation with advancements on one side not affecting the other. In this paper, we rectify this by dissecting and, in a sense, unify the two views.

Introduction

The first annual Computer Poker Competition was held in 2007 providing a testbed for adversarial decision-making with imperfect information. Though incredible advancements have been made since then, the solution of an abstract game still remains a critical component of the top agents. The strength of such a strategy correlates with how well the abstract game models the intractably large full game (Johanson et al. 2011). Algorithmic improvements in equilibrium computation enable teams to solve larger abstract games and thus improve the overall strength of their agents.

The first agents used off-the-shelf linear programming packages to solve the abstract games. Shortly after, two specialized equilibrium-finding techniques emerged drastically reducing resource requirements by allowing implicit game representations. One method, counterfactual regret minimization (CFR), combines many simple no-regret learners together to minimize overall regret and as a whole they converge to an equilibrium (Zinkevich et al. 2008). The other, an application of Nesterov’s excessive gap technique (EGT, 2005), is a gradient method that attacks the convex-concave saddle-point formulation directly (Gilpin et al. 2007).

Currently, the two techniques are thought of as simply different and competing. Though both improved, the advancements and the methods remained isolated. At this point, CFR has more widespread adoption due to its simplicity and the power of the sampling schemes available to it.

In this paper we connect CFR and EGT. Specifically, both define Bregman divergences with the same structure. This viewpoint allows us to make important connections between the two as well as current research on convex optimization and no-regret learning. We show that CFR can be thought of as smoothed fictitious play and its dual weights are a function of the opponent’s average strategy; and that a convergence rate of $O(1/T)$ can be achieved while sampling.

Zero-Sum Extensive-form Games

A **extensive-form game** is a tuple $\Gamma = (N, \mathcal{H}, p, \sigma_c, \mathcal{I}, u)$ (see, e.g., Osborne and Rubinstein). The game has N players, the set of which is $[N]$. The set of **histories**, \mathcal{H} , form a tree rooted at ϕ , the empty history. For a history $h \in \mathcal{H}$, we denote the set of available **actions** by $A(h)$. For any action $a \in A(h)$, the history $ha \in \mathcal{H}$ is a child of h . The tree’s leaves, or **terminal histories**, are denoted by $\mathcal{Z} \subseteq \mathcal{H}$. At a **non-terminal** history the **player choice function**, $p : \mathcal{H} \setminus \mathcal{Z} \rightarrow [N] \cup \{c\}$, determines who is to act, either a player or nature. Nature’s policy is denoted σ_c defines a distribution over actions when it is to act $\sigma_c(\cdot|h) \in \Delta_{A(h)}$. The **information partition**, $\mathcal{I} = \cup_{i \in [N]} \mathcal{I}_i$, is a partition the players’ histories. All histories in an **information set** are indistinguishable to the player to act. We have $\forall I \in \mathcal{I}, h, h' \in I$, that $p(h) = p(h')$ and $A(h) = A(h')$. Finally, at a terminal history, $z \in \mathcal{Z}$, the **utility for player i** , $u_i : \mathcal{Z} \rightarrow \mathbb{R}$, determines the reward for player i reaching terminal z .

Each player plays the game by means of a **behavioral strategy**. A behavioral strategy for player i , $\sigma_i \in \Sigma_i$, maps histories to distributions over actions, $\sigma_i(\cdot|h) \in \Delta_{A(h)}$. When the player is to act, their choice is drawn from this distribution. A strategy must respect the information partition, $\forall I \in \mathcal{I}_i, h, h' \in I, \sigma_i(\cdot|h) = \sigma_i(\cdot|h')$. We call the tuple of N strategies, $(\sigma_1, \dots, \sigma_N)$, a **strategy profile**.

There are two additional properties we require of the extensive-form games we examine. First, we consider two-player **zero-sum** games. That is, $N = 2$ and $u_2(z) = -u_1(z)$; what one player wins the other loses.

Second, we will consider games of **perfect recall**—neither player is forced forget any information they once knew. Mathematically this requirement is that all histories in an information set share the same sequence of information sets and actions from the point-of-view of the acting player.

With these additional restrictions, we can conveniently

represent a game in its **sequence form**, $\Gamma = (A, E, F)$. Following from perfect recall, any sequence of information set/action pairs, henceforth simply sequence, is uniquely identified by its final element. Consequently, we can represent behavioral strategies as vectors, called **realization plans**, such that the expected utility of the game is a bilinear product $x^T Ay$. In particular, a realization plan for the row player is a non-negative vector x indexed by sequences such that $\sum_{a \in A(I)} x(I, a) = x(\text{parent}(I))$ for all $I \in \mathcal{I}_i$ and where $x(\phi) = 1$. In words, the probability mass flowing out of an information set equals the probability of playing to reach that information set. The constraint on the empty sequence, ϕ , normalizes the probability. We represent these linear equality constraints with the matrix E , i.e., $Ex = e_1$, and thus $\Sigma_1 = \{x \mid Ex = e_1, x \geq 0\}$. For the column player, we have corresponding concepts y , F , and Σ_2 . The matrix A is the **payoff matrix**. Entry $a_{i,j}$ is the expected payoff to the row player over all terminals reached by sequences i and j .

A pair of strategies, (x, y) is said to be an ε -**Nash equilibrium** if neither player can benefit by more than ε by deviating to another strategy. In particular,

$$\begin{aligned} x^T Ay + \varepsilon &\geq x' Ay, & \forall x' \in \Sigma_1 \\ -x^T Ay + \varepsilon &\geq -x^T Ay', & \forall y' \in \Sigma_2 \end{aligned}$$

Remarkably, a Nash equilibrium always exists and in our case we can efficiently find ε -equilibria. In the next sections we will discuss and relate large-scale methods to do so.

Counterfactual Regret Minimization

Online learning is a powerful framework for analyzing the performance of adaptive algorithms. At time $t \in [T]$, an online algorithm chooses a policy $x^t \in \Sigma$ and then receives reward vector $u^t \in \mathcal{K}$. It aims to minimize **external regret**,

$$\max_{x^* \in \Sigma} \sum_{t=1}^T u^t \cdot x^* - u^t \cdot x^t,$$

its utility relative to the best fixed policy in hindsight. An algorithm is said to be **no-regret** if its regret grows sublinear in T for any sequence of u 's from bounded set \mathcal{K} . That is, if the bound on its time-averaged regret approaches zero in the worst-case (Cesa-Bianchi and Lugosi 2006).

There is an important connection between no-regret learning and zero-sum equilibrium computation. Two no-regret algorithms in self-play converge to a Nash equilibrium. Operationally, the row player gets reward $u^t = Ay^t$ and the column player $u^t = -A^T x^t$.

Theorem 1. *If two no-regret algorithms play a zero-sum game against one and other for T iterations and have time-averaged regret less than ε , then their average strategies (\bar{x}, \bar{y}) are a 2ε -Nash equilibrium. Here $\bar{x} = \sum_{t=1}^T x^t / T$.*

We provide a proof for completeness.

Proof. For any $x' \in \Sigma_1$ and $y' \in \Sigma_2$,

$$\frac{1}{T} \sum_{t=1}^T x' \cdot Ay^t - x^t \cdot Ay^t \leq \varepsilon, \text{ and}$$

$$\frac{1}{T} \sum_{t=1}^T (-x^t \cdot Ay') - (-x^t \cdot Ay^t) \leq \varepsilon$$

adding the two inequalities together

$$\frac{1}{T} \sum_{t=1}^T x' \cdot Ay^t - x^t \cdot Ay' \leq 2\varepsilon$$

substituting in the definitions of \bar{x} and \bar{y}

$$x' \cdot A\bar{y} - \bar{x} \cdot Ay' \leq 2\varepsilon$$

choosing $y' = \bar{y}$, we get for all $x' \in \Sigma_1$

$$x' \cdot A\bar{y} \leq \bar{x} \cdot A\bar{y} + 2\varepsilon$$

Similarly, if instead we choose $x' = \bar{x}$, we get the second inequality in the definition of a 2ε -Nash. \square

The space of mixed strategies, Σ_1 and Σ_2 , though structured, is complicated. Zinkevich et al. overcome this and describe a no-regret algorithm over the space of realization plans. Their approach minimizes a new notion of regret—counterfactual regret—at each information set using simple no-regret algorithms for the probability simplex. They show that counterfactual regret bounds external regret, thus their approach computes an equilibrium in self-play.

The **counterfactual utility** for action a at information set I is the expected utility given that the player tries to and successfully plays action a . That is, we weight a terminal by the total probability of the opponent and chance reaching it, but only by the remaining probability for the acting player. Computing counterfactual utilities is done by traversing the game tree, or one sparse matrix-vector product.

There are numerous no-regret algorithms operating over the probability simplex, $\Sigma = \Delta$. Let us present two. The first, regret-matching (Hart and Mas-Colell 2000), is also known as the polynomially-weighted forecaster. It is the most commonly used algorithm for zero-sum equilibrium computation. Notationally, we call $r^t = u^t - u^t \cdot x^t e$ the **instantaneous regret** and $R^t = \sum_{i=1}^t r^i$ the **cumulative regret** at time t . Let $L \geq \sup_{u \in \mathcal{K}} \|u\|_\infty^2$.

Definition 1 (Regret-matching). *Choose $x^{t+1} \propto (R^t)_+$.*

Here, $(x)_+ = \max\{0, x\}$.

Theorem 2 (from (Hart and Mas-Colell 2000)). *If x^t 's are chosen using regret-matching, then the external regret is no more than \sqrt{LNT} .*

The second algorithm is Hedge (Freund and Schapire 1997). It is also known as the exponentially-weighted forecaster, exponential weights, or weighted majority (Littlestone and Warmuth 1994).

Definition 2 (Hedge). *Choose $x^{t+1} \propto \exp(\eta R^t)$.*

Theorem 3 (from (Freund and Schapire 1997)). *If x^t 's are chosen using Hedge with $\eta = \sqrt{2 \log(N)/T}/L$, then the external regret is no more than $\sqrt{2LT \log N}$.*

Algorithm 1 CFR Update with Hedge

```
function UPDATEREGRETI(R, g)
  for  $a \in A(I)$  do
    for  $I' \in \text{child}(I, a)$  do
       $u', R_{I'} \leftarrow \text{UPDATEREGRET}_{I'}(R_{I'}, g_{I'})$ 
       $g_a \leftarrow g_a + u'$ 
    end for
  end for
   $x_a \propto \exp(R_{I,a})$ 
  for  $a \in A(I)$  do
     $R_{I,a} \leftarrow R_{I,a} + g_a - g \cdot x$ 
  end for
  return  $g \cdot x, R$ 
end function
function UPDATEREGRET(R, g)
  for  $I \in \text{child}(\Phi)$  do
     $\rightarrow, R_I \leftarrow \text{UPDATEREGRET}_{I'}(R_I, g_I)$ 
  end for
end function
```

For equilibrium-finding, the regret-matching algorithm of Hart and Mas-Colell is common place. Though Hedge has a slightly better theoretical bound than regret-matching, it is computationally more expensive due to \exp and choosing η can be tricky in practice. As we will see, the use of Hedge here leads to an insightful connection between the two equilibrium-finding approaches in question. We show the counterfactual regret update in Algorithm 1.

Connection to Convex Optimization

A different view from the no-regret learning approach is simply to write the equilibrium computation as a non-smooth minimization. This minimization is convex, as both the objective and the set of realization plans are.

Theorem 4. *Let $f(x) = \max_{y \in \Sigma_2} -x^T A y$ then f is convex on Σ_1 . Furthermore, if $x^* \in \operatorname{argmin}_{x \in \Sigma_1} f(x)$ then x^* is a minimax optimal strategy for the row player.*

Proof. First, let $y' \in \operatorname{argmax}_{y \in \Sigma_2} -x^T A y$. Claim $f'(x) = -A y' \in \partial f(x)$. For any $x' \in \Sigma_1$,

$$\begin{aligned} f(x') - f(x) &= \max_{y \in \Sigma_2} -x' \cdot A y - \max_{y \in \Sigma_2} -x \cdot A y \\ &\geq -x' \cdot A y' - \max_{y \in \Sigma_2} -x \cdot A y \\ &= -x' \cdot A y' + x \cdot A y' \\ &= (-A y') \cdot (x' - x) \end{aligned}$$

This is an instantiation of Danskin's theorem (Bertsekas 1999). By the optimality of x^* , for any $\hat{x} \in \Sigma_1$:

$$x^* \cdot A y^* = -f(x^*) \geq -f(x') = \hat{x} \cdot A \hat{y}.$$

where y^* maximizes $-x^* \cdot A y$ and \hat{y} maximizes $-\hat{x} \cdot A y$. \square

Note, the subgradient computation is precisely y 's best response. The CFR-BR technique is exactly CFR applied to this non-smooth optimization (Johanson et al. 2012).

As f is convex, and we can efficiently evaluate its subgradient via a best response calculation, we can use any subgradient method to find an ε -equilibrium strategy. Unfortunately, the most basic approach, the projected subgradient method, requires a complicated and costly projection onto the set of realization plans. We can avoid this projection by employing the proper Bregman divergence. In particular, if $h : \mathcal{D} \rightarrow \mathbb{R}$ is a strongly convex such that we can quickly solve the minimization

$$\operatorname{argmin}_{x \in \mathcal{D}} g \cdot x + h(x) \quad (1)$$

we say h fits \mathcal{D} . In these cases, we can often use h in place of the squared l_2 distance and avoid any projections.

Hoda, Gilpin, and Peña (2008) describe a family of divergences, or distance generating functions, for the set of realization plans. They construct their family of distance generating functions inductively. One such h is as follows:

$$\begin{aligned} h_I(x, y) &= \sum_{a \in A(I)} x_{I,a} \log x_{I,a} + \quad \forall I \in \mathcal{I}_i \\ &\quad \sum_{I' \in \text{child}(I, a)} x_{I,a} [h_{I'}(y_{I'}/x_{I,a})] \\ h(x) &= \sum_{I \in \text{child}(\phi)} h_I(x_I) \end{aligned}$$

A few things worth noting. First, by $\text{child}(x)$ we mean the set of information sets that are immediate children of sequence x . Second, we slightly abuse of notation above in that $h_I(x, y)$ depends on the immediate realization weights, x , and all child weights y . We denote the child weights belonging to information set I' as $y_{I'}$. Second, due to perfect recall, this recursion does bottom out; there are information sets with no children and there are no cycles. At a terminal information set, h_I is the negative entropy function. The recursion makes use of the dilation or perspective operator.

Gordon (2006) introduces the same function in his supplementary materials, but does not provide a closed-form solution to its minimization. The closed-form solution is shown in Algorithm 2. Note that this algorithm is the same as a best response calculation where we replace the max operator with the softmax operator. The normalization step afterwards restores the sequence form constraints, *i.e.*, converts from a behavioral strategy to a realization plan.

The computational structure of Algorithm 1 and 2 are identical. CFR, too, must normalize when computing an equilibrium to properly average the strategies. Both use the softmax operator to on the expected utility to define the current policy. In particular, if $R = 0$, then the computation is equivalent with the exception of the initial sign of g .

Dual Averaging

Nesterov's dual averaging (2009) is a subgradient method, and thus we can use to find equilibria. As we shall see, it has close connections to counterfactual regret when equipped with Hoda, Gilpin, and Peña distance function.

Dual averaging defines two sequences, x^t , the query points, and g^t , the corresponding sequence of subgradients. The averages \bar{x} and \bar{g} converge to a primal-dual solution.

Algorithm 2 Smoothed Best Response

```
function MINIMIZE  $h_I(u)$ 
  for  $a \in A(I)$  do
    for  $I' \in \text{child}(I, a)$  do
       $u', x_{I'} \leftarrow \text{MINIMIZE } h_{I'}(u_{I'})$ 
       $u_{I,a} \leftarrow u_{I,a} + u'$ 
    end for
  end for
   $x_{I,a} \propto \exp(u_{I,a})$  ▷ softmax instead of max
  return  $u_I \cdot x, x$ 
end function
function NORMALIZE $_I(x, Z)$ 
  for  $a \in A(I)$  do
     $x_{I,a} \leftarrow x_{I,a}/Z$ 
    for  $I' \in \text{child}(I, a)$  do
      NORMALIZE $_{I'}(x_{I'}, x_{I,a})$ 
    end for
  end for
end function
function MINIMIZE  $h(g)$ 
  for  $I \in \text{child}(\Phi)$  do
     $-, x_I \leftarrow \text{MINIMIZE } h_{I'}(-g_I)$ 
    NORMALIZE $_I(x, 1)$ 
  end for
  return  $x$ 
end function
```

Definition 3. Let $\beta^t > 0$ be a sequence of step sizes and $h : \mathcal{D} \rightarrow \mathbb{R}$ be a strongly convex distance generating function. The sequence x^t is given by

$$\begin{aligned} x^{t+1} &= \operatorname{argmin}_{x \in \mathcal{D}} \frac{1}{t} \sum_{i=1}^t [f(x^i) + f'(x^i) \cdot (x - x^i)] + \beta^t h(x) \\ &= \operatorname{argmin}_{x \in \mathcal{D}} \bar{g}^t \cdot x + \beta^t h(x) \end{aligned} \quad (2)$$

The method converges for step sizes $O(\sqrt{t})$, or for an appropriately chosen constant step size when the number of iterations is known ahead of time.

Interestingly, Hedge and regret-matching over the simplex are operationally equivalent to dual averaging.

Theorem 5. Hedge is equivalent to Dual Averaging over the simplex with $\sigma(x) = \sum_{i=1}^n x_i \log x_i$, $\beta^t = 1/(t\eta)$ and $g^t = -u^t$.

Proof. The dual averaging minimization can be written as the convex conjugate of the negative entropy (Boyd and Vandenberghe 2004):

$$\begin{aligned} \min_{x \in \Delta} \bar{g}^t \cdot x + \beta^t \sum_{i=1}^n x_i \log x_i \\ &= -\beta^t \max_{x \in \Delta} -\bar{g}^t \cdot x / \beta^t - \sum_{i=1}^n x_i \log x_i \\ &= -\beta^t \log \sum_{i=1}^n \exp(-\hat{g}_i / \beta^t) \end{aligned}$$

The gradient of the conjugate minimizes the objective (Rockafeller 1970),

$$\begin{aligned} x^{t+1} &\propto \exp(-\bar{g}^t / \beta^t) = \exp\left(-\eta \sum_{i=1}^t g^i\right) \\ &\propto \exp\left(\eta \sum_{i=1}^t u^i - u^t \cdot x^t e\right) = \exp(\eta R^t) \end{aligned}$$

The last step follows from $\exp(x) \propto \exp(b) \exp(x) = \exp(bx)$ for any vector x and constant b . \square

In particular, note subtracting the expected utility in the regret update does not at all alter the iterates. We need only accumulate counterfactual utility when using Hedge.

Theorem 6. Regret-matching is equivalent to Dual Averaging over the simplex with $\sigma(x) = \|x_+\|_2^2/2$, $\beta^t = e^T R_+^t/t$ and $g^t = u^t \cdot x^t e - u^t$.

Proof. Consider the dual averaging minimization without the normalization constraint:

$$\begin{aligned} x^{t+1} &= \operatorname{argmin}_{x \geq 0} \bar{g}^t \cdot x + \beta^t \|x_+\|_2^2/2 \\ &= \frac{(-\bar{g}^t)_+}{\beta^t} = \frac{t(\frac{1}{t} \sum_{i=1}^t u^i - u^i \cdot x^i e)_+}{e^T R_+^t} \\ &= \frac{R_+^t}{e^T R_+^t} \end{aligned}$$

Note that by construction x^{t+1} sums to one, therefore the normalization constraint holds. In order for dual averaging to converge, we need $1/\beta^t \in O(\sqrt{t})$. This follows from the no-regret bound on regret-matching, $e^T R_+^t \leq \sqrt{LNT}$. \square

Following from Theorem 6, we see that CFR with regret-matching is dual averaging with a Hoda, Gilpin, and Peña-style divergence built on $\|x_+\|_2^2$. Note that the step sizes must be chosen appropriately to avoid projection. That is, this divergence may not be appropriate for other gradient methods that rely on more stringent step size choices.

Let us explicitly instantiate dual averaging for solving the convex-concave equilibrium saddle-point optimization.

$$\begin{aligned} x^{t+1} &= \operatorname{argmin}_{x \in \Sigma_1} \frac{1}{t} \sum_{i=1}^t -Ay^i + \beta^t h(x), \\ &= \operatorname{argmin}_{x \in \Sigma_1} -A\bar{y}^t + \beta^t h(x), \\ y^{t+1} &= \operatorname{argmin}_{y \in \Sigma_2} \frac{1}{t} \sum_{i=1}^t A^T x^i + \beta^t h(y), \\ &= \operatorname{argmin}_{y \in \Sigma_2} A^T \bar{x}^t + \beta^t h(y). \end{aligned}$$

In words, dual averaging applied to the saddle-point problem can be thought of as fictitious play with a smoothed best response as opposed to an actual best response (Brown 1951).

Dual averaging and Hedge are operationally equivalent at terminal information sets, that is, ones where all sequences

have no children. At a non-terminal information set, dual averaging responds as if its current policy is played against the opponent’s average strategy in future decisions. Counterfactual regret minimization, on the other hand, plays against the opponent’s current policy. The no-regret property of the algorithm guarantees that these two quantities remain close. In rough mathematical terms, we have $\sum_{t=1}^T x^t \cdot Ay^t \approx x^T \cdot \sum_{t=1}^T Ay^t$ within $O(\sqrt{T})$. Conceptually, we can conclude that counterfactual regret minimization, too, is roughly equivalent to smoothed fictitious play.

Operationally, the difference is counterfactual regret minimization propagates *and* accumulates the expected utilities from the child information sets in the regrets. Dual averaging, in spirit, is lazy and re-propagates these utilities on each iteration. We note that this re-propagation is not algorithmically necessary when we have sparse stochastic updates as the expected value of an information set only changes if play goes through it. That is, we can memoize and update this value in a sparse fashion.

This understanding of CFR hints at why it outperforms its bound in practice and why unprincipled speedups may indeed be reasonable. In particular, we can achieve faster rates of convergence, $O(\sqrt{L}/T)$ as opposed to $O(\sqrt{L/T})$, when minimizing smooth functions with gradient descent and the appropriate step size. Two no-regret learners in self-play are essentially smoothing the non-smooth objective for one and other. The smooth objective itself is changing from iteration to iteration, but this suggests we can choose more aggressive step sizes than necessary. Further evidence of this is that the convergence behavior for CFR-BR, a minimization of a known non-smooth objective, exhibits more volatile behavior that is closer to the true CFR regret bound.

Dual averaging with the Hoda, Gilpin, and Peña divergence is itself a no-regret learner over the set of realization plans (Xiao 2010). The regret bound itself is a function of the strong convexity parameter of the distance function. The bound on which appears to be quite loose. The above analysis suggests that it should be similar, or perhaps in some cases better, than the counterfactual regret bound on external regret. This is not shocking as counterfactual regret minimization is agnostic to no-regret algorithm in use.

Initialization with a Prior

When computing an equilibrium or in an online setting, typically the initial strategy is uniformly random. Though the initial error does drop quite rapidly, it is often the case that we have available good priors available. Particularly in an online setting, it is preferable start with a good strategy instead of essentially learning both how to play the game and how to exploit an opponent at the same time. From the optimization perspective, we now discuss sound approaches.

First, let us investigate how to initialize the dual weights—the cumulative counterfactual regrets. In dual averaging, the dual weights to x , $\bar{y} = A\bar{y}$, is the utilities to x of the opponent’s optimal strategy. If we have guesses at x^* and y^* , we can use those to initialize the dual weights. This view of the dual weights is a simplification what is being done by Brown and Sandholm (2014), where counterfactual

regret minimization is started from a known good solution. From the convex optimization view-point this is immediate.

An appealing property of this is that the representation of the opponent’s policy need not be known. That is, so long as we can play against it we can estimate the dual weights. In some cases, the variance of this estimate may be quite poor. With more structured representations, or domain knowledge it may be possible to greatly improve the variance.

It is quite common when considering convex optimization problems to *recenter* the problem. In particular, note that dual averaging starts with the policy $x_0 = \operatorname{argmin}_{x \in \mathcal{D}} h(x)$. By instead choosing $h'(x) = \nabla h(x) - \nabla h(x') \cdot x$, we shift the starting policy to x' . In the case of the negative entropy over the simplex, this is equivalent to instead using a Kullback-Leibler divergence. Note, the step size schedule is an important practical factor that needs to be considered more carefully. In particular, smaller step sizes keep iterates closer to the initial policy.

Convergence of the Current Policy

Some work has considered using CFR’s final policy as opposed to the average policy (Johanson et al. 2012). Experimentally, it is shown that the current policy works quite well in practice despite the lack of bounds on its performance.

A large assortment of recent work on stochastic convex optimization has considered this problem in depth exploring different averaging and step size selection schemes. When we view the problem from the convex optimization view-point these results transfer without modification. Using a step size decreasing like $1/\sqrt{t}$ will lead to convergence of the current iterate for non-smooth optimization (Shamir and Zhang 2013). That is, if the opponent plays a best response, like in CFR-BR, we need not store the average policy reducing the memory requirement by 50%.

Acceleration and Sampling

An important reason that no-regret algorithms have emerged as the dominant approach for large-scale equilibrium computation is their amenability to a various forms of sampling. At a high level, by introducing stochasticity we can drastically reduce the computation on each iteration by introducing different types of sparsity while only marginally increasing the number of necessary iterations.

Theorem 7 (from (Cesa-Bianchi and Lugosi 2006)). *Let the sequence x^t be chosen according to a no-regret algorithm with regret bound \sqrt{CT} and let $\tilde{x}^t \sim x^t$. For all $\delta \in (0, 1)$, with probability at least $1 - \delta$ the regret of the sequence \tilde{x}^t is bounded by $\sqrt{CT} + \sqrt{T/2 \log 1/\delta}$.*

By sampling \tilde{y} from y , now we choose $u^{t+1} = A\tilde{y}^t$. That is, \tilde{y}^t is a standard basis vector and u^{t+1} is just a column of A (Lanctot et al. 2009). This has a number of important computational consequences. First, we no longer require a matrix-vector multiplication reducing the complexity of each iteration to linear from quadratic. In addition to the asymptotic reduction in complexity, we improve the constants since selecting a column of A requires no arithmetic

operations. In fact, we may not even need to store u^{t+1} if we can index directly into A .

A second important computational improvement is we can now often use integral numbers in place of floating-point numbers (Gibson 2013). In particular, note that $r^t = u^t - u^t \cdot \tilde{x}^t e$ is integral so long as u^t is. Furthermore, \tilde{x} can be represented as a vector of counts—the number of times each action is sampled. By using regret-matching, even computing x^t can be done completely with integers as no floating-point math is required to convert R^T to x^{t+1} .

Another form of sampling is possible when nature participates in the game (Zinkevich et al. 2008). For example, imagine that $A = \sum_{i=1}^p A_i$ and we can implicitly access each A_i . This is the case when nature rolls a die or draws a card from a shuffled deck. Instead of explicitly forming and storing A , or examining each A_i on every iteration, we can choose one randomly. That is, $\tilde{A}^t = pA_{i^t}$, where $i^t \sim \text{Uniform}([p])$. When we cannot store A this form of sampling reduces each iteration’s asymptotic complexity from linear in p to constant.

Nesterov’s excessive gap technique (2005) cannot handle stochastic gradients. This is the primary reason that it and the Hoda, Gilpin, and Peña divergence fell out of favor for equilibrium computation. As we note here, the divergence itself has nothing to do with the inability to handle stochasticity. It is a power tool enabling us to consider a wide variety of gradient methods. The stochastic mirror prox (SMP) algorithm of Juditsky, Nemirovski, and Tauvel (2011) is an extension of the extra-gradient method that achieves a rate of $O(\sqrt{L}/T + \sigma/\sqrt{T})$ on saddle-point problems like ours. Here, σ is the variance of the gradient. This rate is optimal. Specifically, it enables us to trade off between the variance of the sampling technique and the slower $1/\sqrt{T}$ rate that CFR achieves. Current sampling procedures favors low computation with sparse updates. It is unlikely that using them along side SMP will work well out-of-the-box. Further investigation is needed to determine if a successful compromise can practically improve performance.

Acknowledgments

This work is supported by the Office of Naval Research MURI grant N00014-09-1-1052.

References

Bertsekas, D. 1999. *Nonlinear Programming*. Athena Scientific.

Boyd, S., and Vandenberghe, L. 2004. *Convex Optimization*. Cambridge University Press.

Brown, N., and Sandholm, T. 2014. Regret transfer and parameter optimization. In *AAAI Conference on Artificial Intelligence (AAAI)*.

Brown, G. W. 1951. Iterative solutions of games by fictitious play. *Activity Analysis of Production and Allocation*.

Cesa-Bianchi, N., and Lugosi, G. 2006. *Prediction, Learning and Games*. Cambridge University Press.

Freund, Y., and Schapire, R. E. 1997. A decision-theoretic generalization of online learning and an application to boost-

ing. *Journal of Computer and System Sciences* 55(1):119–139.

Gibson, R. 2013. *Regret Minimization in Games and the Development of Champion Multiplayer Computer Poker-Playing Agents*. Ph.D. Dissertation, University of Alberta.

Gilpin, A.; Hoda, S.; Peña, J.; and Sandholm, T. 2007. Gradient-based algorithms for finding Nash equilibria in extensive form games. In *International Workshop on Internet and Network Economics (WINE)*.

Gordon, G. J. 2006. No-regret algorithms for online convex programs. In *Neural Information Processing Systems (NIPS)*.

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

Hoda, S.; Gilpin, A.; and Peña, J. 2008. Smoothing techniques for computing Nash equilibria of sequential games. *Optimization Online*.

Johanson, M.; Waugh, K.; Bowling, M.; and Zinkevich, M. 2011. Accelerating best response calculation in large extensive games. In *International Joint Conference on Artificial Intelligence (IJCAI)*.

Johanson, M.; Bard, N.; Burch, N.; and Bowling, M. 2012. Finding optimal abstract strategies in extensive form games. In *Proceedings of the Twenty-Sixth Conference on Artificial Intelligence (AAAI)*.

Juditsky, A.; Nemirovski, A.; and Tauvel, C. 2011. Solving variational inequalities with stochastic mirror-prox algorithm. *Stochastic Systems* 1(1):17–58.

Lanctot, M.; Waugh, K.; Zinkevich, M.; and Bowling, M. 2009. Monte carlo sampling for regret minimization in extensive games. In *Advances in Neural Information Processing Systems (NIPS)*, 1078–1086.

Littlestone, N., and Warmuth, M. 1994. The weighted majority algorithm. *Information and Computation* 108:212–261.

Nesterov, Y. 2005. Excessive gap technique in nonsmooth convex minimization. *SIAM Journal on Optimization* 16(1):235–249.

Nesterov, Y. 2009. Primal-dual subgradient methods for convex problems. *Mathematical Programming* 120(1):221–259.

Osborne, M., and Rubinstein, A. 1994. *A Course On Game Theory*. MIT Press.

Rockafeller, R. T. 1970. *Convex Analysis*. Princeton University Press.

Shamir, O., and Zhang, T. 2013. Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes. In *International Conference on Machine Learning (ICML)*.

Xiao, L. 2010. Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning Research* 11:2543–2596.

Zinkevich, M.; Johanson, M.; Bowling, M.; and Piccione, C. 2008. Regret minimization in games with incomplete information. In *Advances in Neural Information Processing Systems 20 (NIPS)*, 905–912.