

Dynamic Thresholding and Pruning for Regret Minimization

Noam Brown and Christian Kroer and Tuomas Sandholm

Carnegie Mellon University, Computer Science Department,
noamb@cmu.edu, ckroer@cs.cmu.edu, sandholm@cs.cmu.edu

Abstract

Regret minimization is widely used in determining strategies for imperfect-information games and in online learning. In large games, computing the regrets associated with a single iteration can be slow. For this reason, pruning – in which parts of the decision tree are not traversed in every iteration – has emerged as an essential method for speeding up iterations in large games. The ability to prune is a primary reason why the Counterfactual Regret Minimization (CFR) algorithm using regret matching has emerged as the most popular iterative algorithm for imperfect-information games, despite its relatively poor convergence bound. In this paper, we introduce dynamic thresholding, in which a threshold is set at every iteration such that any action in the decision tree with probability below the threshold is set to zero probability. This enables pruning for the first time in a wide range of algorithms. We prove that dynamic thresholding can be applied to Hedge while increasing its convergence bound by only a constant factor in terms of number of iterations. Experiments demonstrate a substantial improvement in performance for Hedge as well as the excessive gap technique.

1 Introduction

We introduce *dynamic thresholding* for online learning algorithms, in which a threshold is set at every iteration such that any action with probability below the threshold is set to zero probability. This enables pruning for the first time in a wide range of algorithms. The theory that we derive applies to each of the two central goals in the area:

1. Regret minimization in any setting, where there can be any number of players in a general-sum game, and our agent may not even know what the game is (except that the agent knows the available actions when it is her turn to move).
2. Converging to an ϵ -Nash equilibrium in a two-player zero-sum game. Results for (1) immediately imply results for this setting by having our algorithm be used by both agents.

We will introduce this first for the application of solving zero-sum imperfect-information games (that is, games like heads-up poker), and then explain how the results directly

carry over to non-zero-sum games and to general regret minimization. Furthermore, the results apply to both extensive-form and normal-form representations.

Imperfect-information extensive-form games are a way to model strategic multi-step interactions between players that have hidden information, such as negotiations, auctions, cybersecurity settings, and medical settings. A Nash equilibrium in relatively small two-player zero-sum games containing around 10^8 nodes can be found precisely using a linear program (Gilpin and Sandholm 2007). For larger games, iterative algorithms are used to converge to an ϵ -Nash equilibrium. There are a number of such iterative algorithms, the most popular of which is *Counterfactual Regret Minimization (CFR)* (Zinkevich et al. 2007). CFR minimizes regret independently at each decision point (called an information set) in the game tree using any regret-minimizing algorithm. By far the most popular regret-minimizing algorithm to use within CFR is *regret matching (RM)* and variants of RM (Hart and Mas-Colell 2000; Gibson et al. 2012; Gibson 2014; Brown, Ganzfried, and Sandholm 2015). CFR+, a variant of CFR with RM, was recently used to essentially solve Limit Texas Hold'em, the largest imperfect-information game ever to be essentially solved and the first that is played competitively by humans (Bowling et al. 2015; Tammelin et al. 2015). That game (after lossless abstraction (Gilpin and Sandholm 2007) as a preprocessor) has over 10^{13} information sets.

When computing strategies for large imperfect-information games, repeatedly traversing the entire game tree with an iterative algorithm may be prohibitively slow. For this reason, pruning—in which parts of the game tree are not traversed on every iteration—has emerged as an essential method for dealing with large games. The ability to prune is a primary reason why the *Counterfactual Regret Minimization algorithm (CFR)* that uses *Regret Matching (RM)* at each information set is the most popular algorithm for imperfect-information games, despite its relatively poor $O(\sqrt{|A|T})$ cumulative regret.

While regret-minimizing algorithms other than RM can be used within CFR, and iterative algorithms other than CFR exist with better convergence bounds in terms of the number of iterations needed (Hoda et al. 2010; Pays 2014; Gilpin, Peña, and Sandholm 2012), CFR with RM exhibits superior empirical performance in large games (Kroer et al.

2015). A primary reason for this is that CFR with RM is able to put zero probability on some actions, and therefore prune large sections of the game tree, particularly in large games. That is, it need not traverse the entire game tree on each iteration. This behavior is shared by some other regret minimizing algorithms, but is relatively uncommon and is considered a desirable property (Luo and Schapire 2014). The ability to prune enables each iteration to be completed far more quickly. While the benefit of pruning varies depending on the game, it can easily be multiple orders of magnitude even in small games (Lanctot et al. 2009; Brown and Sandholm 2015). Moreover, the benefits of pruning typically grow with the size of the game.

In this paper we introduce *dynamic thresholding* that allows pruning to be applied in a wider range of algorithms, and applied more frequently in settings that already support pruning. We focus on *Hedge* (Freund and Schapire 1997; Littlestone and Warmuth 1994), also known as the *exponentially-weighted forecaster*, which is the most popular regret-minimizing algorithm in domains other than extensive-form game solving, on RM, and on the *Excessive Gap Technique* (EGT) (Nesterov 2005; Gilpin, Peña, and Sandholm 2012), which converges to an ϵ -Nash equilibrium in two-player zero-sum games in $O(\frac{1}{\epsilon})$, that is, in significantly fewer iterations than CFR which converges in $O(\frac{1}{\epsilon^2})$.

Dynamic thresholding sets a minimum probability threshold on each iteration, and any action with probability below that threshold is set to zero probability. We decrease the threshold over time, where the decrease is asymptotically slower than the possible decrease of action probabilities in Hedge and EGT. Thus, poor actions may eventually be played with probability below the threshold, and those paths in the game tree can then be pruned using the same methods as are used in CFR with RM (which we will describe in detail later in the paper). We prove that dynamic thresholding increases the convergence bound in Hedge and RM by only a small constant factor, where the factor depends on how aggressively the threshold is set. This holds whether Hedge and RM are used as stand-alone algorithms in any setting, or as subroutines within CFR for game-tree settings.

The remainder of this paper is structured as follows. In the next section, we cover background on imperfect-information extensive-form games, Nash equilibria, and CFR. Then, we formally introduce dynamic thresholding in CFR with Hedge/RM and prove its convergence guarantees. Then, we present experimental results that show that dynamic thresholding leads to a dramatic improvement in the performance of CFR with Hedge and of EGT. Finally, we will conclude and discuss other potential future uses of dynamic thresholding.

2 Background

In this section we present the background needed for the rest of the paper. The first subsection introduces the standard notation. The subsection after that covers CFR, explained in a more general way than usual because we consider other regret minimizing algorithms within CFR than the usual,

which is RM. Finally, the third subsection presents the pruning variants that have been introduced for CFR.

Notation

In an imperfect-information extensive-form game there is a finite set of players, \mathcal{P} . H is the set of all possible histories (nodes) in the game tree, represented as a sequence of actions, and includes the empty history. $A(h)$ is the actions available in a history and $P(h) \in \mathcal{P} \cup c$ is the player who acts at that history, where c denotes chance. Chance plays an action $a \in A(h)$ with a fixed probability $\sigma_c(h, a)$ that is known to all players. The history h' reached after an action is taken in h is a child of h , represented by $h \cdot a = h'$, while h is the parent of h' . More generally, h is an ancestor of h' (and h' is a descendant of h), represented by $h \sqsubseteq h'$, if there exists a sequence of actions from h to h' . $Z \subseteq H$ are terminal histories for which no actions are available. For each player $i \in \mathcal{P}$, there is a payoff function $u_i : Z \rightarrow \mathbb{R}$. If $P = \{1, 2\}$ and $u_1 = -u_2$, the game is two-player zero-sum.

Imperfect information is represented by *information sets* for each player $i \in \mathcal{P}$ by a partition \mathcal{I}_i of $h \in H : P(h) = i$. For any information set $I \in \mathcal{I}_i$, all histories $h, h' \in I$ are indistinguishable to player i , so $A(h) = A(h')$. $I(h)$ is the information set I where $h \in I$. $P(I)$ is the player i such that $I \in \mathcal{I}_i$. $A(I)$ is the set of actions such that for all $h \in I$, $A(I) = A(h)$. $|A_i| = \max_{I \in \mathcal{I}_i} |A(I)|$ and $|A| = \max_i |A_i|$. We define $U(I)$ to be the maximum payoff reachable from a history in I , and $L(I)$ to be the minimum. That is, $U(I) = \max_{z \in Z, h \in I: h \sqsubseteq z} u_{P(I)}(z)$ and $L(I) = \min_{z \in Z, h \in I: h \sqsubseteq z} u_{P(I)}(z)$. We define $\Delta(I) = U(I) - L(I)$ to be the range of payoffs reachable from a history in I . We similarly define $U(I, a)$, $L(I, a)$, and $\Delta(I, a)$ as the maximum, minimum, and range of payoffs (respectively) reachable from a history in I after taking action a . We define $D(I, a)$ to be the set of information sets reachable by player $P(I)$ after taking action a . Formally, $I' \in D(I, a)$ if for some history $h \in I$ and $h' \in I'$, $h \cdot a \sqsubseteq h'$ and $P(I) = P(I')$.

A strategy $\sigma_i(I)$ is a probability vector over $A(I)$ for player i in information set I . The probability of a particular action a is denoted by $\sigma_i(I, a)$. Since all histories in an information set belonging to player i are indistinguishable, the strategies in each of them must be identical. That is, for all $h \in I$, $\sigma_i(h) = \sigma_i(I)$ and $\sigma_i(h, a) = \sigma_i(I, a)$. We define σ_i to be a probability vector for player i over all available strategies Σ_i in the game. A strategy profile σ is a tuple of strategies, one for each player. $u_i(\sigma_i, \sigma_{-i})$ is the expected payoff for player i if all players play according to the strategy profile $\langle \sigma_i, \sigma_{-i} \rangle$. If a series of strategies are played over T iterations, then $\bar{\sigma}_i^T = \frac{\sum_{t \in T} \sigma_i^t}{T}$.

$\pi^\sigma(h) = \prod_{h' \rightarrow a \sqsubseteq h} \sigma_{P(h)}(h, a)$ is the joint probability of reaching h if all players play according to σ . $\pi_i^\sigma(h)$ is the contribution of player i to this probability (that is, the probability of reaching h if all players other than i , and chance, always chose actions leading to h). $\pi_{-i}^\sigma(h)$ is the contribution of all players other than i , and chance. $\pi^\sigma(h, h')$ is the probability of reaching h' given that h has been reached,

and 0 if $h \not\sqsubset h'$. In a *perfect-recall* game, $\forall h, h' \in I \in \mathcal{I}_i$, $\pi_i(h) = \pi_i(h')$. In this paper we focus on perfect-recall games. Therefore, for $i = P(I)$ we define $\pi_i(I) = \pi_i(h)$ for $h \in I$. We define the average strategy $\bar{\sigma}_i^T(I)$ for an information set I to be

$$\bar{\sigma}_i^T(I) = \frac{\sum_{t \in T} \pi_i^{\sigma_i^t} \sigma_i^t(I)}{\sum_{t \in T} \pi_i^{\sigma_i^t}(I)} \quad (1)$$

A *best response* to σ_{-i} is a strategy σ_i^* such that $u_i(\sigma_i^*, \sigma_{-i}) = \max_{\sigma_i' \in \Sigma_i} u_i(\sigma_i', \sigma_{-i})$. A *Nash equilibrium*, is a strategy profile where every player plays a best response. Formally, a Nash equilibrium is a strategy profile σ^* such that $\forall i, u_i(\sigma_i^*, \sigma_{-i}^*) = \max_{\sigma_i' \in \Sigma_i} u_i(\sigma_i', \sigma_{-i}^*)$. We define a *Nash equilibrium strategy* for player i as a strategy σ_i that is part of any Nash equilibrium. In two-player zero-sum games, if σ_i and σ_{-i} are both Nash equilibrium strategies, then (σ_i, σ_{-i}) is a Nash equilibrium. We define an ϵ -*equilibrium* as a strategy profile σ^* such that $\forall i, u_i(\sigma_i^*, \sigma_{-i}^*) + \epsilon \geq \max_{\sigma_i' \in \Sigma_i} u_i(\sigma_i', \sigma_{-i}^*)$.

Counterfactual Regret Minimization (CFR)

Counterfactual Regret Minimization (CFR) is the most popular algorithm for extensive-form imperfect-information games. In CFR, the strategy vector for each information set is determined according to a regret-minimization algorithm (Zinkevich et al. 2007). Typically, *regret matching (RM)* is used as the regret-minimization algorithm in CFR even though Hedge has a better convergence bound (in terms of the number of iterations) (Cesa-Bianchi and Lugosi 2006). One reason is that the vanilla version of Hedge does not support pruning of any paths in extensive-form games because all probabilities in Hedge are strictly positive. In section 3 we introduce a modification to Hedge that allows pruning, so we cover both Hedge and RM in this section.

Our analysis of CFR makes frequent use of *counterfactual value*. Informally, this is the expected utility of an information set given that player i tries to reach it. For player i at information set I given a strategy profile σ , this is defined as

$$v^\sigma(I) = \sum_{h \in I} \left(\pi_{-i}^\sigma(h) \sum_{z \in Z} (\pi^\sigma(h, z) u_i(z)) \right) \quad (2)$$

The counterfactual value of an action a is

$$v^\sigma(I, a) = \sum_{h \in I} \left(\pi_{-i}^\sigma(h) \sum_{z \in Z} (\pi^\sigma(h \cdot a, z) u_i(z)) \right) \quad (3)$$

Let σ^t be the strategy profile used on iteration t . The *instantaneous regret* on iteration t for action a in information set I is

$$r^t(I, a) = v^{\sigma^t}(I, a) - v^{\sigma^t}(I) \quad (4)$$

and the *regret* for action a in I on iteration T is

$$R^T(I, a) = \sum_{t \in T} r^t(I, a) \quad (5)$$

Additionally, $R_+^T(I, a) = \max\{R^T(I, a), 0\}$ and $R^T(I) = \max_a \{R_+^T(I, a)\}$. Regret for player i in the entire game is

$$R_i^T = \max_{\sigma_i' \in \Sigma_i} \sum_{t \in T} (u_i(\sigma_i', \sigma_{-i}^t) - u_i(\sigma_i^t, \sigma_{-i}^t)) \quad (6)$$

In regret matching, a player picks a distribution over actions in an information set in proportion to the positive regret on those actions. Formally, on each iteration $T + 1$, player i selects actions $a \in A(I)$ according to probabilities

$$\sigma^{T+1}(I, a) = \begin{cases} \frac{R_+^T(I, a)}{\sum_{a' \in A(I)} R_+^T(I, a')}, & \text{if } \sum_{a'} R_+^T(I, a') > 0 \\ \frac{1}{|A(I)|}, & \text{otherwise} \end{cases} \quad (7)$$

If a player plays according to regret matching in information set I on every iteration then on iteration T , $R^T(I) \leq \Delta(I) \sqrt{|A(I)| \sqrt{T}}$ (Cesa-Bianchi and Lugosi 2006).

In Hedge, a player picks a distribution over actions in an information set according to

$$\sigma^{T+1}(I, a) = \frac{e^{\eta_T R^T(I, a)}}{\sum_{a' \in A(I)} e^{\eta_T R^T(I, a')}} \quad (8)$$

where η_T is a tuning parameter. There is a substantial literature on how to set η_T for best performance (Cesa-Bianchi and Lugosi 2006; Cesa-Bianchi, Mansour, and Stoltz 2007). If a player plays according to Hedge in information set I on every iteration t and uses $\eta_t = \sqrt{\frac{2 \ln(|A(I)|)}{T}}$ then on iteration T , $R^T(I) \leq \Delta(I) \sqrt{2 \ln(|A(I)|) T}$ (Cesa-Bianchi and Lugosi 2006).

If a player plays according to CFR on every iteration then

$$R_i^T \leq \sum_{I \in \mathcal{I}_i} R^T(I) \quad (9)$$

So, as $T \rightarrow \infty$, $\frac{R_i^T}{T} \rightarrow 0$.

In two-player zero-sum games, if both players' average regret satisfies $\frac{R_i^T}{T} \leq \epsilon$, their average strategies $(\bar{\sigma}_1^T, \bar{\sigma}_2^T)$ form a 2ϵ -equilibrium (Waugh et al. 2009). Thus, CFR constitutes an anytime algorithm for finding an ϵ -Nash equilibrium in zero-sum games.

Pruning Techniques

In this section we discuss pruning techniques that allow parts of the game tree to be skipped within CFR iterations.

(Partial) Pruning Typically, regret is updated by traversing each node in the game tree separately for each player, and calculating the contribution of a history $h \in I$ to $r^t(I, a)$ for each action $a \in A(I)$. If a history h is reached in which $\pi_{-i}^{\sigma^t}(h) = 0$ (that is, an opponent's reach is zero), then from (2) and (3) the strategy at h contributes nothing on iteration t to the regret of $I(h)$ (or to the information sets above it). Moreover, any history that would be reached beyond h would also contribute nothing to its information set's regret because $\pi_{-i}^{\sigma^t}(h') = 0$ for every history h' where $h \sqsubset h'$ and $P(h') = P(h)$. Thus, when traversing the game tree for player i , there is no need to traverse beyond any history h when $\pi_{-i}^{\sigma^t}(h) = 0$. The benefit of this form of pruning, which we refer to as *partial pruning*, varies depending on the game, but empirical results show a factor of 30 improvement in some small games (Lanctot et al. 2009).

Regret-Based Pruning (RBP) While partial pruning allows one to prune paths that an *opponent* reaches with zero probability, the recently introduced *regret-based pruning* (RBP) algorithm allows one to also prune paths that the *traverser* reaches with zero probability (Brown and Sandholm 2015). However, this pruning is necessarily temporary. Consider an action $a \in A(I)$ such that $\sigma^t(I, a) = 0$, and assume for now that it is known action a will not be played with positive probability until some far-future iteration t' (in RM, this would be the case if $R^t(I, a) \ll 0$). Since action a is played with zero probability on iteration t , the strategy played and reward received following action a (that is, in $D(I, a)$) will not contribute to the regret for any information set preceding action a on iteration t . In fact, what happens in $D(I, a)$ has no bearing on the rest of the game tree until iteration t' is reached. So one can “procrastinate” until iteration t' in deciding what happened beyond action a on iteration $t, t + 1, \dots, t' - 1$.

Upon reaching iteration t' , rather than individually making up the $t' - t$ iterations over $D(I, a)$, one can instead do a *single* iteration, playing against the *average* of the opponents’ strategies in the $t' - t$ iterations that were missed, and declare that strategy was played on all the $t' - t$ iterations. This accomplishes the work of the $t' - t$ iterations in a single traversal. Moreover, since player i never plays action a with positive probability between iterations t and $t' - 1$, that means every *other* player can apply partial pruning on that part of the game tree for the $t' - t$ iterations, and skip it completely. This, in turn, means that player i has free rein to play whatever she wants in $D(I, a)$ without affecting the regrets of the other players. In light of that, and of the fact that player i gets to decide what is played in $D(I, a)$ after knowing what the other players have played, player i might as well play a strategy that ensures zero regret for all information sets $I' \in D(I, a)$ in the iterations t to $t' - 1$. For instance, player i can play a best response to the opponents’ average strategy from iterations t to $t' - 1$; this is what we do in the experiments in this paper.

Regret-based pruning only allows a player to skip traversing $D(I, a)$ for as long as $\sigma^t(I, a) = 0$. Thus, in RM if $R^{t_0}(I, a) < 0$ we can prune the game tree beyond action a from iteration t_0 onward in consecutive iterations as long as for the current iteration t_1 we have

$$\sum_{t=1}^{t_0} v^{\sigma^t(I, a)} + \sum_{t=t_0+1}^{t_1} \pi_{-i}^{\sigma^t(I)} U(I, a) \leq \sum_{t=1}^{t_1} v^{\sigma^t(I)} \quad (10)$$

Once this no longer holds, skipping ceases. If we later find another t_0 that satisfies $R^{t_0}(I, a) < 0$, we do another sequence of iterations where we skip traversing after a , and so on.

3 Dynamic Thresholding

The pruning methods described in Section 2 can only be applied when an action is played with zero probability. This makes pruning incompatible with Hedge, because in Hedge all the action probabilities are strictly positive. This motivates our introduction of *dynamic thresholding*, in which low-probability actions are set to zero probability.

In dynamic thresholding for Hedge, when each successive iteration t is computed we set any action with probability less than $\frac{(C-1)\sqrt{\ln(|A(I)|)}}{\sqrt{2|A(I)|^2}\sqrt{t}}$ (where $C \geq 1$) to zero probability and normalize the remaining action probabilities accordingly so they sum to 1. We then use this new probability vector to determine the regrets of iteration $t + 1$. If an action is thresholded, this deviation from what Hedge calls for may lead to worse performance and therefore higher regret. In particular, since the altered probability vectors determine the regrets for future iterations, there is a risk that this error could snowball. However, using the threshold that we just specified above, we ensure that the new regret is within a constant factor C of the traditional regret bound.

Theorem 1. *If player $P(I)$ plays according to Hedge in an information set I for T iterations using threshold $\frac{(C-1)\sqrt{\ln(|A(I)|)}}{\sqrt{2|A(I)|^2}\sqrt{t}}$ with $C \geq 1$ on every iteration t , then $R^T(I) \leq C\sqrt{2}\Delta(I)\sqrt{\ln(|A(I)|)}\sqrt{T}$.*

All proofs can be found in an accompanying online appendix.

To apply the above theorem within CFR, we get from Equation 9 that one can then just sum the regrets of all information sets I to bound the total regret for this player.

Dynamic thresholding can in general be applied to any regret minimization algorithm. We present Theorem 1 specifically for Hedge in order to tailor the threshold for that algorithm, which provides a tighter theoretical bound. In Theorem 2, we also show that dynamic thresholding can be applied to RM. However, it results in very little, if any, additional pruning. This is because RM is very unlikely in practice to put extremely small probabilities on actions. Nevertheless, we prove that dynamic thresholding applies to RM for the sake of completeness and for its potential theoretical applications. Note that the formula for the threshold is now different.

Theorem 2. *If player $P(I)$ plays according to regret matching in an information set I for T iterations using threshold $\frac{C^2-1}{2C|A(I)|^2}\sqrt{t}$ with $C \geq 1$ on every iteration t , then $R^T(I) \leq C\Delta(I)\sqrt{|A(I)|}\sqrt{T}$.*

Again, to apply the above theorem within CFR, we get from Equation 9 that one can then just sum the regrets of all information sets I to bound the total regret for this player.

4 Regret-Based Pruning for Hedge

In this section we describe how dynamic thresholding enables regret-based pruning when using Hedge. To use RBP, it is necessary to determine a lower bound on the number of iterations for which an action will have zero probability. In RM without dynamic thresholding this is simply the minimum number of iterations it would take an action to achieve positive regret, as shown in (10). In Hedge with dynamic thresholding, we instead must determine the minimum number of iterations it would take for an action to reach probability above the dynamic threshold.

Let $R^{T_0}(I, a)$ be the regret for an action a in information set I on iteration T_0 . If $\sigma^{T_0}(I, a) < \frac{(C-1)\sqrt{\ln(|A(I)|)}}{\sqrt{2|A(I)|^2}\sqrt{T_0}}$, where

$\sigma^{T_0}(I, a)$ is defined according to (8), then pruning can begin on iteration T_0 . By Theorem 1, we can prune the game tree following action a on any consecutive iteration T after that if

$$\frac{e^{\eta_T (R^{T_0}(I, a) + U(I, a)(T - T_0))}}{\sum_{a' \in A(I)} e^{\eta_T (R^T(I, a') + \sum_{T'=T_0+1}^T v^{T'}(I, a'))}} < \frac{(C-1)\sqrt{\ln(|A(I)|)}}{\sqrt{2}|A(I)|^2\sqrt{t}} \quad (11)$$

Once this no longer holds, skipping ceases. If we later find another T_0 that satisfies the condition above, we do another sequence of iterations where we skip traversing after a , etc.

5 Experiments

We tested dynamic thresholding with and without RBP on a standard benchmark game called Leduc Hold'em (Southey et al. 2005) and an enlarged variant of Leduc Hold'em featuring more actions, called Leduc-5. Leduc Hold'em is a popular benchmark for imperfect-information game solving due to its feasible size and strategic complexity. In Leduc Hold'em, there is a deck consisting of six cards: two each of Jack, Queen, and King. There are two rounds. In the first round, each player places an ante of 1 chip in the pot and receives a single private card. A round of betting then takes place with a two-bet maximum, with Player 1 going first. A public shared card is then dealt face up and another round of betting takes place. Again, Player 1 goes first, and there is a two-bet maximum. If one of the players has a pair with the public card, that player wins. Otherwise, the player with the higher card wins. In standard Leduc Hold'em, all bets in the first round are 1 chip, while all bets in the second round are 2 chips. In Leduc-5, there are 5 bet sizes to choose from: in the first round the betting options are 1, 2, 4, 8, or 16 chips, while in the second round the betting options are 2, 4, 8, 16, or 32 chips. Leduc Hold'em contains 288 information sets, compared to 34, 224 for Leduc-5.

Hedge requires the user to set the tuning parameter η_t . When proving worst-case regret bounds, the parameter is usually defined as a function of $\Delta(I)$ for an information set I (for example, $\eta_t = \frac{\sqrt{8\ln(|A(I)|)}}{\Delta(I)\sqrt{t}}$) (Cesa-Bianchi and Lugosi 2006). However, this is overly pessimistic in practice, and better performance can be achieved with heuristics while still guaranteeing convergence, albeit at a weaker convergence bound.¹ In our experiments, we set $\eta_t = \frac{\sqrt{\ln(|A(I)|)}}{3\sqrt{\text{VAR}(I)_t}\sqrt{t}}$, where $\text{VAR}(I)_t$ is the observed variance of $v(I)$ up to iteration t , based on a heuristic by Chaudhuri et al. (2009).

In addition to the regret-minimization algorithms which are the main focus of this paper, for comparison we also

¹Convergence is still guaranteed so long as $\Delta(I)$ is replaced with a value that has a constant lower and upper bound, though the worst-case bound may be worse.

experimented with the leading gradient-based algorithm for finding ϵ -equilibrium in zero-sum games, the *excessive gap technique* (EGT) (Nesterov 2005; Hoda et al. 2010), coupled with the distance-generating function from Kroer et al. (2015). It converges to an ϵ -equilibrium in two-player zero-sum games in $O(\frac{1}{\epsilon})$ iterations, that is, in significantly fewer iterations than CFR which converges in $O(\frac{1}{\epsilon^2})$. In this EGT variant the gradient is computed by traversing the game tree. This enables pruning and dynamic thresholding to be implemented in EGT as well (Kroer et al. (2015) similarly computed the gradient through tree traversal in order to enable sampling schemes). In our experiments with EGT, we stop traversing a branch in the game tree when the probability (over nature and the opposing player) of the branch falls below $\frac{c}{T}$ for various values of c . We leave the theoretical verification of this approach as future work.

Figure 1 shows the performance of dynamic thresholding on Hedge and EGT against the vanilla versions of the algorithm as well as against the benchmark algorithms CFR+ and CFR with RM. We present our results with the number of nodes touched on the x axis. Nodes touched is a hardware- and implementation-independent proxy for time. Hedge involves exponentiation when determining strategies, which takes longer than the simple floating point operations of RM. In our implementation, regret matching traverses 36% more nodes per second than Hedge. However, in large-scale multi-core implementations of CFR, memory access is the bottleneck on performance and therefore the penalty for using Hedge should not be as significant.

The two figures on the left show that dynamic thresholding benefits EGT and Hedge, and the relative benefit increases with game size. In Leduc-5, dynamic thresholding improves the performance of EGT by a factor of 2, and dynamic thresholding combined with RBP improves the performance of CFR with Hedge by a factor of 7. The graphs on the right show that, when using thresholding and RBP, Hedge outperforms RM in Leduc, but RM outperforms Hedge in Leduc-5. RM's better performance in Leduc-5 is due to more widespread pruning than Hedge.

While CFR+ exhibits the best performance in Leduc-5, there are several drawbacks to the algorithm that cause it not to be usable in all settings. First, CFR+ is not known to converge when combined with RBP. The noisy performance of CFR+ with RBP in Leduc-5, and the weaker performance of CFR+ with RBP when compared with vanilla CFR+ in Leduc, may be consequences of this. Second, while CFR+ in practice outperforms CFR with RM or Hedge, it has a worse theoretical bound. Moreover, in the long run EGT appears to outperform CFR+. Finally, CFR+ is not known to be compatible with sampling, which is commonly used in large imperfect-information games.

Figure 2 shows the performance of EGT and Hedge with different aggressiveness of dynamic thresholding. For EGT, we threshold by $\frac{c}{T}$, where the number shown in the legend is c . For Hedge, we threshold by $\frac{d\sqrt{\ln(|A|)}}{\sqrt{2}|A|^2\sqrt{T}}$, where d is shown in the legend. The results show that for Hedge the performance is not sensitive to the parameter, and threshold only helps. For EGT, the results using $c = \{0.001, 0.005, 0.01\}$




Figure 1: Performance of EGT, CFR with Hedge, and CFR with RM on Leduc and Leduc-5. CFR with Hedge is shown without any pruning (vanilla Hedge), with dynamic thresholding, and with RBP. EGT is shown without any pruning (vanilla EGT) and with dynamic thresholding. CFR with RM is shown with partial pruning (vanilla RM) and with RBP. Dynamic thresholding on RM resulted in identical performance to vanilla RM, and is therefore not shown separately.




Figure 2: Varying the aggressiveness of dynamic thresholding.

are all similar and beneficial, while using a value of 0.05 is too aggressive, and hurts performance slightly.

6 Conclusion and Future Research

We introduced *dynamic thresholding* for online learning algorithms, in which a threshold is set at every iteration such that any action with probability below the threshold is set to zero probability. This enables pruning for the first time in a wide range of algorithms. We showed that it can be applied to both Regret Matching and Hedge—regardless of whether they are used in isolation for any problem or as subroutines at each information set within Counterfactual Regret Minimization, the most popular algorithm for solving imperfect-information game trees. We proved that the regret bound

increases by only a small constant factor, and each iteration becomes faster due to enhanced pruning. Our experiments demonstrated substantial speedup improvements in Hedge; the relative speedup increases with problem size.

We also developed a version of the leading gradient-based algorithm for solving imperfect-information games, the excessive gap technique coupled with the distance-generating function from Kroer et al. (2015), where we compute the gradient based on a traversal of the game tree and perform dynamic thresholding during this traversal. Experiments again showed that they lead to a significant speedup, and that the relative speedup increases with problem size.

Our results on Hedge might also be useful for boosting when Hedge is used therein (Freund and Schapire 1997;

Luo and Schapire 2014). The idea is that low-weight weak learners and/or low-weight training instances (as an analogy to low-probability actions in our paper) would then not need to be run, which may lead to significant time savings.

Future work also includes studying whether the idea of dynamic thresholding could be applied to other iterative algorithms that place at least some small positive probability on all actions (e.g. Pays (2014) or Daskalakis, Deckelbaum, and Kim (2015)).

Acknowledgments

This material is based on work supported by the NSF under grants IIS-1617590, IIS-1320620, and IIS-1546752, the ARO under award W911NF-16-1-0061, as well as XSEDE computing resources provided by the Pittsburgh Supercomputing Center.

References

- Blackwell, D. 1956. An analog of the minmax theorem for vector payoffs. *Pacific Journal of Mathematics* 6:1–8.
- Bowling, M.; Burch, N.; Johanson, M.; and Tammelin, O. 2015. Heads-up limit hold'em poker is solved. *Science* 347(6218):145–149.
- Brown, N., and Sandholm, T. 2014. Regret transfer and parameter optimization. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- Brown, N., and Sandholm, T. 2015. Regret-based pruning in extensive-form games. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*.
- Brown, N.; Ganzfried, S.; and Sandholm, T. 2015. Hierarchical abstraction, distributed equilibrium computation, and post-processing, with application to a champion no-limit Texas Hold'em agent. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.
- Cesa-Bianchi, N., and Lugosi, G. 2006. *Prediction, learning, and games*. Cambridge University Press.
- Cesa-Bianchi, N.; Mansour, Y.; and Stoltz, G. 2007. Improved second-order bounds for prediction with expert advice. *Machine Learning* 66(2-3):321–352.
- Chaudhuri, K.; Freund, Y.; and Hsu, D. J. 2009. A parameter-free hedging algorithm. In *Advances in neural information processing systems*, 297–305.
- Daskalakis, C.; Deckelbaum, A.; and Kim, A. 2015. Near-optimal no-regret algorithms for zero-sum games. *Games and Economic Behavior* 92:327–348.
- Freund, Y., and Schapire, R. 1997. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences* 55(1):119–139.
- Gibson, R.; Lanctot, M.; Burch, N.; Szafron, D.; and Bowling, M. 2012. Generalized sampling and variance in counterfactual regret minimization. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- Gibson, R. 2014. *Regret Minimization in Games and the Development of Champion Multiplayer Computer Poker-Playing Agents*. Ph.D. Dissertation, University of Alberta.
- Gilpin, A., and Sandholm, T. 2007. Lossless abstraction of imperfect information games. *Journal of the ACM* 54(5).
- Gilpin, A.; Peña, J.; and Sandholm, T. 2012. First-order algorithm with $\mathcal{O}(\ln(1/\epsilon))$ convergence for ϵ -equilibrium in two-person zero-sum games. *Mathematical Programming* 133(1–2):279–298. Conference version appeared in *AAAI-08*.
- Hart, S., and Mas-Colell, A. 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68:1127–1150.
- Hoda, S.; Gilpin, A.; Peña, J.; and Sandholm, T. 2010. Smoothing techniques for computing Nash equilibria of sequential games. *Mathematics of Operations Research* 35(2):494–512. Conference version appeared in *WINE-07*.
- Kroer, C.; Waugh, K.; Kılınç-Karzan, F.; and Sandholm, T. 2015. Faster first-order methods for extensive-form game solving. In *Proceedings of the ACM Conference on Economics and Computation (EC)*.
- Lanctot, M.; Waugh, K.; Zinkevich, M.; and Bowling, M. 2009. Monte Carlo sampling for regret minimization in extensive games. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*, 1078–1086.
- Littlestone, N., and Warmuth, M. K. 1994. The weighted majority algorithm. *Information and Computation* 108(2):212–261.
- Luo, H., and Schapire, R. E. 2014. A drifting-games analysis for online learning and applications to boosting. In *Advances in Neural Information Processing Systems*, 1368–1376.
- Nesterov, Y. 2005. Excessive gap technique in nonsmooth convex minimization. *SIAM Journal of Optimization* 16(1):235–249.
- Pays, F. 2014. An interior point approach to large games of incomplete information. In *AAAI Computer Poker Workshop*.
- Southey, F.; Bowling, M.; Larson, B.; Piccione, C.; Burch, N.; Billings, D.; and Rayner, C. 2005. Bayes' bluff: Opponent modelling in poker. In *Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, 550–558.
- Tammelin, O.; Burch, N.; Johanson, M.; and Bowling, M. 2015. Solving heads-up limit Texas hold'em. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*.
- Waugh, K.; Schnizlein, D.; Bowling, M.; and Szafron, D. 2009. Abstraction pathologies in extensive games. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.
- Zinkevich, M.; Bowling, M.; Johanson, M.; and Piccione, C. 2007. Regret minimization in games with incomplete information. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*.