

Regret-Grow CFR: A Novel Iterative Algorithm to Solve General-Sum Imperfect Information Games

Anonymous Authors

Abstract

General-sum (GS) games are decision problems that the payoffs of players are not necessarily correlated, which opposite is constant-sum (CS) games that the sum of payoffs is constant. Common poker card games can be modeled as CS games, which has aroused much attention. However, GS problems are not studied as thoroughly as CS problems. Moreover, GS games are more complex than CS games, especially when not all information is publicly accessible to all players, which is known as imperfect information games (IIGs). In the theoretical view, CS-IIGs are cases degenerate to bilinear saddle point problems and enjoy various well-form properties, while GS-IIGs are linear-quadratic programming problems, and the payoffs of players may vary at different order. Such differences have made methods like first-order methods and various variants of counterfactual regret (CFR) minimization that are effective in CS-IIGs perform significantly poorly in GS-IIGs. In this paper, we proposed regret-grow CFR, a novel CFR variant for solving GS-IIGs. The core idea of our algorithm is to control the inertia of cumulative regret to grow at a sub-linear rate, in order to adapt the independent updating of bilateral payoffs. In detail, it makes the strategies vary smoothly enough that is capable of payoff discontinuity in GS-IIGs. We provided formal proof of the effectiveness of forgetting the negative instant regret at every iteration. Then in the experiment, we used multiple GS variants of Goofspiel as a test bed to compare regret-grow CFR and existing common methods. It turns out that our algorithm is capable of converging to NE even if other methods fail, and significantly lower the exploitability of players in GS-IIGs.

Introduction

General-sum, imperfect information game is a theorized model to formulate many sequential move real-life decision problems. Imperfect-information games (IIGs) model strategic interactions among a set of participants who make decisions with imperfect information. General-sum (GS) game is a game in which the payoffs of players are not necessarily correlated. While quite a few algorithms like first-order methods and some CFR variants are able to tackle the feature of zero-sum problems, few of them can be applied to GS-IIGs. In a GS setting, the players' payoffs are mutually

independent. For example, even solving a two-player GS-IIG by vanilla CFR is harder than its zero-sum counterpart, because one player's payoff can change considerably while the other player just changes slightly.

Extensive form layout is common in large imperfect information games, which resembles a decision tree in that players make decisions based on probability. For sufficiently large IIGs that are infeasible for a linear program, iterative algorithms are preferred because players' strategies often keep optimizing during a single affordable iteration. The ultimate goal is commonly to find a Nash Equilibrium (Nash 1950) in which no player can improve by unilaterally adjusting the strategy. Some special types of problems like zero-sum IIGs, already enjoy acceptable solutions by multiple algorithms. For example, 6-card and 10-card Leduc can be essentially be solved by First Order Method (Kroer et al. 2015). CFR, proposed in (Martin et al. 2007) and CFR+ proposed in (Tammelin 2014) to solve heads-up limit Texas hold'em problem (Brown and Sandholm 2018, 2019; Brown, Ganzfried, and Sandholm 2015) and find the baseline solution of heads-up no-limit Texas hold'em (HUNL) endgames in Libratus (Silver et al. 2016, 2017, 2018). And for lower space complexity, MCCFR (Marc et al. 2009; Martin et al. 2019) are variants samples the game tree in a probabilistic manner. In order to solve the problem that exhibits ill-condition values, Tuomas Sandholm and Noam Brown (Noam et al. 2019) proposed a discount mechanism for assigning different weights for every iterations. (Silver et al. 2016, 2017, 2018; Schrittwieser et al. 2020)

Two major features have made GS-IIGs challenging. Firstly, the imperfect information makes payoffs of players rely on their opponents. Since the payoff of an action is the expectation from all possible cases, a strategy update from the opponent may cause a change in these probabilities of the cases, and therefore that expectation. Thus players cannot simply take action with the highest payoff. In fact, approximating such a Nash equilibrium is at most PPAD complete (Chen, Deng, and Teng 2009; Song et al. 2019). Meanwhile, although there are known algorithms that can solve zero-sum games efficiently, whether these algorithms are still qualified for GS-IIGs should be speculated. Since the independence of player payoffs, many algorithms, for example, Mini-Max, are no longer providing optimal solutions, because it assumes that payoff of one player is the

opposite of that of the other. Even for algorithms that are capable of approaching a Nash Equilibrium, the behavior of these algorithms near the equilibrium points, i.e. time complexity and stability should be studied rigorously.

In this paper, we proposed regret-grow CFR for GS-IIGs. Firstly, we proposed the selective forgetting of instant regret, and applied it at every iteration to make the cumulative regret grow in expectation. Secondly, the smoothness of the iteration process is shown, which is crucial because payoffs in GS-IIGs have 1-order lower smoothness than that of ZS-IIGs. Thirdly, the updating stride is self-adjusting and self-limiting, which has fewer parameters to tune and makes the algorithm self-adaptive to wide-range payoffs. The sub-linearly yet growing regret value of regret-grow CFR has both theoretical basis and empirical significance. Finally, both regret-grow CFR and other CFR variants are applied to solve the GS variant of Goofspiel, and their behavior both at iterating and near equilibrium is examined. It is shown that there are cases that some existing CFR variants that skip some equilibrium points and deviate toward false values. However, our algorithm on the same customized Goofspiel converges towards Nash equilibrium points correctly, yet eliminates dominated strategy at the same converge rate as the current CFR+ algorithm.

Related Work

Finding a Nash equilibrium is an important, interesting, and well-studied problem. Like two-player zero-sum games can be solved in polynomial time, finding (even an approximate) Nash equilibrium in a two-player GS game is PPAD-complete (Chen, Deng, and Teng 2009), and a correlated equilibrium can be computed efficiently (Jiang and Leyton-Brown 2011). There are popular methods proposed for solving GS-IIGs, most of them can be categorized into vertex enumerating methods, reactive module methods, and no-regret methods.

The polytope vertex enumerating methods focus on extremely small-scaled games that can be converted into norm-form games. One of the methods, proposed by (Avis et al. 2009), is to enumerating all the possible combinations for actions with non-zero probability. The combinations of both sides should contain the same amount of actions. By establishing equations that let actions share the same payoff, the opponent's probability of each action can be calculated. This method is based on the fact that the actions presented in a mixed strategy should have their payoffs equal, and equal to the maximum of all legal actions' payoff. While this algorithm has a rigorous mathematical basis and can be used to examine whether a strategy profile is a Nash equilibrium, two fatal limitations make it not suitable for large GS-IIGs: the algorithm is native to norm-form but not friendly to extensive-form, and even in norm-form, the computational time for an $n \times n$ norm-form is 4^n and not a polynomial time complexity. In principle, one can find all equilibria since the nonlinear equations are polynomials. The idea is to enumerate all supports, solve all roots of the polynomial equations, and select the solutions that correspond to probability distributions. The methods of finding all equilibria are probabilistic, that is, they will find all solutions with a given prob-

ability when they are run for at least some amount of time (which depends on the probability). There are exponentially many supports in the game and there can be exponentially many equilibria. Moreover, the homotopy methods (global Newton, tracing procedure, or quantal response method) are not guaranteed to find all equilibria.

The homotopy methods that use the global Newton method do not converge globally. Govindan and Wilson observe that the iterated polymatrix approximation method typically converges globally but is not failsafe and may get stuck in some games. They find that the problem with homotopy methods is that they need to traverse nonlinear paths and require many small steps in order to obtain reasonable accuracy. They also observe that the homotopy path may have many twists and reversals. Goldberg et al. construct examples where homotopy methods will not only need an exponential number of pivots but also an exponential number of direction reversals. Herings and van den Elzen and Herings and Peeters present a globally convergent homotopy method but note that the triangulations must have very refined mesh and the homotopy path must be traced numerically. Reinforcement learning, focusing on end-to-end solution, are also used in GS-IIGs. Gutierrez et. al (Gutierrez Julian and Michael 2000) studied n-player GS games in which the choices available to players are defined using the Simple Reactive Modules Language (SRML), a subset of Reactive Modules (Alur and Henzinger 1999), a popular and expressive system modeling language that is used in several practical modeling checking systems (e.g., MOCHA (Alur et al. 1998) and Prism (Kwiatkowska, Norman, and Parker 2011)). Reactive Modules support succinct and high-level modeling of concurrent and multi-agent systems. In the games we study, the preferences of system components are specified by associating with each player in the game a temporal logic (LTL) formula that the player desires to be satisfied. Reactive Modules Games with perfect information (where each player can see the entire system state) have been extensively studied (Gutierrez, Harrenstein, and Wooldridge 2015a).

The regret-based methods are proposed in early 2010s, by making strategies proportional to the positive regret values. And by minimizing all actions' regret on all the information sets, the global regret value can be minimized by minimizing the immediate regret values, or simply put, per-information-set no-regret method. This makes the actions easily computed at drastically low computational overhead. However, Zinkevich et al uses average strategy as output, while instant strategy at each iterations (its state variable) varies wildly, which is detrimental of numerical stability and convergence when the algorithm are applied on general-sum games. Also, the naive method of taking average strategy makes it hard and intractable to rule out sub-optimal actions. While the algorithms runs good on zero-sum game, its intrinsic numerical defects brings it into technical difficulties on general-sum games. The perfect information games allows induction, which always go extreme and produces pure strategies as their equilibrium points, which is typically not the case in imperfect games. Although there are proofs shows that CFR can converge to Nash equilibrium in zero-sum games, and

the necessary condition for CFR converges is exactly Nash equilibrium, the proof for that CFR will converge is still absent.

Preliminaries

An **imperfect information game** has both normal-form and extensive-form representations, the latter is preferred because it is native to large problems. An extensive-form game is a decision tree that starts from the initial game state called “root”. There is a set for all the **players** called P . Each node is a game state identified by its **history** h , a sequence of all actions lead to that game state. If there are any actions available, $P(h)$ is the player at this node who take action **action** a from $A(h)$, such actions lead to **child nodes** that represent the new game states. If h is **prefix** of another node’s history h' , then it is called $h \sqsubseteq h'$. Let H be the set of all the histories, for histories identified as h and h' , For the **leaf nodes** who has no available actions and terminates the game, their history sequences are not any prefix of other histories, we use $Z \in H$ for all these nodes. All players will receive a **payoff** or reward when the game reaches leaf nodes, for each $i \in P$, that is $u_i(z)$. We denote the range of payoffs in the game by Δ , and Δ_i represent the difference between maximal and minimal payoffs for player $i \in P$.

In imperfect-information games, it is common that different nodes of same player appear identical to that player. This is because not all actions are publicly observable by every players. Such set is one of the **information sets** (or infosets) of that player. The collection of all such infosets of player i is **information collection** I_i . The information sets are the minimum units many strategic related quantities, is called $A(I)$ that all available actions on information sets.

In extensive game, the player choose action by a stochastic manner, at each information set S , all the players assign a distribution on each available action a . Every player has its **strategy** σ_i that is a mapping that maps every information sets $S \in I_i$ with a vector $\mathbf{R}^{|A(S)|}$, namely, $\sigma_i(S) \in \mathbf{R}^{|A(S)|}$. It is common that σ_{-i} is used as other players’ strategy. The set of all players’ strategy, σ , is called strategy profile.

Nash Equilibrium

Let $u_i(\sigma_i, \sigma)$ to be the player i ’s payoff. The Nash equilibrium is a strategy profile σ that every unilateral changes in σ_i profile will not increase $u_i(\sigma_i, \sigma_{-i})$, i.e.

$$\forall i, u_i(\sigma_i, \sigma_{-i}) = \max_{\sigma'_i} u_i(\sigma'_i, \sigma_{-i}) \quad (1)$$

For measuring how far the players are deviating from the equilibrium, exploitability is defined as:

$$ep_i(\sigma_i, \sigma_{-i}) = \max_{\sigma'_i} u_i(\sigma'_i, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i}) \quad (2)$$

$$ep(\sigma) = \sum_{i \in P} ep_i(\sigma_i, \sigma_{-i}) \quad (3)$$

By definition, we have:

$$ep(\sigma^*) = \sum_{i \in P} ep_i(\sigma_i^*, \sigma_{-i}^*) = 0 \quad (4)$$

The two-player norm form games focus on single information set. Each player’s payoff can be defined as entries in two matrices, let (A, B) be a binary tuple of m-by-n matrices, therefore m and n are numbers of their available actions. Let x and y as strategies of both players. both have their entities non-negative and sum to be exactly 1. As a natural result, u and v as their payoffs.

Best response condition

Let x and y be the mixed strategies of both players. Then those actions whose probabilities are non-zero have their payoffs to be maximal, and therefore mutually equal to others.

$$x_i > 0 \iff a_i^T y = u = \max_i (a_i^T y) \quad (5)$$

where a_i are row vectors of matrix A , and:

$$y_j > 0 \iff b_j^T x = v = \max_j (b_j^T x) \quad (6)$$

where b_j are column vectors of matrix B .

What it has alleviated is the infinite mixed strategy problem to finite-dimension inequalities formation, which however at the cost of numerical behavior of best responses. The collection of best responses would drop at almost all but one of its elements even if an opponent’s strategy deviates a little. Nevertheless, the algorithm can be used reversely, say, for example, player 1, not to find what pure actions are the collection I of best pure strategies should play 1 player against player 2’s y , but when I is potentially possible to become the collections of player’s y .

Methodology

To make both eq(5) and eq(6) have a unique solution, for example, if there are k non-zero entries in x , namely x_1 , and the rest of zero-entries x_0 , the linear problem should contain exactly k equations. Let y_1 be a non-zero part of y . From C_n^k possible different y_1 s, they form $k \times k$ linear equation, which is required by the uniqueness of the solution.

Clearly, this method provides Nash equilibrium points at the cost of NP-hard, by enumerating all the 1 to $\min(m, n)$, it requires all the $2^{\min(m, n)}$. Conversely, the counterfactual minimization method provides $\mathcal{O}(\frac{1}{\epsilon^2})$. So the NonZeroSum-Matrix method is only tractable in small sized information sets, and should be act as benchmark to test whether other algorithms could find equilibrium in test-size problem.

In a nonzero-sum game, minimax is no longer optimal, because it wrongly assumes that both play-ers use the same payoff function. Nonetheless, A’s minimax does guarantee the worst-case outcome for A, because it proceeds as if B would always choose the worst possible moves against A. Therefore, minimax is used as the baseline for comparisons in our examples. More generally, we consider imperfect information nonzero-sum games, in which players can have incomplete mutual knowledge and thus SPE does not apply.

Inspired the existing algorithms, the counterfactual regret minimization is slightly different naively applying gradient-based optimization method, but used convex-combinations

Algorithm 1: Regret Grow CFR

Input: A, B
Output: strategy profile x, y

```

1: function General-Sum-Matrix ( $A, B$ ):
2:  $m, n = A.shape$ 
3: for  $k \leftarrow 1$  to  $\min(m, n)$  do
4:   for  $I : sum(I) = k, I \in R^m, I_i \in 0, 1$  do
5:     for  $J : sum(J) = k, J \in R^n, J_j \in 0, 1$  do
6:        $y = A_{[I, J]}^{-1} \mathbf{1}$ 
7:       if  $0 \leq y \leq 1, Ay \leq 1$  then
8:          $I_{best} = y$ 
9:       end if
10:    end for
11:   end for
12:   for  $J : sum(J) = k, J \in R^n, J_j \in 0, 1$  do
13:     for  $I : sum(I) = k, I \in R^m, I_i \in 0, 1$  do
14:        $x = B_{[I, J]}^{-1} \mathbf{1}$ 
15:       if  $0 \leq x \leq 1, B^T x \leq 1$  then
16:          $J_{best} = x$ 
17:       end if
18:     end for
19:   end for
20:   for  $I : sum(I) = k, I \in R^m, I_i \in 0, 1$  do
21:     for  $J : sum(J) = k, J \in R^n, J_j \in 0, 1$  do
22:        $J' = I_{best-J} = where(I_{best} > 0)$ 
23:        $I' = J'_{best-I} = where(J'_{best} > 0)$ 
24:       if  $I' = I$  then
25:          $x = I_{best}$ 
26:          $y = J_{best}$ 
27:       end if
28:     end for
29:   end for
30: end for
31: return  $x, y$ 

```

instead. For any x

$$\sum_{i=1}^m x_i = 1, 0 \leq x \leq 1 \quad (7)$$

and similar y , there exist:

$$u(x, y) = x^T A y \quad (8)$$

$$v(x, y) = x^T B y \quad (9)$$

In typical CFR+ algorithm, since Taylor expansions is valid when iteration T approaching to infinity, let R as regret on all the actions, as the rule of the iteration have:

$$x' = \frac{R + r}{\sum R + \sum r} \quad (10)$$

let

$$p = \frac{r}{\sum r} \quad (11)$$

$$x' = \frac{\sum R x + \sum r p}{\sum R + \sum r} \quad (12)$$

$$x' = \frac{\sum R x + \sum r x - \sum r x + \sum r p}{\sum R + \sum r} \quad (13)$$

$$x' = x + \frac{\sum r}{\sum R + \sum r} p \quad (14)$$

$$x' \approx x + \frac{\sum r}{\sum R} p \quad (15)$$

Since the all the regret vector r comes from strictly positive actions that has better response for opponent, the payoff functions $u(x', y) > u(x, y)$ always holds.

$$R^{T+1}(I, a) = \begin{cases} R_+^T(I, a) + r^T(I, a), & \text{if } r^T(I, a) \geq 0 \\ R_+^T(I, a) + \epsilon r^T(I, a), & \text{otherwise} \end{cases} \quad (16)$$

where $\epsilon \in [0, 1]$, and R is truncated if negative, i.e.

$$R_+^{T+1}(I, a) = \begin{cases} R^{T+1}(I, a), & \text{if } R^{T+1}(I, a) \geq 0 \\ 0, & \text{otherwise} \end{cases} \quad (17)$$

Definition 1 Let f_i from actions $i \in \{1, 2, \dots, m\}$ to be approximation of best strategy y , player's strategy \hat{p} is convex combination of f_i , and a non-negative loss function $\ell(\hat{p}, y)$. Then the instantaneous regret value for \hat{p} deviating away from f_i is defined as $r_i = \ell(\hat{p}) - \ell(f_i)$.

Definition 2 If approximation above is repeated for n times, then the cumulative loss functions for player and advisor $i \in \{1, 2, \dots, m\}$ are defined as $\hat{L}_n = \sum_{t=1}^n \ell(\hat{p}_t, y_t)$ and

$L_{i,n} = \sum_{t=1}^n \ell(f_{i,t}, y_t)$ respectively, and the cumulative regret

is defined as $R_{i,n} = \sum_{t=1}^n r_{i,t} = \hat{L}_n - L_{i,n}$

Theorem 1 Let ϕ to be function from R to R_+ is a non-negative, convex and increasing function, then

$$\sup_{y_t} \sum_{i=1}^N r_{i,t} \phi'(R_{i,t-1}) \leq 0$$

Proof:

Since $\phi'(R_{i,t-1}) > 0$, using Jensen's inequality for all y ,

$$\ell(\hat{p}_t, y) = \ell\left(\frac{\sum_{i=1}^N \phi'(R_{i,t-1}) f_{i,t}}{\sum_{j=1}^N \phi'(R_{j,t-1})}, y\right) \leq \frac{\sum_{i=1}^N \phi'(R_{i,t-1}) \ell(f_{i,t}, y)}{\sum_{i=1}^N \phi'(R_{j,t-1})}$$

Lemma 1 Let $\mathbf{r}_t = (r_{1,t}, r_{2,t}, \dots, r_{N,t}) \in R^N$ to be instantaneous regret vector, and cumulative regret vector

$\mathbf{R}_n = \sum_{t=1}^n \mathbf{r}_t$. Then the potential function $\Phi : R^N \rightarrow$

R_+ is defined as $\Phi(\mathbf{u}) = \psi\left(\sum_{i=1}^N \phi(u_i)\right)$, where $\phi :$

$R \rightarrow R_+$ is any non-negative increasing function, and $\psi : R_+ \rightarrow R_+$ is any non-negative function for scaling purpose with strictly increasing and concave properties.

Then $\hat{p}_t = \frac{\nabla \Phi(\mathbf{R}_{t-1}) \mathbf{f}_t}{\sum_{j=1}^N \nabla \Phi(\mathbf{R}_{t-1})_j}$, and theorem 7.3 is equivalent to

$$\sup_{y_t} \mathbf{r}_t \cdot \nabla \Phi(\mathbf{R}_{t-1}) \leq 0$$

Lemma 2 $\Phi(\mathbf{R}_n) \leq \Phi(0) + \frac{1}{2} \sum_{t=1}^n C(\mathbf{r}_t)$, where $C(\mathbf{r}_t) =$

$$\sup_{u \in R^N} \left[\psi' \left(\sum_{i=1}^N \phi(u_i) \right) \sum_{i=1}^N \phi''(u_i) r_{i,t}^2 \right]$$

Proof:

$$\Phi(\mathbf{R}_t) = \Phi(\mathbf{R}_{t-1} + \mathbf{r}_t)$$

$$= \Phi(\mathbf{R}_{t-1}) + \nabla \Phi(\mathbf{R}_{t-1}) \cdot \mathbf{r}_t + \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \frac{\partial^2 \Phi}{\partial u_i \partial u_j} r_{i,t} r_{j,t}$$

$$\leq \Phi(\mathbf{R}_{t-1}) + \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \frac{\partial^2 \Phi}{\partial u_i \partial u_j} r_{i,t} r_{j,t}$$

where the second-order term of Taylor expansion shows that

$$\begin{aligned} & \sum_{i=1}^N \sum_{j=1}^N \frac{\partial^2 \Phi}{\partial u_i \partial u_j} r_{i,t} r_{j,t} \\ & \leq \psi'' \left(\sum_{i=1}^N \psi(\xi_i) \right) \sum_{i=1}^N \sum_{j=1}^N \psi'(\xi_i) \psi'(\xi_j) r_{i,t} r_{j,t} \\ & + \psi' \left(\sum_{i=1}^N \psi(\xi_i) \sum_{i=1}^N \psi''(\xi_i) r_{i,t}^2 \right) \\ & = \psi'' \left(\sum_{i=1}^N \psi(\xi_i) \right) \left(\sum_{i=1}^N \psi'(\xi_i) r_{i,t} \right)^2 \\ & + \psi' \left(\sum_{i=1}^N \psi(\xi_i) \sum_{i=1}^N \psi''(\xi_i) r_{i,t}^2 \right) \\ & \leq C(\mathbf{r}_t) \end{aligned}$$

Theorem 2 For any convex loss function ℓ , if it takes values in $[0, 1]$, if scaling function ψ is polynomial weighted function, then for any sequence y_1, y_2, \dots, y_n the loss function have $\hat{L}_n - \min_{i=1,2,\dots,N} L_{i,n} \leq \sqrt{n(p-1)N^{2/p}}$, which also means that regret value is $o(n)$ when $n \rightarrow \infty$

Proof:

Since $\psi'(x) = (x^{\frac{2}{p}})' = \frac{2}{p x^{(p-2)/p}}$, and $\psi''(x) = (x^{\frac{2}{p}})'' = p(p-1)x_+^{p-2}$, where x_+ floors negative components to zero while keeps the positive component. By Holder inequality,

$$\begin{aligned} & \sum_{i=1}^N \psi''(u_i) r_{i,t}^2 \\ & \leq p(p-1) \left(\sum_{i=1}^N \left((u_i)_+^{p-2} \right)^{p/(p-2)} \right)^{(p-2)/p} \left(\sum_{i=1}^N |r_{i,t}|^p \right)^{2/p} \end{aligned}$$

Thus,

$$\begin{aligned} & \psi \left(\sum_{i=1}^N \psi(u_i) \right) \sum_{i=1}^N \psi''(u_i) r_{i,t}^2 \\ & \leq 2(p-1) \left(\sum_{i=1}^N |r_{i,t}|^p \right)^{2/p} \end{aligned}$$

which means that

$$\Phi_p(\mathbf{R}_n) \leq (p-1) \sum_{i=1}^n \|\mathbf{r}_t\|_p^2$$

$\leq n(p-1)N^{2/p}$ which means that the regret grows only sub-linearly. i.e. $\frac{\mathbf{R}_T}{T} \rightarrow 0$ when $T \rightarrow \infty$, it's asymptotically approaching to best response of player should follow with.

Intuition of the Regret-Grow CFR

What inspired the CFR is the convex combination with better actions, however, what to be maximized is a multiple-objective $u_1(x, y)$ and $u_2(x, y)$. The exploitability provides no-exploitability method just like existing CFR+ algorithm,

which can also reformulate two-player general-sum games into single object optimization problem:

$$\max_{x,y} (u(x, y) - \max_p u(p, y) + v(x, y) - \max_q v(x, q)) \quad (18)$$

For the iterating method, we added exploitability terms for augmenting the existing CFR method:

$$r_{x_i} = u(x_i, y) - u(x, y) \quad (19)$$

$$r_{x_{epy}i} = -\max_q x_i, q \quad (20)$$

$$x' = x + \frac{r_x}{\sum R_x} + \frac{r_{x_{epy}}}{\sum R_{x_{epy}}} \quad (21)$$

$$y' = y + \frac{r_y}{\sum R_y} + \frac{r_{y_{epx}}}{\sum R_{y_{epx}}} \quad (22)$$

This will perform maximize the player's own payoff, and minimize the opponent's exploitability, which degenerates to existing CFR when the problem is just zero-sum cases.

Experiments

We evaluate the performance of regret-grow CFR on Goofspiel. Goofspiel is a bidding card game where players have a hand of cards numbered 1 to N, and take turns secretly bidding on the top point-valued card in a point card stack using cards in their hands. While all of these games have imperfect information and are roughly of similar size, they are a diverse set of games, varying both in the degree (the ratio of the number of information sets to the number of histories) and nature (whether due to chance or opponent actions) of imperfect information.




Figure 1: Strategies and regrets for players to decide their first cards.



Figure 2: Strategies and regrets for players to decide their first cards.




Figure 3: How payoff of pure and mixed strategies evolves at the first card.




Figure 4: How exploitability of A and B and their sum are minimized at the first card.

The equilibrium points

We also used subgames of Goofspiel variation as a test bed for the techniques introduced in this paper. First, for finding the equilibrium points, we used a different assessment of the effectiveness of the algorithm in terms of availability than the original one. The new evaluation criterion is defined as a head-to-head comparison between the adversarial sides, weighted differently, and the contrasting algorithms respectively, counting the final benefit of both sides.

Since that customized variant of Goofspiel was used as a test-bed for the techniques of Regret-Grow-CFR, in this experiment, all the Goofspiel upcards are treated as 1, but weighted as [5.00, 1.33, 2.71, 4.27] for player 1, and [4.10, 6.28, 3.33, 3.84] for player 2. When one player wins a card, this contributes their payoff by how the card weighted by this player, while the other's decrease by how this card weighted by that player.

The first chart reveals how both players make decision at first card, this is a mixed strategy profile, which suggests both the players bet their largest card. The empirical convergence rate at won't take effect on initial few turns, rather, it diverge away final equilibrium point by 0.771 for player 1 and 0.479 for player 2.

The result of how the players deal with their second card is presented in Figure 3 and 4. First, their strategies are converging, and therefore the equilibrium point's strategy profile is found. This can also be verified from the view of regret controlling - for both players and at each point where they make decisions, the sum of regrets for all available actions grows sub-linearly, this is also empirically verified the regret-based theories.

An interesting result from Figure 3 and Figure 4 is, once entered into equilibrium, both players' strategies, typically a mixed one, is consisted of the same number of pure strategies, however, this is a natural result of Nash theorem, since the pure strategy, will almost always have a unique best response from the opponent, which will never make opponent take a mixed strategy.

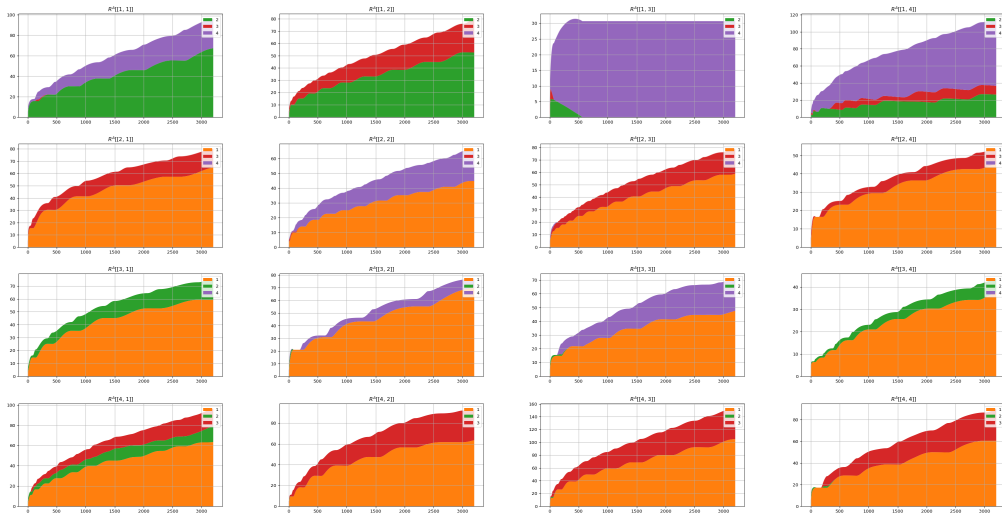


Figure 5: player A, regret values for available actions

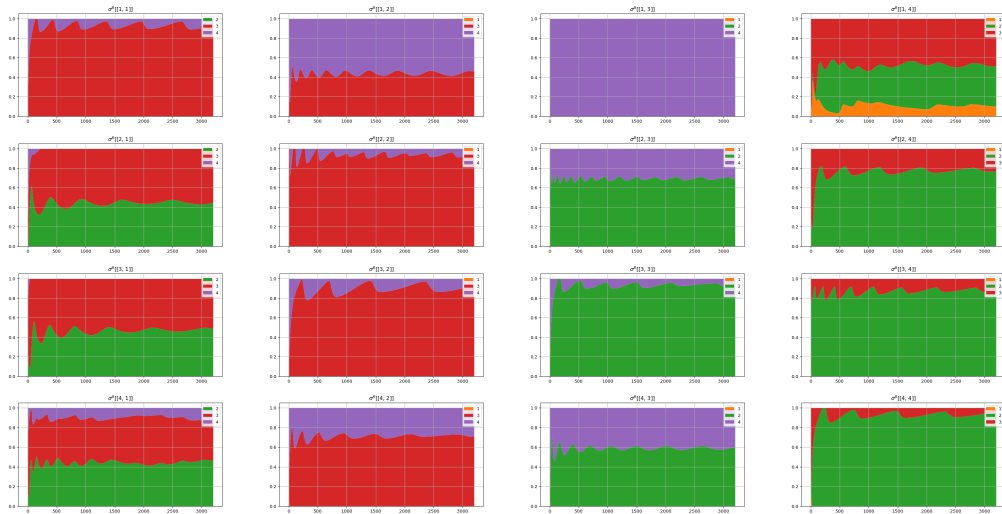


Figure 6: player B, strategy distribution on available actions

For example, an equilibrium that both players' strategies are mixed strategy that contains 2 actions: If player A adopts 2 actions in his mixed strategy, then player B will almost always adopt a mixed strategy that contains 2 actions. The reason can be discussed by considering 3 cases.

(1) If B's strategy contains only one action, i.e. pure strategy, then that is impossible. Because the action that B chose will typically not make 2 actions of A have same payoffs.

(2) Mixed strategy response from B that contain 2 actions is possible, since B's some partitions of probability can make A's two actions have the same payoff.

(3) Mixed strategy response from B that contain 3 actions or beyond is also impossible. If the additional action is inferior to 2 existing actions, this sub-optimal action will be ruled out, resulting in a 2×2 mixed strategy equilibrium. If the additional action is better than the 2 existing actions, then B will prefer the new action, shifting and eventually

landing at a new equilibrium.

In all the cases above, strategies from either side must contain the same number of pure strategies, or actions.

The result of how the players deal with their second card is presented in Figure 3 and 4. First, their strategies are converging, and therefore the equilibrium point's strategy profile is found. This can also be verified from the view of regret controlling - for both players and at each point where they make decisions, the sum of regrets for all available actions grows sub-linearly, this is also empirically verified the regret-based theories.

From figure 3, 7, and 8, it can be verified that the curve of cfv values always try to follow the curve of actions with the highest payoff, using a steady yet fast method to follow. If two or more actions are the best responses that comprised the mixed strategy, their payoffs compete and take turns to lead other actions. Be sure to not confuse it with a

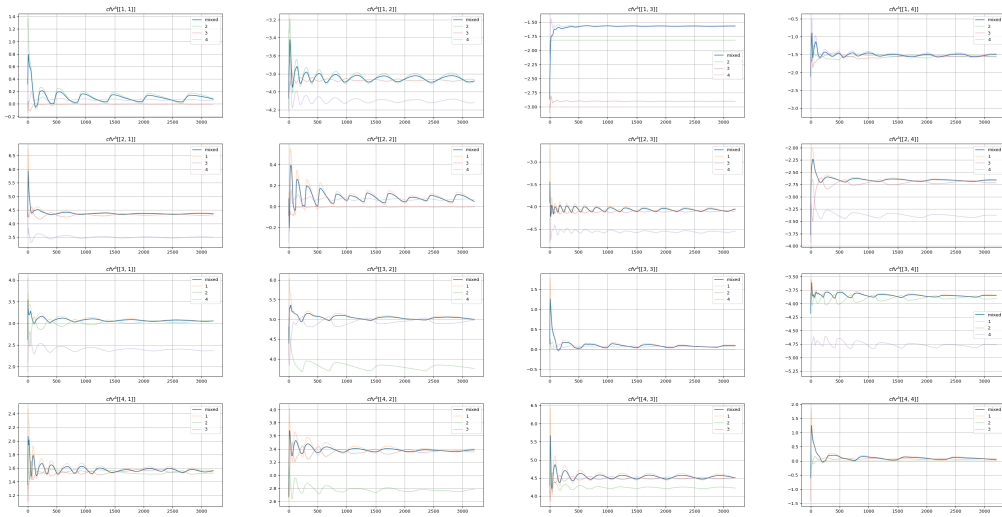


Figure 7: player A in subgames, payoff for strategies and actions

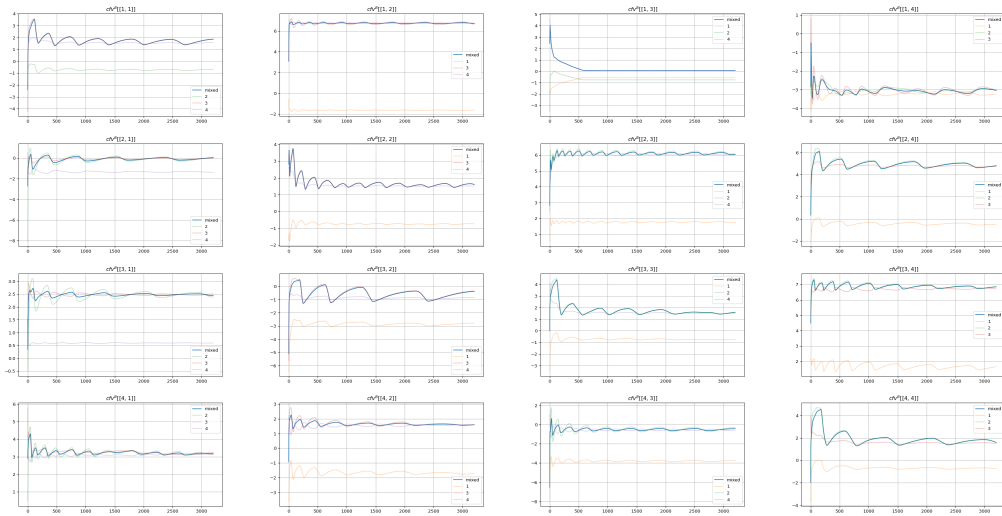


Figure 8: player B in subgames, payoff for strategies and actions

truly sub-optimal action, which disadvantage is permanent and can never be overturned.

While the payoffs of actions can be oscillating wildly, the curve of cvf adopts a fast-yet-smooth pattern to realize payoff maximization and guarantee convergence. This is because the negative instant regret is vanishing and ruling out the sub-optimal action, and also because the regrets sum across the actions' are growing yet sub-linearly, which makes the updating step length smaller across the times, yet allow the significant updates influencing the subsequent iterations. The limited-forgetting mechanism makes the regret grows over time, and limiting the step length of updating strategy; yet still makes the regret grows strictly slower than linear increasing, which may force the strategy not to be a sub-optimal one, and therefore guarantees converging to a Nash equilibrium.

Behavior near Subgame equilibrium

Since we considered the problem of computing an equilibrium solution for general-sum games. The most common solution concept is the Nash equilibrium. For $\epsilon \geq 0$, a strategy profile σ is an ϵ -Nash equilibrium if no player deviates from it.

The importance of deeper analyzing the endgame of Goofspiel variation lies in testing the behavior of Regret-Grow-CFR, especially, at some vicinity of any possible equilibrium points. In our experiments, this results in a different assessment of the effectiveness of the algorithm in terms of availability than the original one. The new evaluation criterion is defined as a head-to-head comparison between the adversarial sides, weighted differently, and the contrasting algorithms respectively, counting the final benefit of both sides.

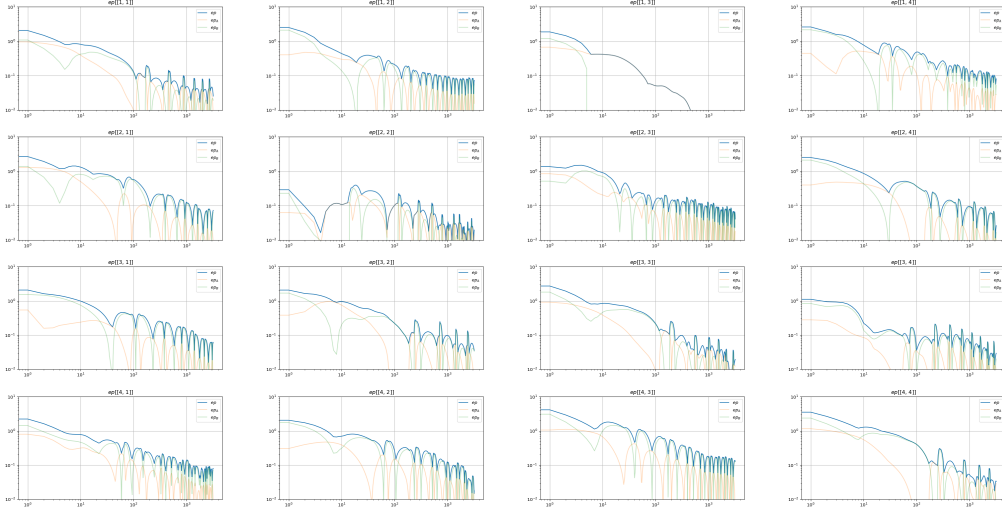


Figure 9: exploitability, for both players and their sum

Figure 5 and 6 have shown 16 distinct subgames in which player A and B deal with high loss, and show how their strategies evolved. They had also shown how the negative influence of bad initial guess is dissolved when the strategy was significantly off-equilibrium, which results can be seen in Figure 7 and 8.

While higher card weight draws much attention from both players, the chaining logic makes players deliberately give up the high-weighted cards by throwing low-ranking cards, this can be verified in the subgame $[(1, 1)]$. For the converging process, the average time lengths for ruling out dominated strategy and perform a fully-updated strategy for both players are 33.94 and 59.67 epochs respectively. More specifically, as Figure 1, 5, and 6 shown, player A bet more often his card-4 for a 5.00 reward, while player B bet a little mixed strategy, which throws card-1 at 34.7%, but concentrates more on his second card for a 6.28 reward. In all the selected 16 subgames, both players' exploitability value and that of the sum are shrinking, and that time complexity is no more than the time complexity of $O(1/T^2)$, which is plotted in dual-logarithm coordinates in Figure 9.

In the experiment, The first chart reveals how both players make a decision at the first card, this is a mixed strategy profile, which suggests both the players bet their largest card. The empirical converge rate at won't take effect on the initial few turns, rather, it diverges away the final equilibrium point by 0.771 for player A and 0.479 for player B.

The game Goofspiel variant is a two-player imperfect game that have multiple subgame nodes, on which all players reveal their actions so that players have full knowledge about their current situations. This makes Goofspiel simple yet capable to catch the complexity of the imperfect games. From equation (15), it is clear that the process of exploitability minimizing is the only necessary condition for a ϵ -Nash equilibrium. This effectively boost the process to both inspect and verify the convergent behavior and their speed near Nash equilibrium.

The exploitability of both players at the subgame $[(1, 3), (2, 3)]$ has three possible equilibrium points, however, only $(1, 0)$ became the converging limit of CFR algorithm.

For example, player A and B at the information set have thrown cards $(1, 3)$ and $(2, 3)$ respectively, the strategy profile approaching to the equilibrium point $(0.825, 0.175)$, $(0.318, 0.682)$ at first, but since it is not a CFR-stable saddle, the CFR iteration process shift away and headed towards $(1, 0)$, $(0, 1)$ instead. The numerical error terms mainly come from our algorithm accumulating regret at the very beginning of the game is played, they should shrink to 0 asymptotically when T approaches infinity.

Conclusion and Future Works

In this paper, we introduced a novel variant of CFR that strategy itself converges to the Nash Equilibrium. In order to make regret-matching ideas fit into the discontinuity of GS-IIG payoffs, regret-grow CFR adopts a partial-forget mechanism to update the cumulative regrets, yet guarantees the cumulative regret vector's element sum increasing. We tested regret-grow CFR on non-zero sum Goofspiel and its subgames, which showed the exploitability values for both players are descending and in the magnitude of $O(1/\sqrt{T})$. Since the regret-grow CFR does not require averaging the previous strategies, dominated strategies can be ruled out, which makes probability distribution on information sets a sparse vector most of the time.

In further research, we would like to examine how such sparsity of payoff can facilitate the algorithms' convergence. Such analysis could offer further algorithmic or theoretical improvements on algorithm stability. For example, it is promising yet challenging to adapt regret-grow CFR with the Monte-Carlo sampling traversing method. If the non-deterministic character does not affect the stability of regret-grow CFR significantly, it would be interesting to compare the approach, in terms of the payoff of both players and exploitability values.

References

- Avis, D.; Rosenberg, G. D.; Savani, R.; and von Stengel, B. 2009. Enumeration of nash equilibria for two-player games. *Economic Theory*, 42(1): 9–37.
- 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 *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS '15*, 7–15. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450334136.
- Brown, N.; and Sandholm, T. 2018. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374): 418–424.
- Brown, N.; and Sandholm, T. 2019. Superhuman AI for multiplayer poker. *Science*, 365(6456): 885–890.
- Chen, X.; Deng, X.; and Teng, S.-H. 2009. Settling the complexity of computing two-player Nash Equilibria. *Journal of the ACM*, 56(3): 1–57.
- Gutierrez Julian, P. G.; and Michael, W. 2000. Imperfect information in reactive modules games. *Association for the Advancement of Artificial Intelligence*.
- Jiang, A. X.; and Leyton-Brown, K. 2011. Polynomial-time computation of exact correlated equilibrium in compact games. *Proceedings of the 12th ACM conference on Electronic commerce - EC '11*.
- Kroer, C.; Waugh, K.; Kiliç-Karzan, F.; and Sandholm, T. 2015. Faster First-Order Methods for Extensive-Form Game Solving. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15*, 817–834. New York, NY, USA: Association for Computing Machinery. ISBN 9781450334105.
- Marc, L.; Kevin, W.; Martin, Z.; and H, B. M. 2009. Monte Carlo Sampling for Regret Minimization in Extensive Games. In *23th Conference on Neural Information Processing Systems*, 1078–1086. Vancouver, British Columbia, Canada.
- Martin, S.; Neil, B.; Marc, L.; Matej, M.; Rudolf, K.; and Michael, B. 2019. Variance reduction in monte carlo counterfactual regret minimization (VR-MCCFR) for extensive form games using baselines. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2157–2164. Honolulu, Hawaii USA: PKP.
- Martin, Z.; Michael, J.; Michael, B.; and Carmelo, P. 2007. Regret minimization in games with incomplete information. *Advances in neural information processing systems*, 20: 1729–1736.
- Nash, J. F. 1950. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1): 48–49.
- Noam, B.; Adam, L.; Sam, G.; and Tuomas, S. 2019. Deep counterfactual regret minimization. In *International conference on machine learning*, 793–802. Long Beach, Calif, USA: International Conference on Machine Learning (ICML).
- Schrittwieser, J.; Antonoglou, I.; Hubert, T.; Simonyan, K.; Sifre, L.; Schmitt, S.; Guez, A.; Lockhart, E.; Hassabis, D.; Graepel, T.; et al. 2020. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839): 604–609.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587): 484–489.
- Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; et al. 2018. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419): 1140–1144.
- Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; et al. 2017. Mastering the game of go without human knowledge. *nature*, 550(7676): 354–359.
- Song, R.; Wei, Q.; Zhang, H.; and Lewis, F. L. 2019. Discrete-time non-zero-sum games with completely unknown dynamics. *IEEE Transactions on Cybernetics*, 51(6): 2929–2943.
- Tammelin, O. 2014. Solving large imperfect information games using CFR+. *arXiv preprint arXiv:1407.5042*.