

Lecture 5

Regret circuits and the Counterfactual Regret Minimization (CFR) paradigm

Instructor: Gabriele Farina*

In Lecture 4 we studied online optimization for *simplex* domains. In this section we show that it is possible to construct an efficient regret minimizer that outputs sequence-form strategies, by combining *any* local regret minimizers at each decision point of the tree-form sequential decision making, including RM and RM⁺ discussed before. The resulting algorithm has been known since the seminal work by [Zinkevich et al. \[2007\]](#) under the name *counterfactual regret minimization (CFR)*.

The analysis we present is based on the formalism of *regret circuits* [[Farina et al., 2019](#)].

1 Regret circuits

Regret circuits are a constructive answer to the question: “*given regret minimizers for two sets \mathcal{X}, \mathcal{Y} , can we compose these regret minimizers for various operations performed on the sets (e.g., intersection, convex hull, Cartesian product), in order to arrive at a regret minimizer for the resulting composite set?*”. In particular, in the following subsections, we will show that once regret minimizers for two sets \mathcal{X} and \mathcal{Y} are known, they can always be efficiently combined to come up with composite regret minimizers for the Cartesian product $\mathcal{X} \times \mathcal{Y}$ and the convex hull $\text{co}\{\mathcal{X}, \mathcal{Y}\}$.

Since, as we say in Lecture 2, any sequence-form strategy space Q can be defined inductively starting from convex hulls and Cartesian products, we can keep combining regret circuits until we construct an algorithm for Q . The resulting family of algorithms is known as counterfactual regret minimization (CFR).

1.1 Regret circuit for Cartesian product

Let $\mathcal{X} \subseteq \mathbb{R}^m$ and $\mathcal{Y} \subseteq \mathbb{R}^n$ be two sets, and let $\mathcal{R}_{\mathcal{X}}$ and $\mathcal{R}_{\mathcal{Y}}$ be regret minimizers for \mathcal{X} and \mathcal{Y} respectively. We can combine $\mathcal{R}_{\mathcal{X}}$ and $\mathcal{R}_{\mathcal{Y}}$ to form a regret minimizer for the Cartesian product $\mathcal{X} \times \mathcal{Y} \subseteq \mathbb{R}^{m+n}$ by doing the following:

- To output the next strategy, we ask $\mathcal{R}_{\mathcal{X}}$ and $\mathcal{R}_{\mathcal{Y}}$ for their next strategies—call them \mathbf{x}^t and \mathbf{y}^t —and return the pair $(\mathbf{x}^t, \mathbf{y}^t) \in \mathcal{X} \times \mathcal{Y}$.
- Any utility vector $\ell \in \mathbb{R}^{m+n}$ can be written as $\ell = (\ell_{\mathcal{X}}, \ell_{\mathcal{Y}}) \in \mathbb{R}^m \times \mathbb{R}^n$. When at time t we receive the utility vector $\ell^t = (\ell_{\mathcal{X}}^t, \ell_{\mathcal{Y}}^t)$ we forward $\ell_{\mathcal{X}}^t$ to $\mathcal{R}_{\mathcal{X}}$ and $\ell_{\mathcal{Y}}^t$ to $\mathcal{R}_{\mathcal{Y}}$.

The algorithm we just described is summarized in Algorithm 1 and given pictorially in Figure 1.

*Computer Science Department, Carnegie Mellon University. ✉ gfarina@cs.cmu.edu.

Algorithm 1: Regret minimizer for $\mathcal{X} \times \mathcal{Y}$

```
1 function NEXTSTRATEGY()  
2    $\mathbf{x}^t \leftarrow \mathcal{R}_{\mathcal{X}}.\text{NEXTSTRATEGY}()$   
3    $\mathbf{y}^t \leftarrow \mathcal{R}_{\mathcal{Y}}.\text{NEXTSTRATEGY}()$   
4   return  $(\mathbf{x}^t, \mathbf{y}^t) \in \mathcal{X} \times \mathcal{Y}$   
  
5 function OBSERVEUTILITY( $\ell^t = (\ell_{\mathcal{X}}^t, \ell_{\mathcal{Y}}^t) \in \mathbb{R}^m \times \mathbb{R}^n$ )  
6    $\mathcal{R}_{\mathcal{X}}.\text{OBSERVEUTILITY}(\ell_{\mathcal{X}}^t)$   
7    $\mathcal{R}_{\mathcal{Y}}.\text{OBSERVEUTILITY}(\ell_{\mathcal{Y}}^t)$ 
```




Figure 1: Regret circuit for the Cartesian product $\mathcal{X} \times \mathcal{Y}$.

Analysis The regret minimizers $\mathcal{R}_{\mathcal{X}}$ and $\mathcal{R}_{\mathcal{Y}}$ output strategies $\mathbf{x}^t, \mathbf{y}^t$ and observe utilities $\ell_{\mathcal{X}}^t, \ell_{\mathcal{Y}}^t$ at all time t . Hence, the regret cumulated by $\mathcal{R}_{\mathcal{X}}$ and $\mathcal{R}_{\mathcal{Y}}$ up to time T is given by

$$R_{\mathcal{X}}^T = \max_{\hat{\mathbf{x}} \in \mathcal{X}} \left\{ \sum_{t=1}^T (\ell_{\mathcal{X}}^t)^\top \hat{\mathbf{x}} - (\ell_{\mathcal{X}}^t)^\top \mathbf{x}^t \right\}$$
$$R_{\mathcal{Y}}^T = \max_{\hat{\mathbf{y}} \in \mathcal{Y}} \left\{ \sum_{t=1}^T (\ell_{\mathcal{Y}}^t)^\top \hat{\mathbf{y}} - (\ell_{\mathcal{Y}}^t)^\top \mathbf{y}^t \right\}.$$

In order to verify that Algorithm 1 indeed yields a regret minimizer we study the regret cumulated by Algorithm 1 up to any time T . Starting from the definition of regret, we have

$$R_{\mathcal{X} \times \mathcal{Y}}^T = \max_{(\hat{\mathbf{x}}, \hat{\mathbf{y}}) \in \mathcal{X} \times \mathcal{Y}} \left\{ \sum_{t=1}^T (\ell_{\mathcal{X}}^t, \ell_{\mathcal{Y}}^t)^\top (\hat{\mathbf{x}}, \hat{\mathbf{y}}) - (\ell_{\mathcal{X}}^t, \ell_{\mathcal{Y}}^t)^\top (\mathbf{x}^t, \mathbf{y}^t) \right\}$$
$$= \max_{\hat{\mathbf{x}} \in \mathcal{X}} \left\{ \sum_{t=1}^T (\ell_{\mathcal{X}}^t)^\top \hat{\mathbf{x}} - \sum_{t=1}^T (\ell_{\mathcal{X}}^t)^\top \mathbf{x}^t \right\} + \min_{\hat{\mathbf{y}} \in \mathcal{Y}} \left\{ \sum_{t=1}^T (\ell_{\mathcal{Y}}^t)^\top \hat{\mathbf{y}} - \sum_{t=1}^T (\ell_{\mathcal{Y}}^t)^\top \mathbf{y}^t \right\}$$
$$= R_{\mathcal{X}}^T + R_{\mathcal{Y}}^T.$$

So, if $\mathcal{R}_{\mathcal{X}}$ and $\mathcal{R}_{\mathcal{Y}}$ guarantee sublinear regret, then so does Algorithm 1. In other words, it is possible to minimize regret on $\mathcal{X} \times \mathcal{Y}$ by simply minimizing it on \mathcal{X} and \mathcal{Y} independently and then combining the decisions.

In summary, we have proven the following result.

Theorem 1.1. The regret circuit of Figure 1 provides a regret minimizer for the Cartesian product $\mathcal{X} \times \mathcal{Y}$ that satisfies the regret bound

$$R_{\mathcal{X} \times \mathcal{Y}}^T = R_{\mathcal{X}}^T + R_{\mathcal{Y}}^T.$$

The extension to the Cartesian product of more than two sets is direct.

1.2 Regret circuit for convex hull

We now show how to combine $\mathcal{R}_{\mathcal{X}}$ and $\mathcal{R}_{\mathcal{Y}}$ for two sets $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^n$ to construct a regret minimizer for the convex hull $\text{co}\{\mathcal{X}, \mathcal{Y}\}$. The construction is harder than in the case of Cartesian products, but some of the ideas carry over. Then, we will use a third regret minimizer \mathcal{R}_{Δ} for the 2-simplex Δ^2 decide how to “mix” \mathbf{x}^t and \mathbf{y}^t . More precisely, we will do the following.

- To output the next strategy, we ask $\mathcal{R}_{\mathcal{X}}$ and $\mathcal{R}_{\mathcal{Y}}$ for their next strategies—call them \mathbf{x}^t and \mathbf{y}^t . Then, we will ask \mathcal{R}_{Δ} for its next strategy $\boldsymbol{\lambda}^t = (\lambda_1^t, \lambda_2^t) \in \Delta^2$, and return the convex combination $\lambda_1^t \mathbf{x}^t + \lambda_2^t \mathbf{y}^t \in \text{co}\{\mathcal{X}, \mathcal{Y}\}$.

- When at time t we receive the utility vector $\ell^t \in \mathbb{R}^n$, we forward ℓ^t to \mathcal{R}_X and \mathcal{R}_Y . Then, we forward the utility vector

$$\ell_\Delta^t := \begin{pmatrix} (\ell^t)^\top \mathbf{x}^t \\ (\ell^t)^\top \mathbf{y}^t \end{pmatrix}. \quad (1)$$

to \mathcal{R}_Δ .

The algorithm we just described is summarized in Algorithm 2 and given pictorially in Figure 2.

Algorithm 2: Regret minimizer for $\text{co}\{\mathcal{X}, \mathcal{Y}\}$

```

1 function NEXTSTRATEGY()
2    $\mathbf{x}^t \leftarrow \mathcal{R}_X.\text{NEXTSTRATEGY}()$ 
3    $\mathbf{y}^t \leftarrow \mathcal{R}_Y.\text{NEXTSTRATEGY}()$ 
4    $\lambda^t = (\lambda_1^t, \lambda_2^t) \in \Delta^2 \leftarrow \mathcal{R}_\Delta.\text{NEXTSTRATEGY}()$ 
5   return  $\lambda_1^t \mathbf{x}^t + \lambda_2^t \mathbf{y}^t \in \text{co}\{\mathcal{X}, \mathcal{Y}\}$ 

```

```

6 function OBSERVEUTILITY( $\ell^t$ )
7    $\mathcal{R}_X.\text{OBSERVEUTILITY}(\ell^t)$ 
8    $\mathcal{R}_Y.\text{OBSERVEUTILITY}(\ell^t)$ 
9    $\mathcal{R}_\Delta.\text{OBSERVEUTILITY}(\ell_\Delta^t := \begin{pmatrix} (\ell^t)^\top \mathbf{x}^t \\ (\ell^t)^\top \mathbf{y}^t \end{pmatrix})$ 

```




Figure 2: Regret circuit for the convex hull $\text{co}\{\mathcal{X}, \mathcal{Y}\}$. The utility vector ℓ_Δ^t is defined in Equation (1).

Analysis The regret minimizers \mathcal{R}_X and \mathcal{R}_Y output strategies $\mathbf{x}^t, \mathbf{y}^t$ and observe utilities ℓ^t at all time t . Hence, the regret cumulated by \mathcal{R}_X and \mathcal{R}_Y up to time T is given by

$$R_X^T = \max_{\hat{\mathbf{x}} \in \mathcal{X}} \left\{ \sum_{t=1}^T (\ell^t)^\top \hat{\mathbf{x}} - (\ell^t)^\top \mathbf{x}^t \right\} \quad (2)$$

$$R_Y^T = \max_{\hat{\mathbf{y}} \in \mathcal{Y}} \left\{ \sum_{t=1}^T (\ell^t)^\top \hat{\mathbf{y}} - (\ell^t)^\top \mathbf{y}^t \right\}. \quad (3)$$

The regret minimizer \mathcal{R}_Δ outputs strategies $\lambda^t = (\lambda_1^t, \lambda_2^t)$ at all times t and observes utility vectors ℓ_Δ^t at all times t . So, the regret cumulated by \mathcal{R}_Δ is given by

$$R_\Delta^T := \max_{\hat{\lambda} \in \Delta^2} \left\{ \left(\sum_{t=1}^T \hat{\lambda}_1 (\ell^t)^\top \mathbf{x}^t + \hat{\lambda}_2 (\ell^t)^\top \mathbf{y}^t \right) \right\} - \left(\sum_{t=1}^T \lambda_1^t (\ell^t)^\top \mathbf{x}^t + \lambda_2^t (\ell^t)^\top \mathbf{y}^t \right).$$

In order to verify that Algorithm 2 indeed yields a regret minimizer we study the regret it cumulates up to any time T . Starting from the definition of regret, we have

$$\begin{aligned} R^T &= \max_{\substack{\hat{\lambda} \in \Delta^2 \\ \hat{\mathbf{x}} \in \mathcal{X}, \hat{\mathbf{y}} \in \mathcal{Y}}} \left\{ \sum_{t=1}^T (\ell^t)^\top (\hat{\lambda}_1 \hat{\mathbf{x}} + \hat{\lambda}_2 \hat{\mathbf{y}}) - (\ell^t)^\top (\lambda_1^t \mathbf{x}^t + \lambda_2^t \mathbf{y}^t) \right\} \\ &= \max_{\substack{\hat{\lambda} \in \Delta^2 \\ \hat{\mathbf{x}} \in \mathcal{X}, \hat{\mathbf{y}} \in \mathcal{Y}}} \left\{ \hat{\lambda}_1 \sum_{t=1}^T (\ell^t)^\top \hat{\mathbf{x}} + \hat{\lambda}_2 \sum_{t=1}^T (\ell^t)^\top \hat{\mathbf{y}} \right\} - \left(\sum_{t=1}^T \lambda_1^t (\ell^t)^\top \mathbf{x}^t + \lambda_2^t (\ell^t)^\top \mathbf{y}^t \right). \end{aligned} \quad (4)$$

Now, we make two important observations. First,

$$\max_{\substack{\hat{\lambda} \in \Delta^2 \\ \hat{\mathbf{x}} \in \mathcal{X}, \hat{\mathbf{y}} \in \mathcal{Y}}} \left\{ \hat{\lambda}_1 \sum_{t=1}^T (\ell^t)^\top \hat{\mathbf{x}} + \hat{\lambda}_2 \sum_{t=1}^T (\ell^t)^\top \hat{\mathbf{y}} \right\} = \max_{\hat{\lambda} \in \Delta^2} \left\{ \hat{\lambda}_1 \max_{\hat{\mathbf{x}} \in \mathcal{X}} \left\{ \sum_{t=1}^T (\ell^t)^\top \hat{\mathbf{x}} \right\} + \hat{\lambda}_2 \max_{\hat{\mathbf{y}} \in \mathcal{Y}} \left\{ \sum_{t=1}^T (\ell^t)^\top \hat{\mathbf{y}} \right\} \right\}, \quad (5)$$

since all components of $\hat{\lambda} = (\hat{\lambda}_1, \hat{\lambda}_2) \in \Delta^2$ are non-negative. Second, the inner minimization problem over \mathcal{X} is related to the cumulative regret $R_{\mathcal{X}}^T$ (2) of the regret minimizer $\mathcal{R}_{\mathcal{X}}$ that observes the utility functions ℓ^t as follows:

$$\max_{\hat{\mathbf{x}} \in \mathcal{X}} \left\{ \sum_{t=1}^T (\ell^t)^\top \hat{\mathbf{x}} \right\} = R_{\mathcal{X}}^T + \sum_{t=1}^T (\ell^t)^\top \mathbf{x}^t. \quad (6)$$

Similarly, for \mathcal{Y} we have from (3) that

$$\max_{\hat{\mathbf{y}} \in \mathcal{Y}} \left\{ \sum_{t=1}^T (\ell^t)^\top \hat{\mathbf{y}} \right\} = R_{\mathcal{Y}}^T + \sum_{t=1}^T (\ell^t)^\top \mathbf{y}^t. \quad (7)$$

Plugging Equations (5), (6), and (7) into (4), we obtain

$$\begin{aligned} R^T &= \max_{\hat{\lambda} \in \Delta^2} \left\{ \hat{\lambda}_1 \max_{\hat{\mathbf{x}} \in \mathcal{X}} \left\{ \sum_{t=1}^T (\ell^t)^\top \hat{\mathbf{x}} \right\} + \hat{\lambda}_2 \max_{\hat{\mathbf{y}} \in \mathcal{Y}} \left\{ \sum_{t=1}^T (\ell^t)^\top \hat{\mathbf{y}} \right\} \right\} - \left(\sum_{t=1}^T \lambda_1^t (\ell^t)^\top \mathbf{x}^t + \lambda_2^t (\ell^t)^\top \mathbf{y}^t \right) \\ &= \max_{\hat{\lambda} \in \Delta^2} \left\{ \hat{\lambda}_1 \left(R_{\mathcal{X}}^T + \sum_{t=1}^T (\ell^t)^\top \mathbf{x}^t \right) + \hat{\lambda}_2 \left(R_{\mathcal{Y}}^T + \sum_{t=1}^T (\ell^t)^\top \mathbf{y}^t \right) \right\} - \left(\sum_{t=1}^T \lambda_1^t (\ell^t)^\top \mathbf{x}^t + \lambda_2^t (\ell^t)^\top \mathbf{y}^t \right) \\ &= \max_{\hat{\lambda} \in \Delta^2} \left\{ \left(\sum_{t=1}^T \hat{\lambda}_1 (\ell^t)^\top \mathbf{x}^t + \hat{\lambda}_2 (\ell^t)^\top \mathbf{y}^t \right) + (\hat{\lambda}_1 R_{\mathcal{X}}^T + \hat{\lambda}_2 R_{\mathcal{Y}}^T) \right\} - \left(\sum_{t=1}^T \lambda_1^t (\ell^t)^\top \mathbf{x}^t + \lambda_2^t (\ell^t)^\top \mathbf{y}^t \right). \end{aligned}$$

No matter the choice of $(\hat{\lambda}_1, \hat{\lambda}_2) \in \Delta^2$, the convex combination $\hat{\lambda}_1 R_{\mathcal{X}}^T + \hat{\lambda}_2 R_{\mathcal{Y}}^T$ satisfies

$$\hat{\lambda}_1 R_{\mathcal{X}}^T + \hat{\lambda}_2 R_{\mathcal{Y}}^T \leq \max\{R_{\mathcal{X}}^T, R_{\mathcal{Y}}^T\}.$$

Hence, we can write

$$\begin{aligned} R^T &\leq \max_{\hat{\lambda} \in \Delta^2} \left\{ \left(\sum_{t=1}^T \hat{\lambda}_1 (\ell^t)^\top \mathbf{x}^t + \hat{\lambda}_2 (\ell^t)^\top \mathbf{y}^t \right) + \max\{R_{\mathcal{X}}^T, R_{\mathcal{Y}}^T\} \right\} - \left(\sum_{t=1}^T \lambda_1^t (\ell^t)^\top \mathbf{x}^t + \lambda_2^t (\ell^t)^\top \mathbf{y}^t \right) \\ &= \max_{\hat{\lambda} \in \Delta^2} \left\{ \left(\sum_{t=1}^T \hat{\lambda}_1 (\ell^t)^\top \mathbf{x}^t + \hat{\lambda}_2 (\ell^t)^\top \mathbf{y}^t \right) \right\} + \max\{R_{\mathcal{X}}^T, R_{\mathcal{Y}}^T\} - \left(\sum_{t=1}^T \lambda_1^t (\ell^t)^\top \mathbf{x}^t + \lambda_2^t (\ell^t)^\top \mathbf{y}^t \right) \\ &= R_{\Delta}^T + \max\{R_{\mathcal{X}}^T, R_{\mathcal{Y}}^T\}. \end{aligned} \quad (8)$$

Equation (8) guarantees that if all three regrets $\{R_{\Delta}^T, R_{\mathcal{X}}^T, R_{\mathcal{Y}}^T\}$ grow sublinearly, then so does R^T . Figure 2 shows the regret circuit that corresponds to our construction above.

In summary, we have proven the following result.

Theorem 1.2. The regret circuit of Figure 2 provides a regret minimizer for the convex hull $\text{co}\{\mathcal{X}, \mathcal{Y}\}$ that satisfies the regret bound

$$R^T \leq R_{\Delta^2}^T + \max\{R_{\mathcal{X}}^T, R_{\mathcal{Y}}^T\}.$$

Extension to multiple sets The construction shown in Algorithm 2 and Figure 2 can be readily extended to handle the convex hull $\text{co}\{\mathcal{X}_1, \dots, \mathcal{X}_n\}$ of n sets as follows.

- At each time instant t , we let n regret minimizers $\mathcal{R}_{\mathcal{X}_1}, \dots, \mathcal{R}_{\mathcal{X}_n}$ (for $\mathcal{X}_1, \dots, \mathcal{X}_n$, respectively) output their next strategies $\mathbf{x}_1^t, \dots, \mathbf{x}_n^t$. Those strategies are combined according to the convex combination coefficients λ^t output by the additional regret minimizer \mathcal{R}_{Δ} for the n -simplex Δ^n to form the convex combination strategy $\lambda_1^t \mathbf{x}_1^t + \dots + \lambda_n^t \mathbf{x}_n^t \in \text{co}\{\mathcal{X}_1, \dots, \mathcal{X}_n\}$.

- The utility vector ℓ^t is fed into all the regret minimizers $\mathcal{R}_{\mathcal{X}_i}$ ($i = 1, \dots, n$). Then, the utility vector ℓ_{Δ}^t , defined as

$$\ell_{\Delta}^t = \begin{pmatrix} (\ell^t)^\top \mathbf{x}_1^t \\ \vdots \\ (\ell^t)^\top \mathbf{x}_n^t \end{pmatrix},$$

is input into a regret minimizer for the n -dimensional simplex Δ^n .

With the same argument used earlier, one can easily show that the regret bound in that case is

$$R^T \leq R_{\Delta^n}^T + \max\{R_{\mathcal{X}_1}^T, \dots, R_{\mathcal{X}_n}^T\}.$$

2 Counterfactual Regret Minimization

We have seen in Lecture 2 that a sequence-form strategy space can be characterized recursively by composing convex hulls and Cartesian products operations. By applying the regret circuits described above, we can then construct a regret minimizers for any sequence-form strategy space. The resulting regret minimizer is called CFR. In a nutshell, CFR decomposes the problem of minimizing regret on the whole tree-form decision problem into local regret minimization problems at each of the individual decision points $j \in \mathcal{J}$. Any regret minimizer \mathcal{R}_j for simplex domains can be used to solve the local regret minimization problems. Popular options are the regret matching algorithm, and the regret matching plus algorithm (Lecture 4).

Before giving pseudocode, we recall and introduce a bit of notation to deal with tree-form sequential decision processes.

Notation for tree-form sequential decision processes We use the following notation for dealing with tree-form sequential decision processes (TFSDPs), most of which was already introduced in Lecture 2. The notation is also summarized in Table 1.

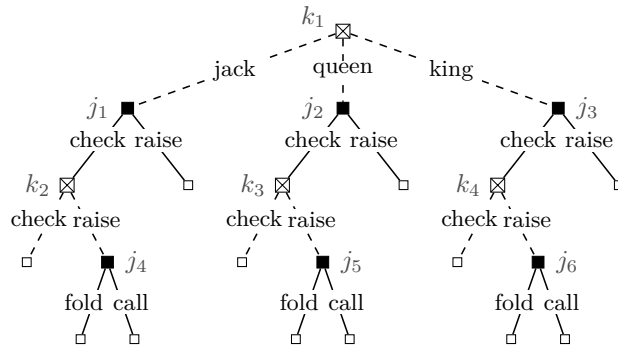


Figure 3: Tree-form sequential decision making process of the first acting player in the game of Kuhn poker.

- We denote the set of decision points in the TFSDP as \mathcal{J} , and the set of observation points as \mathcal{K} . At each decision point $j \in \mathcal{J}$, the agent selects an action from the set A_j of available actions. At each observation point $k \in \mathcal{K}$, the agent observes a signal s_k from the environment out of a set of possible signals S_k .
- **(new!)** We denote by ρ the transition function of the process. Picking action $a \in A_j$ at decision point $j \in \mathcal{J}$ results in the process transitioning to $\rho(j, a) \in \mathcal{J} \cup \mathcal{K} \cup \{\perp\}$, where \perp denotes the end of the decision process. Similarly, the process transitions to $\rho(k, s) \in \mathcal{J} \cup \mathcal{K} \cup \{\perp\}$ after the agent observes signal $s \in S_k$ at observation point $k \in \mathcal{K}$.
- A pair (j, a) where $j \in \mathcal{J}$ and $a \in A_j$ is called a *sequence*. The set of all sequences is denoted as $\Sigma := \{(j, a) : j \in \mathcal{J}, a \in A_j\}$. For notational convenience, we will often denote an element (j, a) in Σ as ja without using parentheses.

- Given a decision point $j \in \mathcal{J}$, we denote by p_j its *parent sequence*, defined as the last sequence (that is, decision point-action pair) encountered on the path from the root of the decision process to j . If the agent does not act before j (that is, j is the root of the process or only observation points are encountered on the path from the root to j), we let $p_j = \emptyset$.

Symbol	Description
\mathcal{J}	Set of decision points
A_j	Set of legal actions at decision point $j \in \mathcal{J}$
\mathcal{K}	Set of observation points
S_k	Set of possible signals at observation point $k \in \mathcal{K}$
ρ	Transition function: <ul style="list-style-type: none"> • given $j \in \mathcal{J}$ and $a \in A_j$, $\rho(j, a)$ returns the next point $v \in \mathcal{J} \cup \mathcal{K}$ in the decision tree that is reached after selecting legal action a in j, or \perp if the decision process ends; • given $k \in \mathcal{K}$ and $s \in S_k$, $\rho(k, s)$ returns the next point $v \in \mathcal{J} \cup \mathcal{K}$ in the decision tree that is reached after observing signal s in k, or \perp if the decision process ends
Σ	Set of sequences, defined as $\Sigma := \{(j, a) : j \in \mathcal{J}, a \in A_j\}$
p_j	Parent sequence of decision point $j \in \mathcal{J}$, defined as the last sequence (decision point-action pair) on the path from the root of the TFSDP to decision point j ; if the agent does not act before j , $p_j = \emptyset$

Table 1: Summary of notation in TFSDPs.

As an example, consider the TFSDP faced by Player 1 in the game of Kuhn poker [Kuhn, 1950], depicted in Figure 3, which we already introduced in Lecture 2. There, we have that $\mathcal{J} = \{j_1, \dots, j_6\}$ and $\mathcal{K} = \{k_1, \dots, k_4\}$. $A_{j_1} = S_{k_4} = \{\text{check}, \text{raise}\}$. $A_{j_5} = \{\text{fold}, \text{call}\}$. $S_{k_1} = \{\text{jack}, \text{queen}, \text{king}\}$. $\rho(k_3, \text{check}) = \rho(j_2, \text{raise}) = \perp$. $\rho(k_1, \text{king}) = j_3$. $\rho(j_2, \text{check}) = k_3$. $p_{j_4} = (j_1, \text{check})$. $p_{j_6} = (j_3, \text{check})$. $p_{j_1} = p_{j_2} = p_{j_3} = \emptyset$.

Notation for the components of vectors Any vector $\mathbf{x} \in \mathbb{R}^{|\Sigma|}$ has, by definition, as many components as sequences Σ . The component corresponding to a specific sequence $ja \in \Sigma$ is denoted as $\mathbf{x}[ja]$. Similarly, given any decision point $j \in \mathcal{J}$, any vector $\mathbf{x} \in \mathbb{R}^{|A_j|}$ has as many components as the number of actions at j . The component corresponding to a specific action $a \in A_j$ is denoted $\mathbf{x}[a]$.

CFR algorithm Pseudocode for CFR is given in Algorithm 3. Note that the implementation is parametric on the regret minimization algorithms \mathcal{R}_j run locally at each decision point. Any regret minimizer \mathcal{R}_j for simplex domains can be used to solve the local regret minimization problems. Popular options are the regret matching algorithm, and the regret matching plus algorithm (Lecture 4).

It can be shown that the regret cumulated by the CFR algorithm satisfies the following bound.

Proposition 2.1. Let R_j^T ($j \in \mathcal{J}$) denote the regret cumulated up to time T by each of the regret minimizers \mathcal{R}_j . Then, the regret R^T cumulated by Algorithm 3 up to time T satisfies

$$R^T \leq \sum_{j \in \mathcal{J}} \max\{0, R_j^T\}.$$

It is then immediate to see that if each R_j^T grows sublinearly in T , then so does R^T .

Algorithm 3: CFR regret minimizer

Data: \mathcal{R}_j , one regret minimizer for $\Delta^{|A_j|}$; one for each decision point $j \in \mathcal{J}$ of the TFSDP.

```
1 function NEXTSTRATEGY()
  [▷ Step 1: we ask each of the  $\mathcal{R}_j$  for their next strategy local at each decision point]
2  for each decision point  $j \in \mathcal{J}$  do
3    |  $\mathbf{b}_j^t \in \Delta^{|A_j|} \leftarrow \mathcal{R}_j.\text{NEXTSTRATEGY}()$ 
  [▷ Step 2: we construct the sequence-form representation of the strategy that plays according to
  the distribution  $\mathbf{b}_j^t$  at each decision point  $j \in \mathcal{J}$ ]
4   $\mathbf{x}^t = \mathbf{0} \in \mathbb{R}^{|\Sigma|}$ 
5  for each decision point  $j \in \mathcal{J}$  in top-down traversal order in the TFSDP do
6    | for each action  $a \in A_j$  do
7      | | if  $p_j = \emptyset$  then
8        | | |  $\mathbf{x}^t[ja] \leftarrow \mathbf{b}_j^t[a]$ 
9      | | else
10     | | |  $\mathbf{x}^t[ja] \leftarrow \mathbf{x}^t[p_j] \cdot \mathbf{b}_j^t[a]$ 
  [▷ You should convince yourself that the vector  $\mathbf{x}^t$  we just filled in above is a valid sequence-form
  strategy, that is, it satisfies the required consistency constraints we saw in Lecture 2. In symbols,
   $\mathbf{x}^t \in Q$ ]
11  return  $\mathbf{x}^t$ 
```

```
12 function OBSERVEUTILITY( $\ell^t \in \mathbb{R}^{|\Sigma|}$ )
  [▷ Step 1: we compute the expected utility for each subtree rooted at each node  $v \in \mathcal{J} \cup \mathcal{K}$ ]
13   $V^t \leftarrow$  empty dictionary [▷ eventually, it will map keys  $\mathcal{J} \cup \mathcal{K} \cup \{\perp\}$  to real numbers]
14   $V^t[\perp] \leftarrow 0$ 
15  for each node in the tree  $v \in \mathcal{J} \cup \mathcal{K}$  in bottom-up traversal order in the TFSDP do
16    | if  $v \in \mathcal{J}$  then
17    | | Let  $j = v$ 
18    | |  $V^t[j] \leftarrow \sum_{a \in A_j} \mathbf{b}_j^t[a] \cdot (\ell^t[ja] + V^t[\rho(j, a)])$ 
19    | | else
20    | | Let  $k = v$ 
21    | |  $V^t[k] \leftarrow \sum_{s \in S_k} V^t[\rho(k, s)]$ 
  [▷ Step 2: at each decision point  $j \in \mathcal{J}$ , we now construct a local utility vector  $\ell_j^t$  called
  counterfactual utility]
22  for each decision point  $j \in \mathcal{J}$  do
23    |  $\ell_j^t \leftarrow \mathbf{0} \in \mathbb{R}^{|A_j|}$ 
24    | for each action  $a \in A_j$  do
25    | |  $\ell_j^t[a] \leftarrow \ell^t[ja] + V^t[\rho(j, a)]$ 
26    | |  $\mathcal{R}_j.\text{OBSERVEUTILITY}(\ell_j^t)$ 
```

References

Martin Zinkevich, Michael Bowling, Michael Johanson, and Carmelo Piccione. Regret minimization in games with incomplete information. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*, 2007.

Gabriele Farina, Christian Kroer, and Tuomas Sandholm. Regret circuits: Composability of regret minimizers. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019.

H. W. Kuhn. A simplified two-person poker. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, volume 1 of *Annals of Mathematics Studies*, 24, pages 97–103. Princeton University Press, Princeton, New Jersey, 1950.