

## Lecture 2: October 30

## Online Convex Programming

*Lecturer: Yishay Mansour**Scribe: Assaf Yifrach, Ofek Kirzner*

## 2.1 Online Convex Programming Model

Online convex programming is a framework of online optimization problems. The model is defined as below:

**Input:** A convex set  $S$

At time  $t \in [1, 2, \dots]$ :

1. Receive input  $x_t$
2. Select a vector  $w_t \in S$
3. Receive a convex loss function  $f_t : S \rightarrow \mathfrak{R}$
4. Suffer loss  $f_t(w_t)$

**Total loss:**  $\sum_{t=1}^T f_t(w_t)$

### 2.1.1 Motivation

In online linear regression, we observe in time  $t$  the vector  $x_t$ . After choosing the weight vector  $w_t$ , we observe the result  $y_t$ . The loss we suffer, assuming quadratic error, is  $(w_t^T x_t - y_t)^2$ . So, we can think of the loss function as  $f_t(w) = (w^T x_t - y_t)^2$  which is indeed a convex function. The regret, as we defined it in Lecture 1, is:

$$\text{Regret} = \underbrace{\sum_{t=1}^T f_t(w_t)}_{\text{online solution}} - \underbrace{\min_{w \in S} \sum_{t=1}^T f_t(w)}_{\text{classic regression solution, knowing the entire input in advance}}$$

## 2.1.2 Regret Relative to Competing Vectors

Sometimes, we analyze the regret of an algorithm with respect to a different set than the input set  $S$ . For example, in linear regression we can compare to  $U = \{w : \|w\|_2 \leq 1\}$

**Definition** Regret with regard to a vector  $u \in U$  at time  $T$

$$\forall u \in U, \text{Regret}_T(u) = \sum_{t=1}^T f_t(w_t) - \sum_{t=1}^T f_t(u)$$

.

**Definition** Regret with regard to a set  $U$  at time  $T$

$$\text{Regret}_T(U) = \max_{u \in U} \text{Regret}_T(u)$$

.

## 2.2 Convexification

While it may be desirable to use convex optimization, the loss functions  $f_t$  might be non-convex, thus we introduce two techniques to adopt the non-convex loss function to the convex optimization model.

- Randomization
- Surrogate Loss

Given some online optimization problem, we solve it by translating it into an online convex optimization (OCC) problem and solving the OCC efficiently with some solver.

### 2.2.1 Convexification by Randomization

We explain the method using an example from Lecture 1: the experts problem, where there are  $d$  experts, each suffers some loss from  $[0, 1]$  in every discrete time  $t$ . We wish to design an online algorithm that suffers a loss not much worse than the loss of the best expert. Notations:

The experts  $\Delta = \{1, \dots, d\}$ .

The loss of expert  $i$  at time  $t$ :  $y_t[i] \in [0, 1]$

We referred to the experts problem as a special case of classification, where each expert  $i$  represents hypothesis  $h_i$  and the loss of an expert is the loss of its prediction:  $y_t[i] = l(h_i(x_t), y_t)$

If we limit the algorithm to choose a single expert in every step, the loss function will not be convex. In addition, as we saw in previous lecture, Cover's impossibility result guarantees that no algorithm can attain low regret for such scenario.

To address this issue, we let the online algorithm to randomize an expert, without letting the adversary know the randomization's realization. On time  $t$ , the algorithm selects a distribution:

$$S = \{w \in \mathfrak{R}^d : w \geq 0, \|w\|_1 = 1\}$$

Choosing a single expert  $p_t \in \Delta$  can be implemented using the selected distribution:

$$Pr[p_t = i] = w_t[i]$$

Now, the expected loss is:

$$E[y_t[p_t]] = \sum_{i=1}^d Pr[p_t = i]y_t[i] = w_t^T y_t$$

**Note:** We assume  $y_t$  is independent with  $p_t$ . While selecting  $y_t$ , the adversary might know the distribution  $w_t$ , but it does not know its realization  $p_t$ .

Now, define the loss function to be the expected loss  $f_t(w) := w^T y_t$ , which is convex.

The function  $f_t$  depends on  $y_t$ , yet it is consistent with our model, since the loss  $f_t$  can depend on  $y_t$  (but not the realization  $p_t$ ).

We would like to compare our algorithm to all the single experts. That is, to the set of unit vectors:  $U = \{(1, 0, \dots, 0), (0, 1, 0, \dots), \dots\}$ .

In this specific example, the total loss of a constant strategy  $w$  is  $\sum_{t=1}^T f_t(w)$  which is linear in  $w$ . Therefore the minimum is attained by some point  $u \in U$ . Which implies that  $Regret_T(U) = Reget_T(S)$

## 2.2.2 Convexification with Surrogate Loss

Similarly to using a hinge-loss as an upper bound for 0 – 1 loss, we upper bound the loss function with a convex function.

First, we define a convex set  $S$  of hypotheses. Then, we optimize a convex upper bound  $\bar{l} : S \rightarrow \mathfrak{R}$  over the original loss. Denote the original loss as  $l$ . We do not specify here in what sense  $\bar{l}$  is an upper bound of  $l$ . We require:

1.  $\bar{l} \geq l$  in some useful sense.
2.  $\bar{l}$  is convex
3.  $\bar{l} = l$  for the optimal solution (or at least the difference is small)

**Example: Online Classification with Finite Hypothesis Class**

We demonstrate surrogate loss on the problem of online classification with finite hypothesis class which we solved in previous lecture using the halving algorithm. We assumed a realizable model, where the adversary chooses some hypothesis  $h^* \in H$  and returns  $y_t = h^*(x_t)$ . Using the halving algorithm we proved an upper bound on the number of errors of  $\log_2(|H|)$ . We will get a similar result using convex programming with surrogate loss.

The hypothesis class:

$$H = \{h_1, \dots, h_d\}$$

The convex set  $S$ , is the distributions over  $H$ :

$$S = \{w \in \mathfrak{R}^d : w \geq 0, \|w\|_1 = 1\}$$

In time  $t$ , we observe the predictions of all hypothesis:

$$v_t = (h_1(x_t), \dots, h_d(x_t))$$

The algorithm predicts:

$$p_t = \begin{cases} 1 & \text{if } w_t^T v_t \geq \frac{1}{2} \\ 0 & \text{if } w_t^T v_t < \frac{1}{2} \end{cases}$$

Thus, one can think of  $S$  as an alternative convex class of hypotheses.

We mark the error timestamps  $M = \{t : p_t \neq y_t\}$ . And now, given the real outcomes  $y_1, \dots, y_t$  we define a convex function, with regard to  $w$ , that upper bounds the original loss  $|y_t - p_t|$ :

$$f_t(w) := \begin{cases} 2|w^T v_t - y_t| & \text{if } t \in M \\ 0 & \text{if } t \notin M \end{cases}$$

**Note:**  $f_t$  depends on  $M$  and thus on  $w_1, \dots, w_t$ . This is consistent with the online convex optimization model since  $f_t$  is given only after the algorithm chooses  $w_t$ .

The algorithm chooses  $w_t \in S$ . We do not specify how the algorithm selects  $w_t$  at time  $t$ , there are various ways to do that, we will address this issue later.

**Properties of  $f_t$  :**

1.  $f_t(w)$  is *convex*: if  $t \notin M$ ,  $f_t = 0$  is convex. Otherwise,  $f_t$  is convex as a composition of a convex function (absolute value) over an affine function  $w^T v_t - y_t$ .
2.  $f_t$  upper bounds the original loss: We show that  $f_t(w_t) \geq |p_t - y_t|$ .  
If  $t \notin M$ ,  $f_t(w_t) = 0 = |p_t - y_t|$ . Otherwise,  $t \in M$ . Since  $p_t \neq y_t$ , then  $|w_t^T v_t - y_t| \geq \frac{1}{2}$ .  
So  $f_t(w_t) = 2|w_t^T v_t - y_t| \geq 1 = |p_t - y_t|$ .

Next, we bound the regret of our solution. For that purpose, we state what the unspecified online convex programming algorithm guarantees.

Preliminary definition - Lipschitz constant is a bound on how fast a function  $f(w)$  can change.

**Definition** A function  $f : S \rightarrow \Re$  is  $L$ -Lipschitz with regard to some norm if:

$$\forall w_1, w_2 \in S : |f(w_1) - f(w_2)| \leq L_t \|w_1 - w_2\|$$

In the next section, we will show an algorithm for online convex programming which achieves the following bound:

$$\forall u \in S : \sum_{t=1}^T f_t(w_t) \leq \sum_{t=1}^T f_t(u) + \frac{\log(d)}{\eta} + 2\eta \sum_{t=1}^T L_t$$

where  $\eta$  is a parameter, and  $L_t$  is Lipschitz constant of  $f_t$  regarding  $L_1$  norm.

In order to use that bound, we calculate a Lipschitz constant for  $f_t$  defined above: If  $t \notin M$ , then  $|f_t(w_1) - f_t(w_2)| = 0$ , therefore  $L_t = 0$ . Otherwise,  $t \in M$ , we show that  $L_t = 1$ :

$$\begin{aligned} |f_t(w_1) - f_t(w_2)| &= |2|w_1^T v_t - y_t| - 2|w_2^T v_t - y_t|| \\ &\leq 2|(w_1 - w_2)^T v_t| && \text{triangle inequality} \\ &\leq 2 \sum_{i:w_1[i] \geq w_2[i]} w_1[i] - w_2[i] && \text{choosing } v_t \text{ which maximizes the expression} \\ &= 2 \frac{1}{2} \|w_1 - w_2\|_1 && \|w_1\|_1 = \|w_2\|_1 = 1 \end{aligned}$$

Using  $\eta = \frac{1}{4}$  we get:

$$\begin{aligned} |M| &\leq \sum_{t=1}^T f_t(w_t) && \text{surrogate loss - upper bound property} \\ &\leq \sum_{t=1}^T f_t(u) + 4\log(d) + \frac{1}{2} \sum_{t=1}^T L_t \end{aligned}$$

Since  $\sum_{t=1}^T L_t = |M|$ , we have

$$\frac{1}{2}|M| \leq \sum_{t=1}^T f_t(u) + 4\log(d)$$

$$|M| \leq 2 \sum_{t=1}^T f_t(u) + 8 \log(d)$$

This is true for every  $u \in S$ , and since we are in the realizable case, there exists a unit vector  $u \in S$  with  $\forall t : u^T v_t = y_t$ . Therefore  $f_t(u) = 2|u^T v_t - y_t| = 0$  also for  $t \in M$ . So we get  $\sum_{t=1}^T f_t(u) = 0$  and overall:

$$|M| \leq 8 \log(d)$$

## 2.3 Follow the Leader (FTL)

We introduce an algorithm for online convex programming that minimizes the retrospective loss. Notice the only missing piece in the convex programming model, is the specification of how to choose  $w_t$ .

**FTL**

$$\forall t : w_t = \arg \min_{w \in S} \sum_{i=1}^{t-1} f_i(w) \text{ (break ties arbitrarily)}$$

**Lemma 2.1** *Let  $w_1, w_2, \dots$  be the vectors produced by FTL, then  $\forall u \in S$ :*

$$\text{Regret}(u) = \sum_{t=1}^T (f_t(w_t) - f_t(u)) \leq \sum_{t=1}^T (f_t(w_t) - f_t(w_{t+1}))$$

**Proof:** We will show that:

$$\sum_{t=1}^T f_t(w_{t+1}) \leq \sum_{t=1}^T f_t(u)$$

The proof is by induction on  $T$ . For the basis of the induction, for  $T = 1 : f_1(w_2) \leq f_1(u)$  from FTL definition. Assume the inequality holds for  $T-1$ :

$$\forall u : \sum_{t=1}^{T-1} f_t(w_{t+1}) \leq \sum_{t=1}^{T-1} f_t(u)$$

Adding  $f_T(w_{T+1})$ :

$$\sum_{t=1}^{T-1} f_t(w_{t+1}) + f_T(w_{T+1}) \leq \sum_{t=1}^{T-1} f_t(u) + f_T(w_{T+1})$$

This is true for all  $u$ , specifically  $u = w_{T+1}$ :

$$\sum_{t=1}^T f_t(w_{t+1}) \leq \sum_{t=1}^T f_t(w_{T+1}) = \min_{w \in S} \sum_{t=1}^T f_t(w) \leq \sum_{t=1}^T f_t(u)$$

The equality follows from FTL definition.  $\square$

### 2.3.1 Online Quadratic Optimization

We will use FTL to solve online quadratic optimization, which is an online convex optimization where  $f_t(w) = \frac{1}{2} \|w - z_t\|_2^2$ .

Assume  $w \in S = \mathfrak{R}^d$ .

In this case, it is possible to derive steps of FTL in a closed form.

Let  $F_t(w) = \sum_{i=1}^{t-1} f_i(w) = \frac{1}{2} \sum_{i=1}^{t-1} \|w - z_i\|_2^2$

This function  $F_t$  is convex, so it enough to solve  $\nabla F_t(w) = 0$ :

$$\nabla F_t(w) = (t-1)w - \sum_{i=1}^{t-1} z_i = 0$$

Therefore  $w_t = \frac{1}{t-1} \sum_{i=1}^{t-1} z_i$

To apply Lemma 2.1, we express  $w_{t+1}$  as a function of  $w_t$ .

$$w_{t+1} = \frac{1}{t}(w_t(t-1) + z_t) = \left(1 - \frac{1}{t}\right)w_t + \frac{z_t}{t}$$

$$w_{t+1} - z_t = \left(1 - \frac{1}{t}\right)(w_t - z_t)$$

To get the bound on the regret, we need to analyze  $f_t(w_t) - f_t(w_{t+1})$

$$f_t(w_t) - f_t(w_{t+1}) = \frac{1}{2} \|w_t - z_t\|_2^2 - \frac{1}{2} \|w_t - z_{t+1}\|_2^2 = \frac{1}{2} \left(1 - \left(1 - \frac{1}{t}\right)^2\right) \|w_t - z_t\|_2^2 \leq \frac{1}{t} \|w_t - z_t\|_2^2$$

Define  $L = \max_i \|z_i\|_2^2$ , and since  $w_t$  is the average of  $z_1, \dots, z_{t-1}$  then,  $\|w_t\|_2^2 \leq L$ . From the triangle inequality:  $\|w_t - z_t\|_2^2 \leq (2L)^2$

To summarize:

$$\sum_{t=1}^T (f_t(w_t) - f_t(w_{t+1})) \leq 4L^2 \sum_{t=1}^T \frac{1}{t} \leq 4L^2 (\ln(T) + 1)$$

From Lemma 2.1, the regret is logarithmic in the number of time steps.

### 2.3.2 Bad Example - FTL is unstable

Consider online linear optimization with the following setting:

$$S = [-1, 1]$$

$$(z_t)_{t=1}^T = \left(-\frac{1}{2}, 1, -1, 1, -1, \dots\right)$$

The cumulative loss at time  $t$  is:

$$F_t(w) = \sum_{i=1}^{t-1} f_i(w) = w^T \left( \sum_{i=1}^{t-1} z_i \right) = \begin{cases} \frac{1}{2}w & \text{if } t \text{ is odd} \\ -\frac{1}{2}w & \text{if } t \text{ is even} \end{cases}$$

In every step,  $w_t$  is determined by the sign of  $\sum_{i=1}^{t-1} z_i$ . Which means  $w_1 = 0$ ,  $w_2 = 1$ ,  $w_3 = -1$ ,  $w_4 = 1$  etc. So  $\forall t > 1 : w_t^T z_t = 1$ .

However, for any constant strategy  $u \in [-1, 1]$ , the loss is at most  $\frac{1}{2}|u| \leq \frac{1}{2}$ . So FTL incurs a regret of linear magnitude as function of  $T$ .

## 2.4 Follow the Regularized Leader (FoReL)

As we seen om previous section, due to its instability FTL might incur regret of a linear magnitude. The problem is that the algorithm is unstable, it performs dramatic steps.

We address this problem by adding regularization. Instead of just selecting the best vector (the one that minimizes the loss retrospectively), we turn to minimize also a regularization function.

Recall the standard problem of linear regression. For an almost singular data matrix, the solution is unstable - for that case, ridge regression stabilizes the solution and stabilizes the bound over the loss.

Let  $R : S \rightarrow \mathfrak{R}_+$  be some convex regularization function.

Now, we minimize the retrospective loss, PLUS the regularization term.

### FoReL

$$\forall t : w_t := \operatorname{argmin}_{w \in S} \sum_{i=1}^{t-1} f_i(w) + R(w)$$

In the upcoming lessons, we will see that different regularization functions  $R$  provide different algorithms with different guarantees. For example, we will see that we can induce the multiplicative update rule that we have seen in *RMW* as a step of *FoReL*.

### 2.4.1 FoReL Example - Online linear optimization with l2-norm regularization

Consider  $l_2$  - norm regularization.

$$R(w) := \frac{1}{2\eta} \|w\|_2^2$$

Recall that the loss function of time  $t$  is

$$f_t(w) = z_t^T w$$

The function we minimize at time  $t$  is

$$F_t(w) = \sum_{i=1}^{t-1} f_i(w) + \frac{1}{2\eta} \|w\|_2^2$$

Let's solve for  $w_t$ . Since  $F_t$  is convex as a sum of convex functions, therefore it is enough to solve  $\nabla F_t(w) = 0$ . We have,

$$\nabla F_t(w) = \sum_{i=1}^{t-1} z_i + \frac{1}{2\eta} 2w_t = 0,$$

therefore

$$w_t = -\eta \sum_{i=1}^{t-1} z_i$$

For future lessons, note that  $w_t = w_{t-1} - \eta z_{t-1} = w_{t-1} - \eta \nabla f_{t-1}(w_{t-1})$ . I.e. the update rule is practically online gradient descent.

### 2.4.2 Bounding the loss of FoReL

Next, we bound the loss of FoReL. For that purpose, we apply Lemma 2.1. In order to apply Lemma 2.1, start the time at  $t = 0$  (instead of  $t = 1$ ) and define the regularization as the loss of time 0.  $f_0(w) := R(w)$ .

$$\begin{aligned} \forall u \in S : \sum_{t=0}^T (f_t(w_t) - f_t(u)) &\leq \sum_{t=0}^T (f_t(w_t) - f_t(w_{t+1})) \\ f_0(w_0) - f_0(u) + \sum_{t=1}^T (f_t(w_t) - f_t(u)) &\leq f_0(w_0) - f_0(w_1) + \sum_{t=1}^T (f_t(w_t) - f_t(w_{t+1})) \\ \sum_{t=1}^T (f_t(w_t) - f_t(u)) &\leq f_0(u) - f_0(w_1) + \sum_{t=1}^T (f_t(w_t) - f_t(w_{t+1})) \end{aligned}$$

**Corollary 2.2** *FoReL with a convex regularization function  $R$ , guarantees  $\forall u \in S$ :*

$$\text{Regret}_T(u) \leq R(u) - R(w_1) + \sum_{t=1}^T (f_t(w_t) - f_t(w_{t+1}))$$

Apply the corollary to online linear optimization with l2-regularization to get the following theorem.

**Theorem 2.3** *Solving online linear optimization by FoReL with l2-norm regularization, that is*

$$\begin{aligned} f_t(w) &= \langle z_t, w \rangle \\ S &= \mathfrak{R}^d \\ R(w) &= \frac{1}{2\eta} \|w\|_2^2 \end{aligned}$$

gives  $\forall u \in \mathfrak{R}^d$  s.t.  $\|u\| \leq B$ :

$$\text{Regret}_T(u) \leq \frac{1}{2\eta} B^2 + \eta \sum_{i=1}^T \|z_t\|_2^2$$

**Proof:** From corollary 2.2, we obtain:

$$\begin{aligned} \text{Regret}_T(u) &\leq R(u) - R(w_1) + \sum_{t=1}^T (f_t(w_t) - f_t(w_{t+1})) \\ &\leq \frac{1}{2\eta} \|u\|_2^2 - 0 + \sum_{t=1}^T (w_t - w_{t+1})^T z_t \end{aligned}$$

We have seen that  $w_t - w_{t+1} = \eta z_t$ . So,

$$\text{Regret}_T(u) \leq \frac{1}{2\eta} \|u\|_2^2 + \sum_{t=1}^T \eta \|z_t\|_2^2 \leq \frac{1}{2\eta} B^2 + \eta \sum_{t=1}^T \|z_t\|_2^2$$

□

In order to analyze the complexity of the bound above, we bound the right term with a parameter  $L$  and optimize  $\eta$ . Assume the average norm of  $z_t$  is bounded  $\frac{1}{T} \sum_{t=1}^T \|z_t\|_2^2 \leq L^2$  and obtain

$$\text{Regret}_T(u) \leq \frac{1}{2\eta} B^2 + \eta T L^2$$

The optimal  $\eta$  is  $\eta = \frac{B}{L\sqrt{2T}}$ , and we get:

$$\text{Regret}_T(u) \leq \frac{L\sqrt{2T}}{2B} B^2 + \frac{B}{L\sqrt{2T}} L^2 = \frac{LB\sqrt{T}}{\sqrt{2}} + \frac{LB\sqrt{T}}{\sqrt{2}} = LB\sqrt{2T}$$

In terms of complexity, the regret grows as square root  $t$ .

### 2.4.3 The doubling trick - $T$ is unknown

Notice, that the optimal  $\eta$  in the solution to the previous problem (online linear optimization with  $l_2$ -regularization) depends on  $T$  and is required for actually running the algorithm. However,  $T$  might be unknown ahead. This might happen generally in other algorithms as well.

A standard solution for that problem called the *doubling trick*. The idea is to split the time line into phases of lengths  $1, 2, 4, 8, \dots$ . In the  $m^{\text{th}}$  phase, initialize the algorithm and run it for  $2^m$  steps while assuming  $T = 2^m$ .

Assume we have an algorithm  $A$  that gives  $\text{Regret}_T(u) \leq \alpha\sqrt{T}$ , assuming  $T$  is known. If we run the doubling trick over phases for  $T$  steps we obtain:

$$\text{Regret}_T(u) \leq \sum_{m=0}^{\lceil \log_2(T) \rceil} \alpha\sqrt{2^m} \approx \alpha \frac{\sqrt{2}^{\log_2(T)} - 1}{\sqrt{2} - 1} = C\sqrt{T}$$

Thus, the regret grows like  $\sqrt{T}$  as we wanted. However, this algorithm does not seem reasonable at all. In every phase transition we forget everything we learned so far. In future lessons, we will have a parameter  $\eta_t$  that depends rather smoothly on the current time  $t$ .

**Note** (about FoReL without the doubling trick): We run the algorithm with  $\eta = \frac{D}{\sqrt{T}}$  (for some constant  $D$ ). So during the first  $\sqrt{T}$  steps, the regularization dominates the loss function  $F_t(w) = \sum_{i=1}^{t-1} f_i(w) + \sqrt{T}\|w\|_2^2$  meaning, that we practically "ignore" the inputs during these steps.