DSC 211 Introduction to Optimization Winter 2024 Instructor: Jun-Kun Wang
Scribe: Yunzhou Yan February 20, 2024
Editors/TAs: Merlin Chang, Marialena Sfyraki

Lecture 13: Online convex optimization

1 Introduction to online convex optimization

Online convex optimization (OCO) lies at the intersection between learning and convex optimization. Most OCO scenarios follow the following general protocol.

Protocol/Setting

1: **for** $t = 1, 2, \ldots$ **do**
2: Commit a point z_t with its convex decision space $Z \subset R^d$.
3: Receive a convex loss function $\ell_t(\cdot) : Z \to \mathbb{R}$ and incurs a loss $\ell_t(z_t)$.
4: **end for**

The goal of OCO is to compete with any fixed comparator in a convex decision space Z. Formally, the regret of the algorithm relative to any fixed benchmark z^* in Z is defined as

$$\text{Regret}_T(z_*) := \sum_{t=1}^{T} l_t(z_t) - \sum_{t=1}^{T} l_t(z_*)$$

We want to achieve sub-linear regret (e.g. $\frac{\text{Regret}_T(z_*)}{T} \to 0$,as $T \to \infty$). That is, the average regret $\frac{\text{Regret}_T(z_*)}{T}$ is vanishing.

Q: Why not compete with the best action at each around? i.e. why not consider

$$\sum_{t=1}^{T} l_t(z_t) - \sum_{t=1}^{T} \min_z l_t(z). \quad (1)$$

A: Recall we observe $\ell_t(\cdot)$ only after committing a point z_t in the decision space Z. For the metric defined on (1), when we do not commit a point that is the argmin of a loss function $\ell_t(\cdot)$ that we are going to receive, we will suffer some loss at each round, and consequently the value of (1) will grow linearly with T. This motivated us to consider "regret" with respect to a fixed benchmark action z_*:

1.1 Online linear optimization

Online linear optimization (OLO) is a special case of online convex optimization where the cost function l_t is linear at each timestep. Let $l_t(z) = c_t^T z_t$.

Protocol/Setting

for $t = 1, 2, \ldots$ **do**
  Commit a point z_t with its convex decision space $Z \subset R^d$.
  Receive a loss function $\ell_t(z) = c_t^\top z_t$ and incurs a loss $\ell_t(z_t)$.
end for

We can reduce online convex optimization to online linear optimization. We can bound the regret of OCO by the regret of OLO. Recall g_x is a subgradient of $f(x) : \mathcal{Z} \to \mathbb{R}$ at x if

$$f(y) \geq f(x) + \langle g_x, y - x \rangle, \quad \forall y \in C$$

Therefore,

$$\underbrace{l_t(z_t) - l_t(z_*)}_{\text{per-round regret of OCO}} \leq \underbrace{\langle z_t - z_*, c_t \rangle}_{\text{per-round regret of OLO}},$$

where c_t is the subgradient of $l_t(\cdot)$ at z_t.

2 Follow-the-Leader (FTL)

In the context of FTL, during the initial time step, we select the initial vector, and subsequently, at each subsequent time step, we opt for the vector with the minimum loss across all preceding rounds.

Algorithm 1 Follow-the-Leader(FTL)

Input a convex decision space $Z \in R^d$ and $z_{\text{init}} \in Z$.
  for $t = 1, 2, \ldots$ **do**
    if $t = 1$ **then**
      Commit $z_t = z_{\text{init}}$
    else
      Commit $z_t = \arg\min_{z \in Z} \sum_{s=1}^{t-1} l_s(z)$.
    end if
    Receive a convex loss function $l_t(\cdot) : Z \to R$ and suffer loss $l_t(z_t)$.
  end for

Lemma 1. *Let $z_1, z_2, \dots$ be the sequences of points generated by FTL. Then, for any benchmark $z_* \in Z$:*

$$Regret_T(z_*) = \sum_{t=1}^{T} l_t(z_t) - \sum_{t=1}^{T} l_t(z_*) \leq \sum_{t=1}^{T} l_t(z_t) - \sum_{t=1}^{T} l_t(z_{t+1})$$

Now, we would like to prove the Lemma. Before we show the proof, it is clear to see that if we subtract the term $\sum_{t=1}^{T} l_t(z_t)$ from both sides, the lemma will be equivalent to $\sum_{t=1}^{T} l_t(z_{t+1}) \leq \sum_{t=1}^{T} l_t(z_*)$

Proof. We prove this inequality by induction. The base case of T = 1 follows directly from the definition of z_{t+1}. i.e. $z_2 = \arg\min_{z \in Z} l_1(z) \Rightarrow z_2 \leq z_*$
Assume the inequality holds for $t = T - 1$, then for all $z_* \in Z$, we have

$$\sum_{t=1}^{T-1} l_t(z_{t+1}) \leq \sum_{t=1}^{T-1} l_t(z_*)$$

Now adding $l_T(z_{T+1})$ to both sides, then we get

$$l_T(z_{T+1}) + \sum_{t=1}^{T-1} l_t(z_{t+1}) \leq l_T(z_{T+1}) + \sum_{t=1}^{T-1} l_t(z_*)$$

$$\Leftrightarrow \sum_{t=1}^{T} l_t(z_{t+1}) \leq l_T(z_{T+1}) + \sum_{t=1}^{T-1} l_t(z_*).$$

The above holds for all $z_* \in Z$ and let $z_* = z_{T+1}$. Then, the above becomes

$$\sum_{t=1}^{T} l_t(z_{t+1}) \leq \sum_{t=1}^{T} l_T(z_{T+1}) \tag{2}$$

$$\leq \sum_{t=1}^{T} l_T(z_*), \qquad \forall z_* \in Z, \tag{3}$$

where the last inequality we used that $z_{T+1} = \arg\min_{z \in Z} \sum_{t=1}^{T} l_T(z)$. This concludes our inductive argument. □

Consider $l_t(z) = \|z - c_t\|_2^2$

Theorem 1. *Assume* $\max_t \|\mathbf{c}_t\| \leq L$. *Then, Follow-The-Leader (FTL) has regret at most*

$$Regret_T(\mathbf{z}^*) \leq 4L^2(\log(T) + 1),$$

for any $\mathbf{z}^* \in \mathbb{R}^d$.

Proof. We assume that $Z = \mathbb{R}^d$.
Using the Follow-The-leader (FTL) rule,

$$\mathbf{z}_t = \arg\min_{\mathbf{z}\in Z}(\sum_{s=1}^{t-1} \|\mathbf{z} - c_s\|_2^2) = \arg\min_{\mathbf{z}\in Z} F(\mathbf{z})$$

Since $\mathbf{F}(\mathbf{z})$ is convex, we have

$$\frac{\partial F(\mathbf{z}_t)}{\partial \mathbf{z}} = \sum_{s=1}^{t-1} 2(\mathbf{z}_t - c_s) = 0 \Rightarrow \mathbf{z}_t = \frac{1}{t-1}\sum_{s=1}^{t-1} \mathbf{c}_s$$

namely, $\mathbf{z}_t$ is the average of $\mathbf{c}_1, \ldots, \mathbf{c}_{t-1}$. Note that we can rewrite

$$\mathbf{z}_{t+1} = \frac{1}{t}(\mathbf{c}_t + (t-1)\mathbf{z}_t) = \left(1 - \frac{1}{t}\right)\mathbf{z}_t + \frac{1}{t}\mathbf{c}_t,$$

which yields

$$\mathbf{z}_{t+1} - \mathbf{c}_t = \left(1 - \frac{1}{t}\right)(\mathbf{z}_t - \mathbf{c}_t).$$

Therefore,

$$l_t(\mathbf{z}_t) - l_t(\mathbf{z}_{t+1}) = \frac{1}{2}\|\mathbf{z}_t - \mathbf{c}_t\|_2^2 - \frac{1}{2}\|\mathbf{z}_{t+1} - \mathbf{c}_t\|_2^2 = \frac{1}{2}\left(1 - \left(1 - \frac{1}{t}\right)^2\right)\|\mathbf{z}_t - \mathbf{c}_t\|_2^2 = \frac{1}{2}(\frac{2}{t} - \frac{1}{t^2})\|\mathbf{z}_t - \mathbf{c}_t\|_2^2$$

$$\leq \frac{1}{2}(\frac{2}{t})\|\mathbf{z}_t - \mathbf{c}_t\|_2^2 = \frac{1}{t}\|\mathbf{z}_t - \mathbf{c}_t\|_2^2.$$

Let $L = \max_t \|\mathbf{c}_t\|$. Since $\mathbf{z}_t$ is the average of $\mathbf{c}_1, \ldots, \mathbf{c}_{t-1}$, we have that $\|\mathbf{z}_t\| \leq L$, and therefore, by the triangle inequality, $\|\mathbf{z}_t - \mathbf{c}_t\|_2 \leq 2L$. We have therefore obtained:

$$\sum_{t=1}^{T}(l_t(\mathbf{z}_t) - l_t(\mathbf{z}_{t+1})) \leq (2L)^2 \sum_{t=1}^{T} \frac{1}{t}.$$

Now we want to use induction to prove the inequality

$$\sum_{t=1}^{T} \frac{1}{t} \leq \log(T) + 1,$$

Firstly, let us consider the base case $T = 1$, it is obviously true that $\sum_{t=1}^{1} \frac{1}{t} = 1 <= \log(1) + 1 = 1$
Now let us assume that the inequality holds for $T - 1$. Then we have

$$\sum_{t=1}^{T-1} \frac{1}{t} \leq \log(T-1) + 1,$$

$$\log(T)+1 = \log((T-1)\frac{T}{T-1})+1 = \log(T-1)+\log(\frac{T}{T-1})+1 \geq \sum_{t=1}^{T-1}\frac{1}{t}+\log(\frac{T}{T-1})$$

Using the well-known inequality $\log(x) \geq 1 - \frac{1}{x}$, we have $\log(\frac{T}{T-1}) \geq \frac{1}{T}$. After substitution, we have

$$\log(T) + 1 \geq \sum_{t=1}^{T-1}\frac{1}{t} + \frac{1}{T} = \sum_{t=1}^{T}\frac{1}{t}$$

Combining the above with Lemma 1 and using the inequality

$$\sum_{t=1}^{T}\frac{1}{t} \leq \log(T) + 1,$$

we can conclude that $\text{Regret}_T(\mathbf{z}^*) \leq \sum_{t=1}^{T}(l_t(\mathbf{z}_t) - l_t(\mathbf{z}_{t+1})) \leq (2L)^2 \sum_{t=1}^{T}\frac{1}{t} \leq 4L^2(\log(T)+1)$,which completes the proof. □

Example: Failure of Follow-the-Leader:
Let the decision space $z = [-1, 1] \subseteq \mathbb{R}$. Let $\ell_t(z) = c_t z \in \mathbb{R}$, where c_t is defined as follows:

$$c_t = \begin{cases} -0.5 & \text{,if } t = 1, \\ 1 & \text{,if } t \text{ is even,} \\ -1 & \text{,if } t \text{ is odd.} \end{cases}$$

- Let $z_1 = \theta \in [-1, 1]$, $c_1 = -0.5$.
  We hence have: $l_1(z_1) = c_1 z_1 = -0.5 z_1$.
- $z_2 = \arg\min_z l_1(z) = \arg\min_z c_1 z = \arg\min_{z\in[-1,1]} -0.5z = 1$.
  We hence have: $c_2 = 1, l_2(z) = c_2 z_2 = 1$
- $z_3 = \arg\min_z l_1(z) + l_2(z) = \arg\min_z (c_1 + c_2)z = \arg\min_{z\in[-1,1]} 0.5z = -1$.
  We hence have: $c_3 = -1, l_3(z) = c_3 z_3 = 1$
- $z_4 = \arg\min_z l_1(z)+l_2(z)+l_3(z) = \arg\min_z (c_1+c_2+c_3)z = \arg\min_{z\in[-1,1]} -0.5z = 1$.
  We hence have: $c_4 = 1, l_4(z) = c_4 z_4 = 1$.

From this pattern, we can deduce that if $t > 1$

$$z_t = \arg\min_{z\in[-1,1]}(\sum_{s=1}^{t-1} a_s l_s(z)) = \begin{cases} \arg\min_{z\in[-1,1]} 0.5z = -1 & \text{if } t \text{ is odd,} \\ \arg\min_{z\in[-1,1]} -0.5z = 1 & \text{if } t \text{ is even,} \end{cases}$$

Therefore, the cumulative loss of the FTL algorithm will therefore be $-0.5z_1 + (T-1)$, where $z_1 \in [-1, 1]$. while the cumulative loss of the fixed solution $z_* = 0 \in Z$ is 0. Thus, the regret of FTL is $\Omega(T)$, which is not sub-linear!

3 Follow-the-Regularized-Leader (FTRL)

Follow-the-Regularized-Leader(FTRL) is a natural modification of the basic FTL algorithm in which we minimize the loss on all past rounds plus a regularization term. The goal of the regularization term is to stabilize the solution.

Algorithm 2 Follow-the-Regularized-Leader(FTRL)

Require: a strongly convex regularizer $R(\cdot) : Z \to \mathbb{R}$.
for $t = 1, 2, \ldots$ **do**
  Commit $z_t = \arg\min_{z \in Z}(\sum_{s=1}^{t-1} l_s(z)) + R(z)$.
  Receive a convex loss function $l_t(\cdot) : Z \to \mathbb{R}$ and suffer loss $l_t(z_t)$.
end for

Remark: There are connections between FTRL and Online Gradient Descent. Consider $\ell_t(z) = \langle z, c_t \rangle$ and $R(z) = \frac{1}{2}\eta\|z\|_2^2$. From the update rule of FTRL, we know that $z_t = \arg\min_{z \in Z} \sum_{s=1}^{t-1}\langle z, c_s \rangle + \frac{1}{2\eta}\|z\|_2^2 = \arg\min_{z \in Z} \phi(z)$, where $\phi(z) = \sum_{s=1}^{t-1}\langle z, c_s \rangle + \frac{1}{2\eta}\|z\|_2^2$. Since $\phi(z)$ is convex and differentiable, by the optimaliy condition, z_t satisfies:

$$\nabla\phi(z_t) = \sum_{s=1}^{t-1} c_s + \frac{1}{\eta} z_t = 0 \tag{4}$$

$$\Rightarrow z_t = -\eta \sum_{s=1}^{t-1} c_s = z_{t-1} - \eta c_{t-1} \tag{5}$$

$$= z_{t-1} - \eta \nabla l_{t-1}(z_{t-1}) \tag{6}$$

which becomes Online Gradient Descent!

Lemma 2. *Let $z_1, z_2, \ldots$ be the sequences of points generated by FTRL. Then, for any benchmark $z_* \in Z$:*

$$\text{Regret}_T(z_*) = \sum_{t=1}^{T} l_t(z_t) - \sum_{t=1}^{T} l_t(z_*) \leq R(z_*) - R(z_1) + \sum_{t=1}^{T} l_t(z_t) - \sum_{t=1}^{T} l_t(z_{t+1})$$

Proof. Observe that running FTRL on $l_1, \ldots, l_T$ is equivalent to running FTL on $l_0, l_1, \ldots, l_T$ where $l_0 = R$. Using Lemma 1, we obtain

$$\sum_{t=0}^{T}(l_t(z_t) - l_t(z_*)) \leq \sum_{t=0}^{T}(l_t(z_t) - l_t(z_{t+1})),$$

which leads to

$$
\begin{aligned}
&R(z_0) - R(z_*) + \text{Regret}_T(z_*) \\
&= \sum_{t=0}^{T} (l_t(z_t) - l_t(z_*)) \\
&\leq \sum_{t=0}^{T} (l_t(z_t) - l_t(z_{t+1})) \\
&= R(z_0) - R(z_1) + \sum_{t=1}^{T} l_t(z_t) - \sum_{t=1}^{T} l_t(z_{t+1}),
\end{aligned}
$$

rearranging the terms completes the proof. □

Theorem 2. *Consider running FTRL on a sequence of linear functions, $l_t(z) = \langle z, c_t \rangle$ for all t, with $Z = \mathbb{R}^d$, and with the regularizer $R(z) = \frac{1}{2\eta}\|z\|_2^2$. Then, for all z_* we have*

$$
Regret_T(z_*) \leq \frac{1}{2\eta}\|z_*\|_2^2 + \eta \sum_{t=1}^{T} \|z_t\|_2^2.
$$

In particular, assume $\|z_*\|_2 \leq D$ and let G be such that $\|c_t\|_2 \leq G$ for all t, then by setting $\eta = \frac{D}{G\sqrt{2T}}$ we obtain

$$
Regret_T(z_*) \leq DG\sqrt{2T}.
$$

Proof. Using Lemma 2 and the FTRL update step $z_{t+1} = z_t - \eta c_t$,

$$
Regret_T(z_*) \leq R(z_*) - R(z_1) + \sum_{t=1}^{T} l_t(z_t) - \sum_{t=1}^{T} l_t(z_{t+1}) \leq \frac{1}{2\eta}\|z_*\|_2^2 + \sum_{t=1}^{T} \langle z_t - z_{t+1}, c_t \rangle
$$

$$
= \frac{1}{2\eta}\|z_*\|_2^2 + \eta \sum_{t=1}^{T} \|c_t\|_2^2
$$

Now, if we substitute $\|z_*\|_2 \leq D$ and $\|c_t\|_2 \leq G$, we have

$$
Regret_T(z_*) \leq DG\sqrt{2T}
$$

□

So far, we have focused on Euclidean Regularization $R(x) = \frac{1}{2}\|x\|_2^2$. Now consider general $R(x)$ that is strongly convex with respect to norm $\|\cdot\|$.

Lemma 3. *Let $R(z) : Z \to R$ be a μ-strongly convex function over Z with respect to a norm $\|\cdot\|$. If $\ell_t(\cdot)$ is L_t-Lipschitz with respect to $\|\cdot\|$, then*

$$\ell_t(z_t) - \ell_t(z_{t+1}) \le L_t\|z_t - z_{t+1}\| \le \frac{L_t^2}{\mu}.$$

Remark: The lemma indicates that the difference between consecutive predictions is bounded by $\frac{L_t}{\mu}$!

Proof. We have

$$z_t = \arg\min_{z\in Z} \sum_{s=1}^{t-1} l_s(z) + R(z) = \arg\min_{z\in Z} \sum_{s=1}^{t-1} F_t(z)$$

Since $l_s(z)$ is linear and $R(z)$ is μ-strongly convex, then $F_t(z)$ is also μ-strongly convex!(We have proved it in our homework)
By strong convexity,

$$F_t(z_{t+1}) \ge F_t(z_t) + \langle \nabla F_t(z_t), z_{t+1} - z_t\rangle + \frac{\mu}{2}\|z_{t+1} - z_t\|^2$$

$$F_{t+1}(z_t) \ge F_{t+1}(z_{t+1}) + \langle \nabla F_{t+1}(z_{t+1}), z_t - z_{t+1}\rangle + \frac{\mu}{2}\|z_{t+1} - z_t\|^2$$

Since z_t and z_{t+1} are the minimizers of $F_t(z)$ and $F_{t+1}(z)$, then we have $\nabla F_t(z_t) = 0$ and $\nabla F_{t+1}(z_{t+1}) = 0$. As a result, the inequalities become

$$F_t(z_{t+1}) \ge F_t(z_t) + \frac{\mu}{2}\|z_{t+1} - z_t\|^2$$

$$F_{t+1}(z_t) \ge F_{t+1}(z_{t+1}) + \frac{\mu}{2}\|z_{t+1} - z_t\|^2$$

Combining the above two inequalities, we have

$$F_{t+1}(z_t) - F_t(z_t) \ge F_{t+1}(z_{t+1}) - F_t(z_{t+1}) + \mu\|z_t - z_{t+1}\|^2.$$

$$\Leftrightarrow l_t(z_t) - l_t(z_{t+1}) \ge \mu\|z_t - z_{t+1}\|^2$$

Moreover, by Lipschitzness of $\ell_t(\cdot)$,

$$L_t\|z_t - z_{t+1}\| \ge |\ell_t(z_t) - \ell_t(z_{t+1})| \ge \mu\|z_t - z_{t+1}\|^2$$

$$\frac{L_t}{\mu} \ge \|z_t - z_{t+1}\|.$$

□

Theorem 3. *Suppose each $\ell_1(\cdot), \ell_2(\cdot), \ldots, \ell_T(\cdot)$ is convex, and each $\ell_t(\cdot)$ is L_t-Lipschitz with respect to a norm $\|\cdot\|$. Denote L_0 as the constant such that $L_t \le L_0$. FTRL with a μ-strongly convex regularizer $R(\cdot)$ with respect to $\|\cdot\|$ (e.g. $\frac{\mu}{2}\|\cdot\|^2$) has*

$$Regret_T(z^*) \le R(z^*) - R(z_1) + \frac{TL_0^2}{\mu}.$$

Proof.

$$\begin{aligned}
\mathrm{Regret}_T(z^*) &= \sum_{t=1}^{T} \ell_t(z_t) - \ell_t(z^*) && \text{(By definition)} \\
&\le R(z_*) - R(z_1) + \sum_{t=1}^{T} (\ell_t(z_t) - \ell_t(z_{t+1})) && \text{(By Lemma 2)} \\
&\le R(z_*) - R(z_1) + \sum_{t=1}^{T} \frac{L_t^2}{\mu} && \text{(By Lemma 3)} \\
&\le R(z^*) - R(z_1) + \frac{TL_0^2}{\mu} && \text{(By } L_t \le L_0\text{)}
\end{aligned}$$

□

Bibliographic notes

For more details about Online Convex Optimization, see e.g., [1], [2], and [3].

References

[1] Shai Shalev-Shwartz Online Learning and Online Convex Optimization Foundations and Trends in Machine Learning. 2011.

[2] Elad Hazan. Introduction to Online Convex Optimization. Second Edition. The MIT Press. 2022.

[3] Francesco Orabona A Modern Introduction to Online Learning arXiv:1912.13213. 2023