CSCI699: Theory of Machine Learning	Fall 2021
Lecture 18: Online Convex Optimization	
Instructor: Vatsal Sharan	*Scribe: Grace Zhang*

Recap: Last lecture, we showed the Weighted Majority algorithm achieves bounded regret in the online learning setting for finite hypothesis classes. What if $|\mathcal{H}| = \infty$?

Theorem 1. *For every hypothesis class $\mathcal{H}$, there exists an algorithm for online learning with a regret bound:*

$$\sum_{t=1}^{T} |p_t - y_t| - \min_{h\in\mathcal{H}} \sum_{t=1}^{T} |h(x_t) - y_t| \leq \sqrt{2\mathrm{Ldim}(\mathcal{H}) \log(eT)T}$$

1 Online Convex Optimization

In online convex optimization, at each timestep, the learner chooses $w_t \in S$ (for some domain S). We have convex loss $f_t : S \to \mathbb{R}$ at every time t. The learner suffers $f_t(w_t)$ at time t. We can write regret as:

$$\mathrm{Regret}(T) = \sum_{t=1}^{T} f_t(w_t) - \min_{w\in S} \sum_{t=1}^{T} f_t(w)$$

Example 1. *(linear regression)*

$S = \mathbb{R}^d$

$f_t(w) = (\langle w, x_t\rangle - y_t)^2$*(convex as a function of w)*

Note that the data (x_t, y_t) are baked into f_t.

Example 2. *(experts)*

A set of experts are discrete which is not a convex set. However, we can create a convex set by drawing from a distribution over experts.

S = *simplex over* $\mathbb{R}^d$ $(\delta_d = \{w \in \mathbb{R}^d, w_i \geq 0 \ \forall i \in [d], \sum w_i = 1\})$ $f_t(w) = \langle w, v_t\rangle$ $v_t = (l(h_1(x_t), y_t), l(h_2(x_t), y_t), \cdots, l(h_d(x_t), y_t))$ $(l(h_i(x), y) \in [0, 1])$

Note that f_t is merging the loss function and data.

2 Algorithmic Frameworks (even beyond convex functions)

2.1 Algorithm: Follow-the-Leader (FTL)

At every timestep t, we play

$$w_t \in \arg\min_{w \in S} \sum_{i=1}^{t-1} f_i(w) \tag{1}$$

Lemma 2. *(FTL vs. lookahead oracle)*

Let $w_1, \cdots, w_T$ be produced by the FTL algorithm according to Equation 1. For any $u \in S$, define

$$\text{Regret}(u, T) = \sum_{t=1}^{T} \left[f_t(w_t) - f_t(u) \right]$$

Then $\text{Regret}(u, T) \leq \sum_{t=1}^{T} \big[\underbrace{f_t(w_t) - f_t(w_{t+1})}_{\textit{how stable FTL is}} \big].$

$$\therefore \text{Regret}(T) = \max_{u \in S} \text{Regret}(u, T) \leq \sum_{t=1}^{T} \left[f_t(w_t) - f_t(w_{t+1}) \right]$$

Proof. It suffices to show that $\forall u \in S$,

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

We prove this by induction: The base case can be shown easily. Assume inductive hypothesis on $T - 1$.

$$\begin{aligned}
& \sum_{i=1}^{T-1} f_t(w_{t+1}) \leq \sum_{i=1}^{T-1} f_t(u) && \forall u \in S \\
\implies & \sum_{t=1}^{T} f_t(w_{t+1}) \leq \sum_{i=1}^{T-1} f_t(u) + f_T(w_{T+1}) && \forall u \in S \\
& \qquad \therefore \text{ this holds for } u = w_{T+1} \\
& \sum_{t=1}^{T} f_t(w_{t+1}) \leq \sum_{t=1}^{T} f_t(w_{T+1}) \\
& \qquad \text{since } w_{T+1} \in \arg\min_{w \in S} \sum_{t=1}^{T} f_t(w) \\
& \sum_{t=1}^{T} f_t(w_{t+1}) \leq \sum_{t=1}^{T} f_t(u) && \forall u \in S
\end{aligned}$$

This completes the inductive proof and the Lemma follows. ☐

Next we look at an example of where FTL works and one where it does not.

Example 3. *Quadratic Optimization (FTL works)*

Lemma 3. *Assume* $f_t(w) = \frac{1}{2}\|w - z_t\|_2^2$ *where* $\|z_t\|_2 \leq B\ \forall t \in [T]$. *Then FTL has regret* $O\big(B^2 \log(T)\big)$.

Proof. Let's look at the FTL solution given some examples.

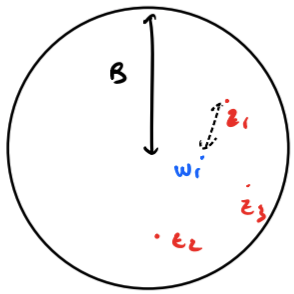

As we can see, FTL has the closed form solution:

$$
\begin{aligned}
w_t &= \frac{1}{t-1}\sum_{i=1}^{t-1} z_i \\
\implies w_{t+1} &= \frac{(t-1)w_t + z_t}{t} \\
&= \left(1 - \frac{1}{t}\right) w_t + \frac{z_t}{t}
\end{aligned}
$$

Using Lemma 2. (FTL vs. lookahead oracle),

$$
\begin{aligned}
\text{Regret}(T) &\leq \sum_{t=1}^{T} \big[f_t(w_t) - f_t(w_{t+1})\big] \\
f_t(w_t) - f_t(w_{t+1}) &= \frac{1}{2}\|w_t - z_t\|_2^2 - \frac{1}{2}\left\|(1 - \frac{1}{t})w_t + \frac{z_t}{t} - z_t\right\|_2^2 \\
&= \frac{1}{2}\Big(1 - (1 - \frac{1}{t})^2\Big)\|w_t - z_t\|_2^2 \\
&\leq \left(\frac{1}{t}\right)\|w_t - z_t\|_2^2 \\
&\leq \frac{4B^2}{t} \text{ (any 2 points are at most } 2B \text{ apart)} \\
\therefore \text{Regret}(T) &\leq 4B^2 \sum_{t=1}^{T}\left(\frac{1}{t}\right)
\end{aligned}
$$

We have the algebraic property that,

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

$$\therefore \text{Regret}(T) = O(B^2 \log(T))$$

$\square$

Example 4. *Linear Optimization (FTL fails)*

Let $S = [-1, 1]$ be FTL's possible predictions. Consider the linear functions $f_t(w) = wv_t$ (in $d = 1$ dimensions) where $(v_1, v_2, \cdots) = (-0.5, 1, -1, 1, -1, \cdots)$

What does FTL do here?

Initialize $w_1 = 0 \to loss(w_1, v_1) = w_1 v_1 = 0$.

Since $f_1(w) = -0.5w \implies w_2 = 1 \to loss(w_2, v_2) = w_2 v_2 = 1$.

Then $f_1(w) + f_2(w) = 0.5w \implies w_3 = -1 \to loss(w_3, v_3) = 1$. And so on and so forth.

$\therefore$ *FTL gets loss 1 on every example except the first. An expert $u = 0$ gets 0 loss.*

$$\therefore \text{Regret}(u, T) = T - 1$$

$$\therefore (T) \geq T - 1$$

Why are these two examples different? For quadratic functions, w_t and w_{t+1} get closer and closer (we get low regret). For linear functions, w_t and w_{t+1} do not get closer (we get high regret).

2.2 Algorithm: Follow the Regularized Leader (FTRL)

Let $\Psi : S \to \mathbb{R}$ be a function called a regularizer. Let $f_1, \cdots, f_T$ be the sequence of loss functions played by the environment.

FTRL algorithm: At every time t, choose,

$$w_t \in \arg\min_{w \in S} \left(\Psi(w) + \sum_{i=1}^{t-1} f_i(w) \right)$$

(FTL is FTRL with $\Psi = 0$)

Next, let's see how FTRL does in the linear case.

Example 5. *Linear f_t, quadratic Ψ*

Theorem 4. *For any $\eta > 0$, FTRL with $S \subseteq \mathbb{R}^d$ a convex set and $\Psi(w) = \dfrac{\|w\|_2^2}{2\eta}$, $f_t(w) = \langle w, v_t \rangle$ satisfies,*

$$\text{Regret}(u, T) \leq \underbrace{\frac{1}{2\eta}\|u\|_2^2}_{\textit{large when } \eta \textit{ is small}} + \underbrace{\eta \sum_{t=1}^{T} \|v_t\|_2^2}_{\textit{large when } \eta \textit{ is large}}$$

If $\|u\|_2 \leq B$ and $\|v_t\|_2 \leq L$, then choosing $\eta = \dfrac{B}{L\sqrt{T}}$ gives $\text{Regret}(T) = O(BL\sqrt{T})$.

Proof. FTRL: $w_t = \operatorname*{arg\,min}_{w \in S} \left(\dfrac{1}{2\eta}\|w\|_2^2 - \displaystyle\sum_{i=1}^{t-1} \langle w, \theta_T \rangle \right)$

where $\theta_t = -\displaystyle\sum_{i=1}^{t-1} v_i$.

By adding $\dfrac{\eta}{2}\|\theta_t\|_2^2$ (independent of w so the arg min is unchanged) to complete the square,

$$w_t = \operatorname*{arg\,min}_{w \in S} \left(\frac{1}{2\eta}\|w - \eta\theta_t\|_2^2 \right)$$

Therefore,

1. for $S = \mathbb{R}^d, w_t = \eta\theta_t$

   Rewriting, $w_t = -\eta(v_1 + v_2 + \cdots + v_{t-1})$

   $\implies w_{t+1} = w_t - \eta v_t$

2. for $S \neq \mathbb{R}^d, w_t = \Pi_S(\eta\theta_t)$

   where $\Pi_S(x) = \operatorname*{arg\,min}_{w \in S} \|w - x\|_2^2$

   This is called a lazy projection step: accumulate gradients θ_t, only project for making predictions (also called Nesterov's dual averaging method).

By using Lemma 2 (FTL vs. one-step lookahead) on the sequence $\Psi, f_1, \cdots, f_T$, for any $u \in S$:

$$\big[\Psi(w_0) - \Psi(u)\big] + \sum_{t=1}^{T} \big[f_t(w_t) - f_t(u)\big] \leq \big[\Psi(w_0) - \Psi(w_1)\big] + \sum_{t=1}^{T} \big[f_t(w_t) - f_t(w_{t+1})\big]$$

Since $\Psi(w_1) \geq 0$,

$$
\begin{aligned}
\mathrm{Regret}(u,T) &= \sum_{t=1}^{T}\left[f_t(w_t) - f_t(u)\right] \\
&\le \frac{1}{2\eta}\left\|u\right\|_2^2 + \sum_{t=1}^{T}\left[f_t(w_t) - f_t(w_{t+1})\right] \\
f_t(w_t) - f_t(w_{t+1}) &= \langle v_t, w_t - w_{t+1}\rangle \\
&\le \left\|v_t\right\|_2 \cdot \left\|w_t - w_{t+1}\right\|_2 \\
&= \left\|v_t\right\|_2 \cdot \left\|\Pi_S(\eta\theta_t) - \Pi_S(\eta\theta_{t+1})\right\|_2
\end{aligned}
$$

Claim 5. *(Projections onto convex sets are contractive)*

If S is convex, then for any $x_1, x_2 \in \mathbb{R}^d$,

$$
\left\|\Pi_S(x_1) - \Pi_S(x_2)\right\|_2 \le \left\|x_1 - x_2\right\|_2
$$

Proof. (by picture)

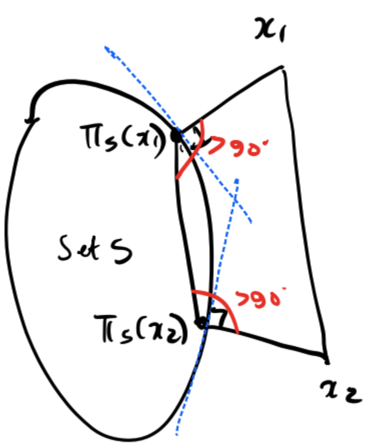

□

$$
\begin{aligned}
f_t(w_t) - f_t(w_{t+1}) &\le \left\|v_t\right\|_2 \cdot \left\|\eta\theta_t - \eta\theta_{t+1}\right\|_2 \\
&= \eta\left\|v_t\right\|_2^2 \\
\therefore \mathrm{Regret}(u,T) &\le \frac{1}{2\eta}\left\|u\right\|_2^2 + \eta\sum_{t=1}^{T}\left\|v_t\right\|_2^2
\end{aligned}
$$

□

3 Beyond linear: Online Convex Optimization

Consider convex functions to avoid intractability. For convex functions, a linear approximation suffices.

3.1 Algorithm: Online Gradient Descent (OGD)

Algorithm 1 Online Gradient Descent

Let $w_1 = 0,\ \theta_1 = 0$
for $t = 1, 2, \ldots, T$ **do**
    Predict w_t, receive f_t
    find gradient $v_t = \nabla f_t(w_t)$
    if $S = \mathbb{R}^d$ **then**
        $w_{t+1} = w_t - \eta \nabla v_t$
    else
        $w_{t+1} = \Pi_S(\eta \theta_{t+1}),\ \theta_{t+1} = \theta_t - v_t$
    end if
end for=0

Note: If f is not differentiable, use subgradients.

Theorem 6. *OGD enjoys the following regret bound for every $w^* \in S$,*

$$\text{Regret}(w^*, T) \leq \frac{\|w^*\|_2^2}{2\eta} + \frac{\eta}{2} \sum_{t=1}^{T} \|v_t\|_2^2$$

If $\|v_t\|_2 \leq \rho\ \forall t$ (f_t is ρ-Lipschitz $\forall t$) and $\|w^*\|_2 \leq B$, then setting $\eta = \dfrac{B}{\rho\sqrt{T}}$ yields,

$$\text{Regret}(T) \leq B\rho\sqrt{T}$$

Proof. Next Lecture □