Online Convex Programming and Gradient Descent

Instructors: Sham Kakade and Ambuj Tewari

1 Online Convex Programming

The *online convex programming problem* is a sequential paradigm where at each round the learner chooses decisions from a convex feasible set $D \subset \mathbb{R}^d$. In this problem, at each round t:

1. the learner chooses a decision $w_t \in D$, where D is a convex subset of $\mathbb{R}^d$.
2. nature chooses a *convex* cost function in $c_t : D \to \mathbb{R}$.
3. the learner incurs the cost $c_t(w_t)$, and the cost function $c_t(\cdot)$ is revealed to the algorithm.

Crucially, the algorithm learns c_t only after the decision w_t is chosen. Hence at time t, the algorithm has knowledge of the previous functions, $\{c_1(\cdot), \ldots, c_{t-1}(\cdot)\}$. Importantly, no statistical assumptions on the sequence of convex functions are made — they should be thought of as an arbitrary sequence unknown apriori to the algorithm.

If algorithm A uses the sequence of decisions

$$\{w_1, \ldots, w_T\}$$

on the sequence

$$\{c_1, \ldots, c_T\} ,$$

then A has regret, at time T in comparison to the best constant decision, defined as:

$$R_T(A) = \sum_{t=1}^{T} c_t(w_t) - \inf_{w \in D} \sum_{t=1}^{T} c_t(w)$$

We are interested in algorithms with little regret.

1.1 Online Gradient Descent

The simplest algorithm to consider here is the gradient descent algorithm. There are two issues we must address. First, we must ensure our decisions are always in the feasible set D. The second is that the gradient may not be defined.

To address, the later issue, we work with a subgradient. A subgradient $\nabla c(w)$ of a convex function $c(\cdot)$ at w satisfies, for all $w' \in D$

$$c(w') - c(w) \geq \nabla c(w) \cdot (w' - w)$$

A subgradient always exists for a convex function, though it may not be unique.

Define the Online Gradient Descent algorithm (GD) with fixed learning rate η is as follows: at $t = 1$, select any $w_1 \in D$, and update the decision as follows

$$w_{t+1} = \Pi_D[w_t - \eta \nabla c_t(w_t)]$$

where $\Pi_D[w]$ is the projection of w back into D, i.e. it is the closest point (under the L_2 norm) in D to w. More formally:

$$\Pi_D[w] \in \text{argmin}_{w' \in D} ||w - w'||_2$$

Hence, $w_{t+1} \in D$.

Theorem 1.1. *Assume that D is convex, closed, non-empty, and bounded. In particular, there exists a constant D_2 such for all $w, w' \in D$,*

$$||w - w'||_2 \leq D_2$$

Also, assume that for all times t, that

$$||\nabla c_t(w_t)||_2 \leq G_2$$

If we set $\eta = \frac{D_2}{G_2}\sqrt{\frac{1}{T}}$, then for all sequences of convex functions $\{c_1, \ldots c_T\}$

$$R_T(GD) \leq D_2\, G_2 \sqrt{T}$$

Note that there is no explicit dimensionality dependence.

We now provide the proof. Throughout, let $\nabla_t = \nabla c_t(w_t)$. Let w^* be a minimizer of $\sum_t^T c_t(w)$ (which exists since D is closed and convex). By convexity, we have

$$R_T(A) = \sum_{t=1}^{T} (c_t(w_t) - c_t(w^*)) \leq \sum_{t=1}^{T} \nabla_t \cdot (w_t - w^*)$$

Now we appeal to the following Lemma.

Lemma 1.2. *Let w^* be an arbitrary point in D. The decisions of GD algorithm satisfy:*

$$\sum_{t=1}^{T} \nabla_t \cdot (w_t - w^*) \leq \frac{1}{2\eta} D_2^2 + \frac{\eta}{2} G_2^2 T$$

Proof. A fundamental property of projections into convex bodies is that for an arbitrary $w' \in \mathbb{R}^d$, we have for all $w \in D$:

$$||\Pi_D[w'] - w||_2^2 \leq ||w' - w||_2^2$$

Using the notation that $|| \cdot ||$ refers to the L_2 norm:

$$\begin{aligned} ||w_t - w^*||^2 - ||w_{t+1} - w^*||^2 &= ||w_t - w^*||^2 - ||\Pi_D[w_t - \eta\nabla_t] - w^*||^2 \\ &\geq ||w_t - w^*||_2^2 - ||w_t - \eta\nabla_t - w^*||_2^d \\ &= 2\eta\nabla_t \cdot (w_t - w^*) - \eta^2 ||\nabla_t||_2^2 \end{aligned}$$

and so

$$\nabla_t \cdot (w_t - w^*) \leq \frac{1}{2\eta}(||w_t - w^*||^2 - ||w_{t+1} - w^*||^2) + \frac{\eta}{2} G_2^2$$

using the definition of G_2. Summing over t,

$$\begin{aligned} \sum_{t=1}^{T} \nabla_t \cdot (w_t - w^*) &\leq \frac{1}{2\eta}(||w_1 - w^*||^2 - ||w_{T+1} - w^*||^2) + \frac{\eta}{2} G_2^2 T \\ &\leq \frac{1}{2\eta} D_2^2 + \frac{\eta}{2} G_2^2 T \end{aligned}$$

which completes the proof. □