Mehryar Mohri
Advanced Machine Learning 2021
Courant Institute of Mathematical Sciences
Homework assignment 1
March 09, 2021
Due: March 23, 2021

A. Online-to-batch conversion

Let $\mathcal{H}$ be a finite hypothesis set of functions mapping from $\mathcal{X}$ to $\mathbb{R}$ and $\ell\colon \mathbb{R} \times \mathcal{Y} \to \mathbb{R}_+$ a convex function bounded by M, convex with respect to its first argument. Let $\mathcal{A}$ be an online learning algorithm that at each round returns a probability distribution p_t over $\mathcal{H}$. The goal of this problem is to study an online-to-batch conversion from these probability distributions into a randomized algorithm.

Let $\mathcal{P}$ be the set of suffixes $\mathcal{P}_t$: $\mathcal{P}_t = \{\mathsf{p}_t, \ldots, \mathsf{p}_T\}$, $t = 1, \ldots, T$. Fix $\delta > 0$. For each $\mathcal{P} \in \mathcal{P}$, we define:

$$\Gamma(\mathcal{P}) = \frac{1}{|\mathcal{P}|} \sum_{\mathsf{p}_t \in \mathcal{P}} \sum_{h \in \mathcal{H}} \mathsf{p}_t(h) \ell(h(x_t), y_t) + M \sqrt{\frac{\log \frac{T}{\delta}}{|\mathcal{P}|}}.$$

The online-to-batch conversion is done in two steps: first, a distribution $\mathcal{P}_\delta$ is selected via $\mathcal{P}_\delta \in \operatorname{argmin}_{\mathcal{P} \in \mathcal{P}} \Gamma(\mathcal{P})$; next, a randomized algorithm is defined via the distribution p over $\mathcal{H}$ defined for any $h \in \mathcal{H}$ by:

$$\mathsf{p}(h) = \frac{1}{|\mathcal{P}_\delta|} \sum_{\mathsf{p}_t \in \mathcal{P}_\delta} \mathsf{p}_t(h).$$

Let h_{rand} be the randomized hypothesis thereby defined.

1. Show that for any $\delta > 0$, with probability at least $1 - \delta$ over the draw of an i.i.d. sample $S = ((x_1, y_1), \ldots, (x_T, y_T))$ from $\mathcal{D}$, the following inequality holds:

$$\mathbb{E}[\ell(h_{\text{rand}}(x, y))] \leq \frac{1}{T} \sum_{t=1}^{T} \mathop{\mathbb{E}}_{h \sim \mathsf{p}_t} [\ell(h(x_t), y_t)] + M \sqrt{\frac{\log \frac{T}{\delta}}{T}}.$$

   Hint: you can apply Azuma's inequality to an appropriately chosen martingale sequence.

2. Let R_T denote the expected regret of the online algorithm $\mathcal{A}$. Then, show that for any $\delta > 0$, with probability at least $1 - \delta$ over the draw of an i.i.d. sample $S = ((x_1, y_1), \ldots, (x_T, y_T))$ from $\mathcal{D}$, the following inequality holds:

$$\mathbb{E}[\ell(h_{\text{rand}}(x, y))] \leq \inf_{h \in \mathcal{H}} \mathbb{E}[\ell(h(x), y)] + \frac{R_T}{T} + 2M\sqrt{\frac{\log \frac{2T}{\delta}}{T}}.$$

B. Mirror Descent

The notation and definitions used are those adopted in lectures.

1. Prove that Mirror Descent coincides with EG when the convex set is the simplex and the unnormalized relative entropy is used as a Bregman divergence. In particular, you should show that the corresponding mirror map Φ is 1-strongly convex with respect to $\|\cdot\|_1$ on the simplex.

2. Consider the scenario where the functions f_t are differentiable and where, when requesting the gradient $\nabla f_t(\mathbf{w})$ of f_t at $\mathbf{w}$, the learner receives only a random variable $g_t(\mathbf{w})$, such that $\mathbb{E}[g_t(\mathbf{w})] = \nabla f_t(\mathbf{w})$. When $\mathbf{w}_t$ itself is a random variable, we have $\mathbb{E}[g_t(\mathbf{w}_t) | \mathbf{w}_t] = \nabla f_t(\mathbf{w}_t)$. Show that MD in this scenario benefits from the following guarantee:

$$\mathbb{E}[R_T(\text{MD})] \leq \frac{\mathsf{B}(\mathbf{w}^* \parallel \mathbf{w}_1)}{\eta} + \frac{\eta\, \mathbb{E}\big[\|g_t(\mathbf{w}_t)\|_*^2\big]}{2\alpha},$$

   and that for an appropriate choice of η, we have

$$\mathbb{E}[R_T(\text{MD})] \leq DG_* \sqrt{\frac{2T}{\alpha}},$$

   when $\mathsf{B}(\mathbf{w}^* \parallel \mathbf{w}_1) \leq D^2$ and $\mathbb{E}\big[\|g_t(\mathbf{w}_t)\|_*^2\big] \leq G_*^2$.

3. In this question, we adopt the same assumptions as in the previous one except: the functions f_t are all equal to f, f is assumed to be convex and β-smooth with respect to $\|\cdot\|$ and, instead of the upper bound $\mathbb{E}\big[\|g(\mathbf{w}_t)\|_*^2\big] \leq G_*^2$, we will assume that the following bound on the variance holds for all $\mathbf{w}$:

$$\mathbb{E}\big[\|\nabla f(\mathbf{w}) - g(\mathbf{w})\|_*^2\big] \leq \sigma^2.$$

(a) Show that the following inequality holds:

$$\sum_{t=1}^{T} f(\mathbf{w}_{t+1}) - f(\mathbf{w}^*)$$
$$\leq \sum_{t=1}^{T} \nabla f(\mathbf{w}_t) \cdot (\mathbf{w}_{t+1} - \mathbf{w}_t) + \frac{\beta}{2}\|\mathbf{w}_{t+1} - \mathbf{w}_t\|^2 + \nabla f(\mathbf{w}_t) \cdot (\mathbf{w}_t - \mathbf{w}^*).$$

(b) Prove the identity $2\mathbf{u} \cdot \mathbf{v} \leq \mu\|\mathbf{u}\|_*^2 + \|\mathbf{v}\|^2/\mu$ valid for any $\mu > 0$ and vectors $\mathbf{u}$ and $\mathbf{v}$. Use that to show the following:

$$\sum_{t=1}^{T} f(\mathbf{w}_{t+1}) - f(\mathbf{w}^*)$$
$$\leq \sum_{t=1}^{T} g(\mathbf{w}_t) \cdot (\mathbf{w}_{t+1} - \mathbf{w}_t) + \nabla f(\mathbf{w}_t) \cdot (\mathbf{w}_t - \mathbf{w}^*)$$
$$+ \sum_{t=1}^{T} \frac{\eta}{2}\left[\|\nabla f(\mathbf{w}_t) - g(\mathbf{w}_t)\|_*^2\right] + \frac{\beta + 1/\eta}{2}\|\mathbf{w}_{t+1} - \mathbf{w}_t\|^2.$$

(c) Use the 1-strong convexity of Φ to show the following:

$$\sum_{t=1}^{T} f(\mathbf{w}_{t+1}) - f(\mathbf{w}^*)$$
$$\leq \sum_{t=1}^{T} g(\mathbf{w}_t) \cdot (\mathbf{w}_{t+1} - \mathbf{w}^*) + [\nabla f(\mathbf{w}_t) - g(\mathbf{w}_t)] \cdot (\mathbf{w}_t - \mathbf{w}^*)$$
$$+ \sum_{t=1}^{T} \frac{\eta}{2}\left[\|\nabla f(\mathbf{w}_t) - g(\mathbf{w}_t)\|_*^2\right] + (\beta + 1/\eta)\mathsf{B}(\mathbf{w}_{t+1} \parallel \mathbf{w}_t).$$

(d) Prove the following inequality:

$$[\nabla\Phi(\mathbf{w}_{t+1}) - \nabla\Phi(\mathbf{v}_{t+1})] \cdot (\mathbf{w}_{t+1} - \mathbf{w}^*) \leq 0.$$

(e) Use the previous question to prove:

$$\frac{1}{\beta + \frac{1}{\eta}} g(\mathbf{w}_t) \cdot (\mathbf{w}_{t+1} - \mathbf{w}^*) \leq \mathsf{B}(\mathbf{w}^* \parallel \mathbf{w}_t) - \mathsf{B}(\mathbf{w}^* \parallel \mathbf{w}_{t+1}) - \mathsf{B}(\mathbf{w}_{t+1} \parallel \mathbf{w}_t).$$

(f) Use the previous results to conclude that the following regret bound holds for MD run with step size $\frac{1}{\beta+1/\eta}$, with $\eta = \frac{D}{\sigma}\sqrt{\frac{2}{T}}$:

$$\mathbb{E}[R_T(\text{MD})] \leq \beta D^2 + \sigma D\sqrt{2T}.$$