Lecture 13: May 28

Lecturer: Yu-Xiang Wang *Scribes: Yichen Zhou*

Note: *LaTeX template courtesy of UC Berkeley EECS dept.*

This lecture's notes illustrate some uses of various LaTeX macros. Take a look at this and imitate.

13.1 Follow the Leader (FTL) Algorithm

In the OCO setting of regret minimization, the most intuitive way is to use at any time the optimal decision in hindsight. Let

$$\mathbf{x}_{t+1} = \arg\min_{\mathbf{x}\in\mathcal{K}} \sum_{\tau=1}^{t} f_\tau(\mathbf{x}). \tag{13.1}$$

The regret of FTL can be linear in some cases.

Example 13.1.1 Consider $\mathcal{K} \in [-1, 1], f_1(x) = \frac{1}{2}x, f_2(x) = -x, f_3(x) = x, f_4(x) = -x, f_5(x) = x, \cdots, f_{2n} = -x, f_{2n+1} = x.$
Thus

$$\sum_{\tau=1}^{t} f_\tau(\mathbf{x}) = \begin{cases} -\frac{1}{2}x & \textit{if } t \textit{ is even;} \\ \frac{1}{2}x & \textit{if } t \textit{ is odd.} \end{cases}$$

FTL will output 0,-1,1,-1,1,$\cdots$, then

$$regret_T \le 0 + 1 + 1 + 1 + \cdots - \sum_{t=1}^{T} f_t(0) = T - 1 = O(T) \tag{13.2}$$

In order to induce stability, FTL can be modified by either randomization or regularization. By adding a regularization term, we can obtain the RFTL (Regularized Follow the Leader) algorithm.

13.2 Regularized Follow the Leader(RFTL) Algorithm

13.2.1 Regularization functions

Definition 13.1 Regularization function *R(x): K*$\mapsto \mathbb{R}.R(x)$ is twice differentiable, smooth and strong convex on $\mathcal{K}$, and usually non-negative.

Definition 13.2 The diameter of the set K relative to the function R is denoted as $D_R = \sqrt{\max_{\mathbf{x},\mathbf{y}\in\mathcal{K}}\{R(\mathbf{x}) - R(\mathbf{y})\}}$

Definition 13.3 Dual norm: $\|\mathbf{y}\|^* = \max_{\|\mathbf{x}\|\le 1}\langle \mathbf{x}, \mathbf{y}\rangle$

Definition 13.4 Matrix norm: $\|\mathbf{x}\|_A = \sqrt{\mathbf{x}^\top A\mathbf{x}}$

Example 13.2.1 $\|\mathbf{x}\|_{\mathbf{y}} = \|\mathbf{x}\|_{\nabla^2 R(\mathbf{y})} = \sqrt{x^\top \nabla^2 R(\mathbf{y})\mathbf{x}}$, $\|\mathbf{x}\|_{\mathbf{y}}^* = \|\mathbf{x}\|_{\nabla^{-2} R(\mathbf{y})} = \sqrt{x^\top \nabla^{-2} R(\mathbf{y})\mathbf{x}}$

Definition 13.5 Bregman divergence: $B_R(\mathbf{x}\|\mathbf{y}) = R(\mathbf{x}) - R(\mathbf{y}) - \nabla R(\mathbf{y})^\top(\mathbf{x}-\mathbf{y}) = \frac{1}{2}\|\mathbf{x}-\mathbf{y}\|_{\mathbf{z}}^2$, $\forall z \in [x, y]$

Example 13.2.2 $B_R(\mathbf{x_t}\|\mathbf{x_{t+1}}) = \frac{1}{2}\|\mathbf{x_t} - \mathbf{x_{t+1}}\|_{\mathbf{x_t}}^2$

Example 13.2.3 $R(x) = \frac{1}{2}\|x\|^2$, $B_R(x\|t) = \frac{1}{2}\|x-t\|^2$

Example 13.2.4 $R(x) =< x, logx >= \sum_i x_i log x_i$, $B_R(x\|t) = \sum_i x_i log\frac{x_i}{y_i} = KL(x\|y)$

13.2.2 RFTL Algorithm

Algorithm 1 Regularized Follow the Leader(RFTL)

Input: $y, R, \mathcal{K}$
Let $x_1 = \arg\min_{\mathbf{x}\in\mathcal{K}} R(x)$
for $t = 1, 2, \cdots, T$ **do**
  Predict x_t
  Observe $\nabla_t = \nabla f_t(\mathbf{x}_t)$
  Update $\mathbf{x}_{t+1} = \arg\min_{\mathbf{x}\in\mathcal{K}} \left\{\eta \sum_{\tau=1}^t \nabla_\tau^\top \mathbf{x} + R(\mathbf{x})\right\}$
end for

Theorem 13.6 RFTL Algorithm attains for every $u \in \mathcal{K}$ the following bound on the regret:

$$\text{regret}_T \leq 2\eta \sum_{t=1}^T \|\nabla_t\|_t^{*2} + \frac{R(\mathbf{u}) - R(\mathbf{x}_1)}{\eta}. \tag{13.3}$$

If we know the upper bound on the local norms such as $\|\nabla_t\|_t^* \leq G_R$, we can obtain the bound on the regret:

$$\text{regret}_T \leq 2D_R G_R \sqrt{2T} \tag{13.4}$$

Lemma 13.7 $\forall \mathbf{u} \in \mathcal{K}$, RFTL algorithm guarantees the following regret bound:

$$\text{regret}_T \leq \sum_{t=1}^T \nabla_t^\top \mathbf{x}_t - \sum_{t=1}^T \nabla_t^\top \mathbf{x}_{t+1} + \frac{1}{\eta} D_R^2 \tag{13.5}$$

Proof:
$\because$ Define $g_0(\mathbf{x}) = \frac{1}{\eta} R(\mathbf{x}), g_t(\mathbf{x}) = \nabla_t^\top \mathbf{x}$
$\therefore regret_T = \sum_{t=1}^T (g_t(x_t) - g_t(u)), \forall \mathbf{u} \in \mathcal{K}$
$\because \sum_{t=1}^T (g_t(x_t) - g_t(u)) = \sum_{t=1}^T (g_t(x_t) - g_t(x_{t+1}) + g_t(x_{t+1}) - g_t(u)) = \sum_{t=1}^T (g_t(x_t) - g_t(x_{t+1})) + \sum_{t=0}^T (g_t(x_{t+1}) - g_t(u)) - g_0(x_1) + g_0(u)$
$\because \sum_{t=0}^T (g_t(x_{t+1}) - g_t(u)) \leq 0$ and $-g_0(x_1) + g_0(u) \leq \frac{1}{\eta} D_R^2$
$\therefore regret_T = \sum_{t=1}^T (g_t(x_t) - g_t(u)) \leq \sum_{t=1}^T (g_t(x_t) - g_t(x_{t+1})) + \frac{1}{\eta} D_R^2 = \sum_{t=1}^T (\nabla_t^\top \mathbf{x}_t - \nabla_t^\top \mathbf{x}_{t+1})) + \frac{1}{\eta} D_R^2$ ∎