Lecture 19 Adaptive online learning

Example ADMM: $\min_{\theta} \|\theta - y\|^2 + \lambda \|D^{(k+1)}\theta\|_1$

$\min \|\theta - y\|^2 + \lambda\|D^{(1)} z\|_1$

s.t. $D^{(k)}\theta = z$

$L = \|\theta - y\|^2 + \lambda\|D^{(1)}z\|_1 + u^T(D^{(k)}\theta - z) + \frac{\rho}{2}\|D^{(k)}\theta - z\|^2$

$[I]$

find

① $\arg\min_{\theta} L(\theta, z, u) = \arg\min_{\theta} \theta^T\left(\frac{\rho}{2}D^{(k)T}D^{(k)} + I\right)\theta + b^T\theta$

$A\theta = b$

in $O(n)$

$O(k)$

~~$O(n^3)$~~

② $\arg\min_{z} L(\theta, z, u) = \mathrm{Prox}_{\|D^{(1)}\cdot\|_1}(b)$

③ $u^+ = u + \rho(D^{(k)}\theta - z)$

Fused Lasso, TV denoising

DP alg in $O(n)$

Recap: OCO.

For $t = 1, 2, 3, \dots, T$

- player chooses $x_t \in K$
- adversary chooses $f_t \in \mathcal{F}$
- player incur loss $f_t(x_t)$
- player receive feedback

$$\left\{ \begin{array}{l} \nabla f_t(x_t) \leftarrow \text{full info setting} \\ f_t(x_t) \leftarrow \text{Bandit setting} \\ \nabla f_t(x_t) + z_t \leftarrow \end{array} \right.$$

$$\text{Regret} = \sum_{t=1}^{T} f_t(x_t) - \min_x \sum_{t=1}^{T} f_t(x)$$

"No Regret" $\iff \lim_{T\to\infty} \frac{\text{Regret}(T)}{T} = 0$

$$\text{Static Regret}(u) = \sum_{t=1}^{T} f_t(x_t) - \sum_{t=1}^{T} f_t(u)$$

(Zinkevich (2003))

Example 1. $f_t(x) = \left(x - \frac{t}{T}\right)^2$ $\quad K = [0,1]$

$$\min_x \sum_{t=1}^{T} f_t(x) = \sum_{t=1}^{T} \left(x - \frac{t}{T}\right)^2$$

$$\nabla = \sum_{t=1}^{T} 2\left(x - \frac{t}{T}\right) = 0$$

$$x = \frac{1}{2T} \frac{2(1+2+3+\dots+T)}{T} = \frac{2 \cdot \frac{1}{2} T(T+1)}{2T^2} = \frac{T+1}{2T} \approx \frac{1}{2}$$

$$\sum_{t=1}^{T} \left(\frac{T+1}{2T} - \frac{t}{T}\right)^2 = \frac{\left(\frac{T+1}{2} - t\right)^2}{T^2} \approx O(T)$$

$$\int_1^T \left(\frac{T+1}{2} - t\right)^2 dt = \left[\frac{1}{3}\left(t - \frac{T+1}{2}\right)^3\right]_1^T \approx T^3$$

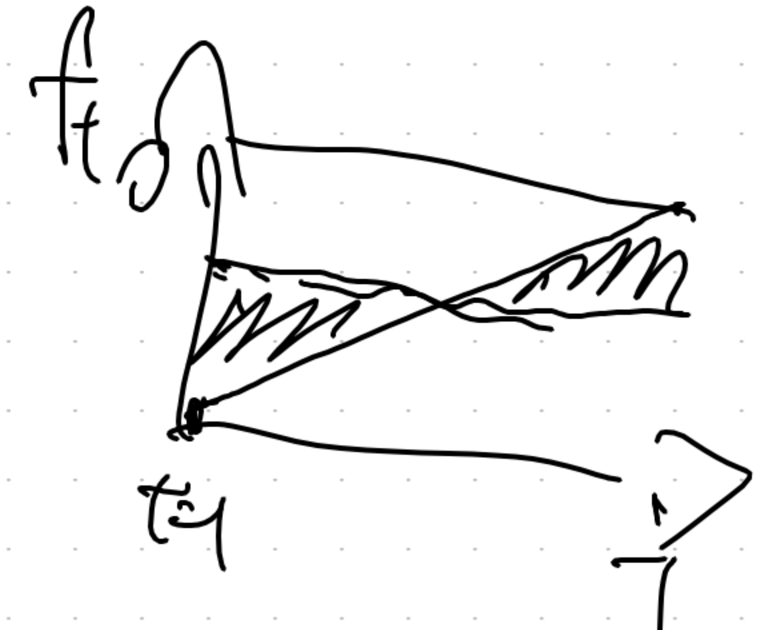

Regret $\approx \log(T)$

objective $\approx O(T)$

Can we do better?

"Dynamic Regret" compete against an arbitrary sequence of comparators in the hindsight

$$\text{D.Regret} = \sum_{t=1}^{T} f_t(x_t) - \min_{(u_1 \cdots u_T)} \sum_{t=1}^{T} f_t(u_t)$$

Example: $f_t(x) = \begin{cases} (x-1)^2 & \text{w.p. } \frac{1}{2} \\ (x+1)^2 & \text{w.p. } \frac{1}{2} \end{cases}$ $\leftarrow$

Best dynamic comparator in the hindsight, $f_1 \cdots f_T$

$u_t = \arg\min_x f_t(x) \quad \min_x f_t(x) = 0$

$$\min_{u_1 \cdots u_T} \sum f_t(u_t) = 0$$

Take time t. $\min_x \mathbb{E}[f_t(x)] = \min_x \frac{1}{2}(x-1)^2 + \frac{1}{2}(x+1)^2 \underset{x=0}{=} 1$

D.Regret of any player $= T$

What to do?

1. Restrict the comparator class $u_1 \cdots u_T$

$P_T(u_1 \cdots u_T)$ $\quad \left\{ u_1 \cdots u_T \;\middle|\; \boxed{\sum_{t=2}^{T} \|u_t - u_{t-1}\|_2 \le P_T} \right\}$ path constraint

Zinkevich (2003)

Zhang et al. (2018) ICML

full info feedback

(Besbes et al. (2013) Operations Research)

(Chen, Wang, W., 2019)

2. Make assumptions about $f_1, f_2, \cdots, f_T$

$f_1, \cdots f_T$ change slowly

e.g. $\sum_{t=2}^{T} \|f_{t-1} - f_t\| \le V_T$ $\quad \left\| \begin{matrix} \|f_1\|_\infty \\ \vdots \\ \|f_T\|_\infty \end{matrix} \right\|_{1}$ $\quad \left\| \begin{matrix} \|f_1\| \\ \vdots \\ \|f_T\| \end{matrix} \right\|_{p,q}$

$p \ge 1 \quad q \ge 1$

$$\boxed{\sum_{t=2}^{T} \sup_x |f_{t-1}(x) - f_t(x)| \le V_T}$$

$\leftarrow$ noisy gradient feedback

$x_t^* = \arg\min_x f_t(x)$

$x_1^*, x_2^*, \cdots, x_T^*$ change slowly $\quad \sum_{t=2}^{T} \|x_{t-1} - x_t\|_p^q \le V_T(p,q) \leftarrow$ (Yang et al. 2016) NeurIPS

$$\text{Regret}(T, (u_1, \cdots, u_T)) = \sum_t f_t(x_t) - \sum_t f_t(u_t) \le \text{function}(T, P_T(u_1 \cdots u_T))$$

$$P_T = \sum_{t=2}^{T} \|u_t - u_{t-1}\|_2$$

Online gradient Descent (OGD)

$$x_{t+1} = \text{Proj}_K (x_t - \eta_t \nabla f_t(x_t))$$

S. Regret $\le O(G \cdot D \cdot \sqrt{T})$

$G \ge \|\nabla f_t(x)\|_2$, f_t is G-Lipschitz, Convex

$D \ge \|u_1 - u_2\|_2$ $\forall u_1, u_2 \in K$

Thm (Thm 2, Section 2.2 of Zinkevich (2003))

$\eta_t = \eta$ then

$$\text{D. Regret}(T, P_T) \le \frac{D^2}{2\eta} + \frac{P_T D}{\eta} + \frac{T\eta G^2}{2}$$

$$= \sqrt{T(D^2 + P_T D) G^2}$$

Proof:

(1) Convexity of f_t, $\forall u \in K$

$$f_t(u) \ge f_t(x_t) + \langle g_t, u - x_t \rangle$$

$$f_t(x_t) - f_t(u) \le \langle g_t, x_t - \underset{\uparrow\, u_t}{u} \rangle$$

(2) By OGD alg.

$$\|x_{t+1} - u\|^2 = \|\Pi_K(x_t - \eta g_t) - u\|^2$$

$$\underset{u \in K}{\le} \|x_t - \eta g_t - u\|_2^2$$

$$= \|x_t - u\|^2 + \eta^2 \|g_t\|^2 - 2\eta \langle g_t, x - u \rangle$$

Take $u = u_t$

$$\|x_{t+1} - u_t\|^2 \le \|x_t - u_t\|^2 + \eta^2 \|g_t\|^2 + 2\eta \langle g_t, x_t - u_t \rangle$$

(3) $\|x_{t+1}-u_t\|^2 = \|x_{t+1}-u_{t+1}+u_{t+1}-u_t\|^2 = \underbrace{\|x_{t+1}-u_{t+1}\|^2 + \|u_{t+1}-u_t\|^2 + 2\langle x-u_{t+1}, u_{t+1}-u_t\rangle}$

$$f_t(x_t)-f_t(u_t) \overset{\text{Step(1)}}{\le} \langle g_t, x_t-u_t\rangle$$

$$\overset{\text{Step(2)}}{\le} \frac{1}{2\eta}\left(\|x_t-u_t\|^2+\eta^2\|g_t\|^2 - \underbrace{\|x_{t+1}-u_t\|_2^2}\right)$$

$$\le \frac{1}{2\eta}\left(\|x_t-u_t\|^2 - \|x_{t+1}-u_{t+1}\|^2 - \underbrace{\|u_{t+1}-u_t\|^2}_{\le 0} - 2\langle x_{t+1}-u_{t+1}, u_{t+1}-u_t\rangle + \underbrace{\eta^2\|g_t\|^2}_{\le \eta^2 G^2}\right)$$

$$\ge -\overbrace{\|x_{t+1}-u_{t+1}\|_2}^{\le D}\|u_{t+1}-u_t\|_2$$

$$\le \frac{1}{2\eta}\left(\|x_t-u_t\|^2 - \|x_{t+1}-u_{t+1}\|^2 + 2D\cdot\|u_{t+1}-u_t\|_2 + \eta^2 G^2\right)$$

Sum up $t=1,\dots,T$, telescope

$$\sum f_t(x_T) - f_t(u_T) \le \frac{1}{2\eta}\left(\|x_1-u_1\|^2 - \underbrace{\|x_{T+1}-u_{T+1}\|^2}_{\le 0} + 2D\sum_{t=2}^{T}\|u_{t+1}-u_t\|_2 + T\eta^2 G^2\right)$$

$$= \frac{D^2}{2\eta} + \frac{\eta G^2 T}{2} + \frac{2DP_T}{2\eta} = \sqrt{TG^2(D^2+DP_T)}$$

choose $\eta = \frac{\sqrt{D^2+DP_T}}{\sqrt{T}\,G}$ □

- This bound is optimal in T, P_T but the OGD cannot compete against all $u_1 \sim u_T$ simultaneously

Adaptive Alg P_T is not an input (Zhang et. al. 2018)

adaptive in P_T, Ader, Hedge

- $P_T = 1$ you get Static-Regret

$P_T = T$ (Example 2) D-Regret $= T$

Non function Variations, Noisy Gradient Feedback, $\sum_{t=2}^{T} \sup_x |f_t(x) - f_{t-1}(x)| \leq V_T$

Feedback is $g_t = \nabla f_t(x_t) + Z_t$ $\quad Z_t$ is independent subgaussian noise

D.Regret: $\mathbb{E}\left[\sum_{t=1}^{T} f_t(x_t)\right] - \sum_{t=1}^{T} f_t(x_t^*)$, $x_t^* = \arg\min_x f_t(x)$

$\leq \text{function}(T, V_T)$

$\boxed{\max_{u_1 \sim u_T} P_T}$

$(|x_1^* - x_2^*| + \cdots + |x_{T-1}^* - x_T^*|)$

OGD with Restarting Schedule

Alg: A: any OCO alg with Static Regret Bound

Δ_T Batch Size

$\Delta_T, \Delta_T, \cdots, \Delta_T \quad \left\lceil \frac{T}{\Delta_T} \right\rceil$

OGD OGD

T

Thm. (Prop 2 in Besbes et.al.)

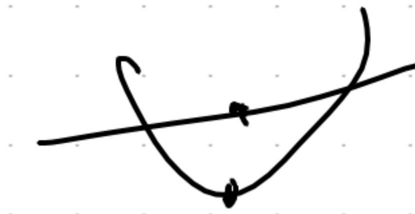

$$\underline{\text{Regret}(T, V_T) \le \left\lceil \frac{T}{\Delta_T} \right\rceil \cdot \text{Regret}_A(\Delta_T) + 2\Delta_T \cdot V_T}$$

Proof. $\sum_{t=1}^{T} \left(f_t(x_t) - f_t(x_t^*)\right)$ $\qquad \underline{x_j^*} = \arg\min \sum_{t_j=1}^{\Delta_T} f_{t_j}(x)$

$$= \sum_{j=1}^{T/\Delta_T} \sum_{t_j=1}^{\Delta_T} \underline{f_{t_j}(x_{t_j}) - f_{t_j}(x_j^*)} + f_{t_j}(x_j^*) - f_{t_j}(x_{t_j}^*)$$

$$\le \frac{T}{\Delta_T} \cdot \text{Regret}_A(\Delta_T) + \underbrace{\sum_{j=1}^{T/\Delta_T} \sum_{t_j=1}^{\Delta_T} \left(f_{t_j}(x_j^*) - f_{t_j}(x_{t_j}^*)\right)}_{(*)}$$

$V_T = \sum_j V_T^{(j)}$

$$(*) = \sum_j \Delta_T \left(\underline{\frac{1}{\Delta_T} \sum_{t_j=1}^{\Delta_T} f_{t_j}}\right)(x_j^*) - \sum_{t_j=1}^{\Delta_T} f_{t_j}(x_{t_j}^*)$$

for $\forall\, i_1, i_2 \in j$th Bin $\quad \|f_{i_1} - f_{i_2}\|_\infty \le V_T^{(j)}$

$\cup$

f_j

$$\le \sum_j \Delta_T \left(\underline{\frac{1}{\Delta_T} \sum_{t_j=1}^{\Delta_T} f_{t_j}}\right)\left(\frac{1}{\Delta_T} \sum x_{t_j}^*\right) - \sum_{t_j=1}^{\Delta_T} f_{t_j}(x_{t_j}^*) \le \sum_j \Delta_T \frac{1}{\Delta_T} \sum_{t_j=1}^{\Delta_T} \left(\left(\underline{\frac{1}{\Delta_T} \sum_{t_j=1}^{\Delta_T} f_{t_j}}\right)(x_{t_j}^*) - f_{t_j}(x_{t_j}^*)\right)$$

x_j^* is optimal ☺

$$(*) \le \sum_j \Delta_T V_T^{(j)} \le \Delta_T V_T$$

$$\text{D.regret} \lesssim \frac{T}{\Delta_T} \cdot \text{Regret}_A(\Delta_T) + \Delta_T V_T$$

If OGD Convex $\quad \frac{T}{\Delta_T}\sqrt{\Delta_T}\, D\sqrt{G^2+\sigma^2} + \Delta_T V_T \asymp T^{\frac{2}{3}} V_T^{\frac{1}{3}}$

If OGD Strongly Convex $\quad \frac{T}{\Delta_T} \frac{\log(1+T) \cdot G^2}{m} + \Delta_T V_T \asymp \sqrt{\frac{T \log T}{m} V_T}\, G$

$$\asymp \sqrt{T V_T}$$

Example 1 : $f_t(x) = (x - \frac{t}{T})^2 = (x-\theta_t)^2$

$g_t = 2(x_t - \theta_t) + z_t$

you know x_t, $\quad y_t = \frac{g_t}{2} + x_t = \theta_t + \text{Noise}$

$$E \sum_t (A_t(y_1 \cdots y_T) - \theta_t)^2 = \text{MSE}$$

θ_t changes slowly

$\sum_t |\theta_t - \theta_{t-1}| \le P_T$ $\quad \vec{\theta} \in \underline{TV}(P_T)$

$\sqrt{\sum_t |\theta_t - \theta_{t-1}|^2} \le P_T$ $\quad \vec{\theta} \in \underline{\text{Sobolev}}(P_T)$

$\left(\sum_t (\theta_t - \theta_{t-1})^p\right)^{\frac{1}{p}} \le P_T$ $\quad p \to \infty$ $\quad \vec{\theta} \in \underline{\text{Holder class}}(P_T)$

Optimal (Offline) Rates $O\left(T^{\frac{1}{3}} P_T^{\frac{2}{3}}\right)$

$O\left(T^{\frac{1}{2}} V_T^{\frac{1}{2}}\right)$

$O\left(T^{\frac{1}{2}} P_T^{\frac{1}{2}}\right)$

Baby and W. (2019)

$\nabla \left(y_t - \theta x_t\right)^2 \Leftrightarrow y_t = \theta_t + \text{Noise}$

Strongly Adaptive Regret (Daniely et al., 2015) Lower bound for Linear Smoother

SA-Regret(τ) = $\max\limits_{I \in [q, q+\tau] \subset [T]}$ Regret$_I(\tau)$