Reading Report on "Online Optimization with Gradual Variations", the Best Student Paper from 12'COLT, by C. Chiang & T. Yang, et al.

Mengfei Cao

April 16

Introduction

- An example of online convex optimization model

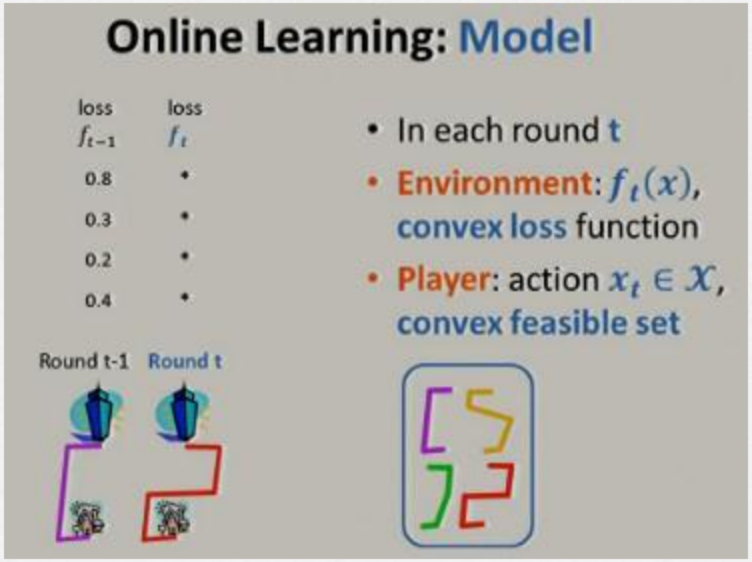

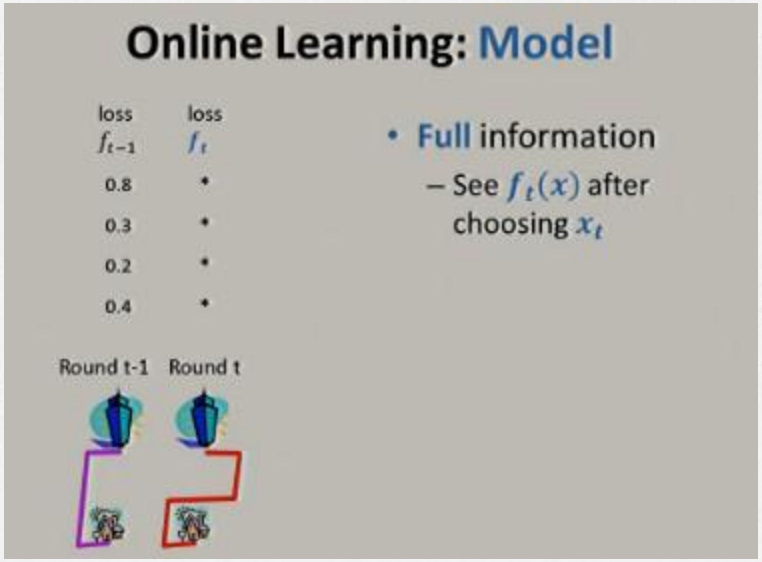

Introduction

- An example of online convex optimization model

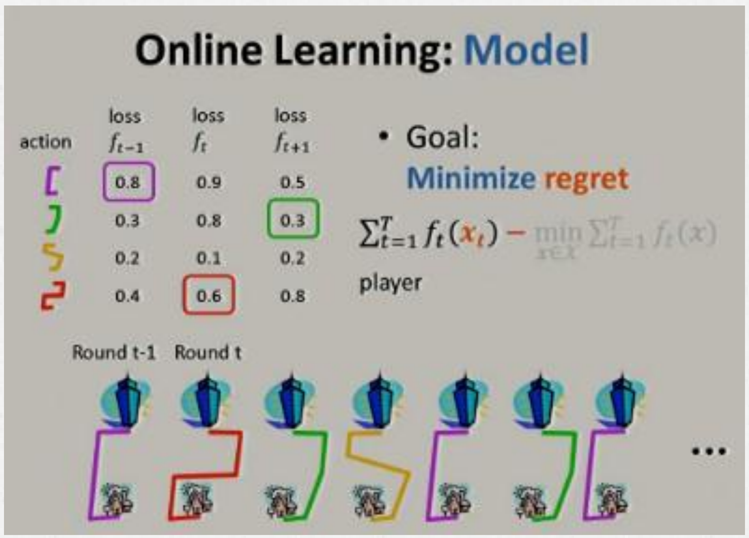

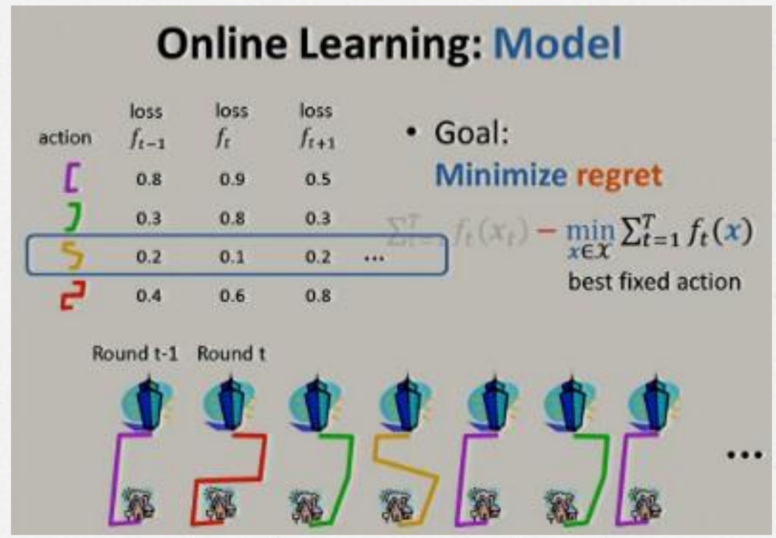

Introduction

- The General Online Convex Optimization Model

Online Convex Optimization (OCO)

Input: A convex set X
for t = 1,2,...
  predict a vector x_t from X
  receive a convex loss function $f_t: X \to R$
  suffer loss $f_t(x_t)$

Minimize: *regret* = $\sum_{t=1}^{T} f_t(x_t) - \min_{\pi \in S} \sum_{t=1}^{T} f_t(\pi)$

Regret Bound: using Gradient Descend Algorithm, achieve $O(\sqrt{TN})$, where N is the dimensionality of the convex set.

Scenario No. 1:

- Online Linear Optimization

Online Linear Optimization

Input: A convex set X
for t = 1,2,...
    predict a vector x_t from X
    receive a convex loss function $f_t: X \to R$
    suffer loss $f_t(x_t)$

Minimize: *regret* = $\sum_{t=1}^{T} f_t(x_t) - \min_{\pi \in S} \sum_{t=1}^{T} f_t(\pi)$

$$f_t(x_t) = < x_t, z_t >, for\ some\ z_t$$

Regret Bound: using Gradient Descend Algorithm, achieve $O(\sqrt{TN})$, where *N* is the dimension of x_t.

Scenario No. 2:

- Prediction with Expert Advice

Prediction with Expert Advice

Input: A convex set X
for t = 1,2,…
  predict a vector x_t from X
  receive a convex loss function $f_t: X \rightarrow R$
  suffer loss $f_t(x_t)$

Minimize: *regret* = $$\sum_{t=1}^{T} f_t(x_t) - \min_{\pi \in S} \sum_{t=1}^{T} f_t(\pi)$$

$f_t(x_t) =< x_t, z_t >, for\ some\ z_t,$ and $\forall x \in X: |x| = 1$

Regret Bound: using Multiplicative Update Algorithm, achieve $O\left(\sqrt{T \ln N}\right)$, where N is the number of experts.

Scenario No. 3:

- Online Convex Optimization*: gradient of each loss function is λ-smooth.

Online Convex Optimization*

Input: A convex set X
for t = 1,2,…
  predict a vector x_t from X
  receive a convex loss function $f_t: X \rightarrow R$
  suffer loss $f_t(x_t)$

Minimize: *regret* = $\sum_{t=1}^{T} f_t(x_t) - \min_{\pi \in S} \sum_{t=1}^{T} f_t(\pi)$

$$\boxed{\|\nabla f_t(x) - \nabla f_t(y)\|_2 \leq \lambda \|x - y\|_2 \,, \ \text{for any } x, y \in \mathcal{X}}$$

Regret Bound: using Gradient Descend Algorithm, achieve $O(\sqrt{TN})$, where *N* is the dimension of x_t.

Scenario No. 4:

- Online Strictly Convex Optimization: $\beta - convex$ loss function

Online Strictly Convex Optimization

Input: A convex set X
for t = 1,2,…
  predict a vector x_t from X
  receive a convex loss function $f_t: X \rightarrow R$
  suffer loss $f_t(x_t)$

Minimize: *regret* = $\sum_{t=1}^{T} f_t(x_t) - \min_{\pi \in S} \sum_{t=1}^{T} f_t(\pi)$

Definition 3 *For $\beta > 0$, we say that a function $f : \mathcal{X} \rightarrow \mathbb{R}$ is β-convex, if for all $x, y \in \mathcal{X}$,*

$$f(x) \geq f(y) + \langle \nabla f(y), x - y \rangle + \beta \langle \nabla f(y), x - y \rangle^2 .$$

Regret Bound: using online Newton Step Algorithm, achieve $O(NlogT)$, where *N* is the dimension of x_t.

Motivation: Variations in cost functions can be small

- What else could the regret bound depend on other than T?
  - better be small so as to be tight and easy to bound
- Variation-bounded regret (Previous work):

Problems	T rounds	Variation V	
Online Linear Optimization	$\sqrt{NT}$ [Zin03]	$\sqrt{V}$ [HK08]	
Prediction with Expert Advice	$\sqrt{T \ln N}$ [LW94]	$\sqrt{V \ln N}$ [HK08]	
Online Convex Optimization	$\sqrt{NT}$ [Zin03]	?	
Online Strictly Convex Optimization	$N \log T$ [HAK07]	$N \log V$ [HK09]	

$$\begin{cases} V = \sum_{t=1}^{T} \|f_t - \mu\|_2^2 \\ \mu = \sum_{t=1}^{T} f_t / T \end{cases}$$

$$\begin{cases} Q = \sum_{t=1}^{T} \|v_t - \mu\|_2^2 \\ \mu = \sum_{t=1}^{T} v_t / T \\ f_t(x) = -\ln \langle v_t, x \rangle \end{cases}$$

Motivation: Variations May Not be Reliable

- Example:

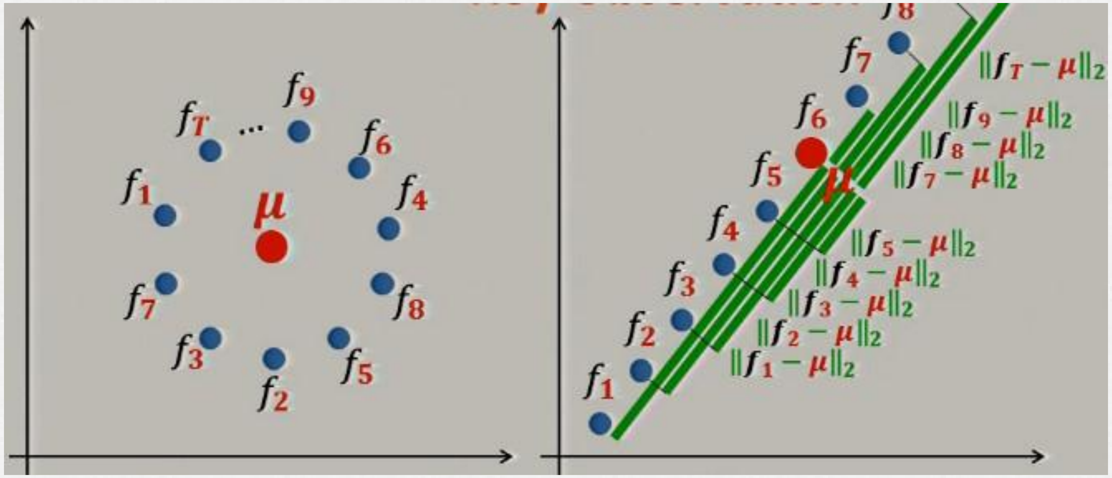

$$V = \sum_{t=1}^{T} \|f_t - \mu\|_2^2$$

In some cases, the consecutive changes between loss functions are smooth and small.

- Deviation-bounded regret (Proposed work):

$$D_p = \sum_{t=1}^{T} \max_{x \in \mathcal{X}} \|\nabla f_t(x) - \nabla f_{t-1}(x)\|_p^2$$

$$D_2 \leq O(V)$$

Motivation: Deviations Can be Very Small

- Example (Linear Online Optimization):

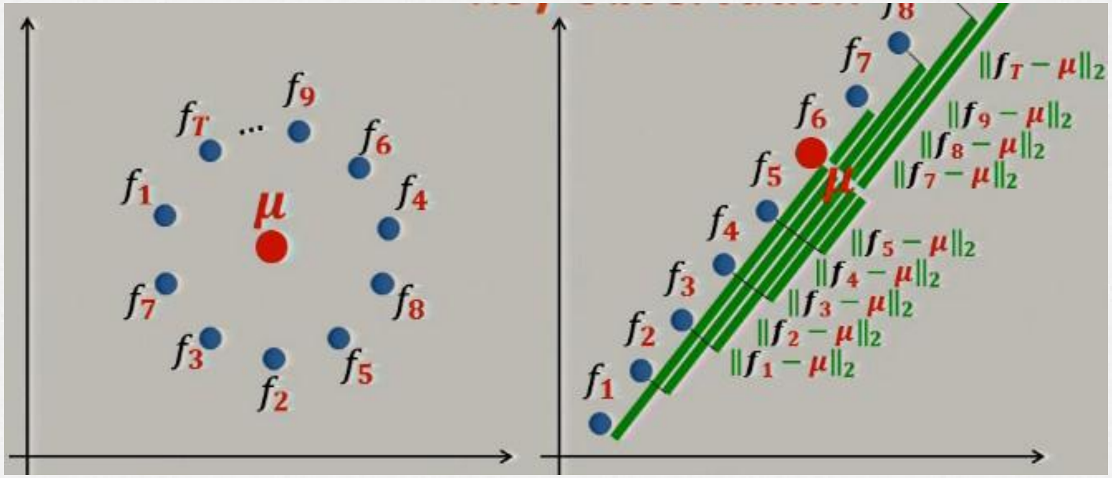

$$\sum_{t=1}^{T} \left(f_t(\hat{x}_t) - f_t(\pi)\right) \leq \eta D_2 + \frac{2}{\eta} \leq O\left(\sqrt{D_2}\right)$$

$$\eta = \sqrt{2/D_2}, \Rightarrow O(\sqrt{D_2})$$

A Unified Algorithm

Algorithm 1 META algorithm

1: Initially, let $x_1 = \hat{x}_1 = (1/N, \ldots, 1/N)^\top$.
2: In round $t \in [T]$:
2(a): Play $\hat{x}_t$.
2(b): Receive f_t and compute $\ell_t = \nabla f_t(\hat{x}_t)$.
2(c): Update

$$x_{t+1} = \arg\min_{x \in \mathcal{X}} \left(\langle \ell_t, x \rangle + \mathcal{B}^{\mathcal{R}_t}(x, x_t) \right),$$
$$\hat{x}_{t+1} = \arg\min_{\hat{x} \in \mathcal{X}} \left(\langle \ell_t, \hat{x} \rangle + \mathcal{B}^{\mathcal{R}_{t+1}}(\hat{x}, x_{t+1}) \right).$$

$$\hat{x}_t = \arg\min_{\hat{x} \in \mathcal{X}} \left(\langle \ell_{t-1}, \hat{x} \rangle + \mathcal{B}^{\mathcal{R}_t}(\hat{x}, x_t) \right)$$

$\mathcal{R} : \mathbb{R}^N \to \mathbb{R}$ *be a differentiable function*

$$\mathcal{B}^{\mathcal{R}}(x, y) = \mathcal{R}(x) - \mathcal{R}(y) - \langle \nabla \mathcal{R}(y), x - y \rangle$$ (Bregman divergence)

Proof for Regret Bound in Scenario No.1 and No.2 (linear)

$$\sum_{t=1}^{T} f_t\left(\hat{x}_t\right) - f_t\left(\pi\right) \le \sum_{t=1}^{T} S_t + \sum_{t=1}^{T} A_t$$

where $S_t = \langle f_t - f_{t-1}, \hat{x}_t - x_{t+1} \rangle$ and $A_t = \mathcal{B}^{\mathcal{R}_t}(\pi, x_t) - \mathcal{B}^{\mathcal{R}_t}(\pi, x_{t+1})$

$$\mathcal{R}_t(x) = \frac{1}{2\eta}\left\|x\right\|_2^2$$

$$\sum_{t=1}^{T} S_t \le \sum_{t=1}^{T} \left\|f_t - f_{t-1}\right\|_*^2 \le \sum_{t=1}^{T} \eta \left\|f_t - f_{t-1}\right\|_2^2 \le \eta D_2$$

$$\sum_{t=1}^{T} A_t = \frac{1}{2\eta}\left(\left\|\pi - x_1\right\|_2^2 - \left\|\pi - x_{T+1}\right\|_2^2\right) \le \frac{2}{\eta}$$

$$\sum_{t=1}^{T} \left(f_t(\hat{x}_t) - f_t(\pi)\right) \le \eta D_2 + \frac{2}{\eta} \le O\left(\sqrt{D_2}\right)$$

by choosing $\eta = \sqrt{2/D_2}$

$$\mathcal{R}_t(x) = \frac{1}{\eta}\sum_{i=1}^{N} x(i)\left(\ln x(i) - 1\right)$$

$$\sum_{t=1}^{T} S_t \le \sum_{t=1}^{T} \left\|f_t - f_{t-1}\right\|_*^2 \le \sum_{t=1}^{T} \eta \left\|f_t - f_{t-1}\right\|_\infty^2 \le \eta D_\infty$$

$$\sum_{t=1}^{T} A_t = \frac{1}{\eta}\left(\mathrm{RE}\left(\pi \| x_1\right) - \mathrm{RE}\left(\pi \| x_{T+1}\right)\right) \le \frac{1}{\eta}\ln N$$

$$\sum_{t=1}^{T} \left(f_t(\hat{x}_t) - f_t(\pi)\right) \le \eta D_\infty + \frac{1}{\eta}\ln N \le O\left(\sqrt{D_\infty \ln N}\right)$$

by choosing $\eta = \sqrt{(\ln N)/D_\infty}$

Proof for Regret Bound in Scenario No.3 (λ-smooth)

$$\sum_{t=1}^{T} \left(f_t\left(\hat{x}_t\right) - f_t\left(\pi\right)\right) \leq \sum_{t=1}^{T} S_t + \sum_{t=1}^{T} A_t - \sum_{t=1}^{T} B_t$$

where $S_t = \langle f_t - f_{t-1}, \hat{x}_t - x_{t+1} \rangle$ and $A_t = \mathcal{B}^{\mathcal{R}_t}(\pi, x_t) - \mathcal{B}^{\mathcal{R}_t}(\pi, x_{t+1})$

$$B_t = \mathcal{B}^{\mathcal{R}_t}(x_{t+1}, \hat{x}_t) + \mathcal{B}^{\mathcal{R}_t}(\hat{x}_t, x_t)$$

$$\mathcal{R}_t(x) = \frac{1}{2\eta} \|x\|_2^2$$

$$\sum_{t=1}^{T} S_t \leq \sum_{t=1}^{T} \|f_t - f_{t-1}\|_*^2 \leq \sum_{t=1}^{T} \eta \|f_t - f_{t-1}\|_2^2$$

$$\sum_{t=1}^{T} A_t = \frac{1}{2\eta}\left(\|\pi - x_1\|_2^2 - \|\pi - x_{T+1}\|_2^2\right) \leq \frac{2}{\eta}$$

$$\textstyle\sum_{t=1}^{T} B_t \geq \frac{1}{4\eta} \sum_{t=1}^{T} \|\hat{x}_t - \hat{x}_{t-1}\|_2^2 - O(1)$$

$$\sum_{t=1}^{T} \left(f_t(\hat{x}_t) - f_t(\pi)\right) \leq O\left(\eta D_2 + \frac{1}{\eta}\right) \leq O\left(\sqrt{D_2}\right)$$

when $\lambda \leq 1/\sqrt{8\eta^2}$ and $\eta = 1/\sqrt{D_2}$.

Proof for Regret Bound in Scenario No.4 ($\beta - convex$, λ-smooth)

$$\sum_{t=1}^{T}\left(f_t\left(\hat{x}_t\right) - f_t\left(\pi\right)\right) \leq \sum_{t=1}^{T} S_t + \sum_{t=1}^{T} A_t - \sum_{t=1}^{T} B_t - \sum_{t=1}^{T} C_t$$

where $S_t = \langle f_t - f_{t-1}, \hat{x}_t - x_{t+1} \rangle$ and $A_t = \frac{1}{2}\left\|\pi - x_t\right\|_{H_t}^2 - \frac{1}{2}\left\|\pi - x_{t+1}\right\|_{H_t}^2$
$B_t = \frac{1}{2}\left\|x_{t+1} - \hat{x}_t\right\|_{H_t}^2 + \frac{1}{2}\left\|\hat{x}_t - x_t\right\|_H^2$ and $C_t = \beta\left\|\pi - \hat{x}_t\right\|_{h_t}^2$ where $h_t = \ell_t\ell_t^\top$

$\mathcal{R}_t(x) = \frac{1}{2}\|x\|_{H_t}^2$, with $H_t = I + \beta\gamma^2 I + \beta\sum_{\tau=1}^{t-1}\ell_\tau\ell_\tau^\top$

$\sum_{t=1}^{T} B_t \geq \frac{1}{4}\sum_{t=1}^{T}\left\|\hat{x}_t - \hat{x}_{t-1}\right\|_2^2 - O(1)$

$$\sum_{t=1}^{T} S_t + \sum_{t=1}^{T} A_t - \sum_{t=1}^{T} C_t \leq O(1+\beta\gamma^2) + \frac{8N}{\beta}\ln\left(1 + \frac{\beta}{4}\sum_{t=1}^{T}\left\|\ell_t - \ell_{t-1}\right\|_2^2\right)$$

$$\sum_{t=1}^{T}\left(f_t(\hat{x}_t) - f_t(\pi)\right) \leq O((N/\beta)\ln D_2)$$

$with\ \lambda \geq 1$ $\beta \leq 1\ and\ D_2 \geq 1$

Conclusion

- By introducing the notion of L_p-deviations, the work derived a tighter bound as proved for four specific online learning scenarios using a unified algorithm META, when the environment follows some stable pattern or the adversary is kind, from the high-level understanding.

Problem	Deviation D_p / Regret	$\mathcal{R}_t$ / Algorithm
Online **Linear** Optimization	$\sum_{t=1}^{T} \|f_t - f_{t-1}\|_2^2$	$\frac{1}{2\eta}\|x\|_2^2$
	$\sqrt{D_2}$	Gradient Descent [Zink03] type
Prediction with **Expert** Advice	$\sum_{t=1}^{T} \|f_t - f_{t-1}\|_\infty^2$	$\frac{1}{\eta}\sum_{i=1}^{N} x(i)(\ln x(i) - 1)$
	$\sqrt{D_\infty \ln N}$	Exp. Weighted Average [CL06] type
Online **Convex** Optimization	$\sum_{t=1}^{T} \max_{x \in \mathcal{X}} \|\nabla f_t(x) - \nabla f_{t-1}(x)\|_2^2$	$\frac{1}{2\eta}\|x\|_2^2$
	$\sqrt{D_2}$	Gradient Descent [Zink03] type
Online **β-convex** Optimization	$\sum_{t=1}^{T} \max_{x \in \mathcal{X}} \|\nabla f_t(x) - \nabla f_{t-1}(x)\|_2^2$	$\frac{1}{2}\|x\|_{H_t}^2$, with $H_t = \lambda I + \beta \sum_{\tau=1}^{t-1} \ell_\tau \ell_\tau^T$
	$N \log D_2$	Online Newton Step [HAK07] type

- Technically, the application of Bregman projections is vital.

Definition 2 *Let* $\mathcal{R} : \mathbb{R}^N \to \mathbb{R}$ *be a differentiable function and* $\mathcal{X} \subseteq \mathbb{R}^N$ *a convex set. Define the Bregman divergence of* $x, y \in \mathbb{R}^N$ *with respect to* $\mathcal{R}$ *by* $\mathcal{B}^{\mathcal{R}}(x, y) = \mathcal{R}(x) - \mathcal{R}(y) - \langle \nabla \mathcal{R}(y), x - y \rangle$. *Define the projection of* $y \in \mathbb{R}^N$ *onto* $\mathcal{X}$ *according to* $\mathcal{B}^{\mathcal{R}}$ *by* $\Pi_{\mathcal{X},\mathcal{R}}(y) = \arg\min_{x \in \mathcal{X}} \mathcal{B}^{\mathcal{R}}(x, y)$.

Theorem 8 *When the L_2-deviation of the loss functions is D_2, the regret of our algorithm is at most $O(\sqrt{D_2})$.*

Theorem 10 *When the L_∞-deviation of the loss functions is D_∞, the regret of our algorithm is at most $O(\sqrt{D_\infty \ln N})$.*

Theorem 11 *When the loss functions have L_2-deviation D_2 and the gradient of each loss function is λ-smooth, with $\lambda \le 1/\sqrt{8D_2}$, the regret of our algorithm is at most $O(\sqrt{D_2})$.*

Theorem 15 *Suppose the loss functions are β-convex and their L_2-deviation is D_2, with $\beta \le 1$ and $D_2 \ge 1$. Furthermore, suppose the gradient of each loss function is λ-smooth, with $\lambda \ge 1$, and has L_2-norm at most γ. Then the regret of our algorithm is at most $O(\beta\gamma^2 + (N/\beta)\ln(\lambda N D_2))$, which becomes $O((N/\beta)\ln D_2)$ for a large enough D_2.*