Unconstrained Online Linear Learning in Hilbert Spaces

Minimax Algorithms and Normal Approximations

H. Brendan McMahan[1] and Francesco Orabona[2]

[1]Google

[2]Toyota Technological Institute at Chicago

Talk @ COLT 2014

Outline

1. Online Learning
2. Minimax Analysis for Unconstrained Games
   - Conditional Value of Unconstrained Games
   - Orthogonal Adversary
   - Parallel Adversary
3. Applications
   - A Power Family of Minimax Algorithms
   - Tight Bounds for Unconstrained Learning

Outline

1. Online Learning
2. Minimax Analysis for Unconstrained Games
   - Conditional Value of Unconstrained Games
   - Orthogonal Adversary
   - Parallel Adversary
3. Applications
   - A Power Family of Minimax Algorithms
   - Tight Bounds for Unconstrained Learning

Online Learning in a Nutshell

- Sequential prediction over rounds $t = 1, 2, \ldots$
- At each round t
  - Learner receives a question $x_t \in \mathbb{X}$
  - Predicts answer $\hat{y}_t$
  - Receives "correct" answer $y_t \in \mathbb{Y}$ and suffers loss $\ell(\hat{y}_t, y_t)$
- Allows x_t and y_t to be generated by an adversary

Online Learning and Regret Bounds

The learner must minimize its loss on the T observed samples

$$\sum_{t=1}^{T} \ell(\hat{y}_t, y_t)$$

Online Learning and Regret Bounds

The learner must minimize its loss on the T observed samples, compared to the loss of the best fixed predictor

$$\sum_{t=1}^{T} \ell(\hat{y}_t, y_t) - \sum_{t=1}^{T} \ell(u(x_t), y_t)$$

where $u \in \mathcal{C}$

Online Learning and Regret Bounds

The learner must minimize its loss on the T observed samples, compared to the loss of the best fixed predictor

$$\sum_{t=1}^{T} \ell(\hat{y}_t, y_t) - \sum_{t=1}^{T} \ell(u(x_t), y_t) = \text{Regret}(u)$$

where $u \in \mathcal{C}$

Online Learning and Regret Bounds

The learner must minimize its loss on the T observed samples, compared to the loss of the best fixed predictor

$$\sum_{t=1}^{T} \ell(\hat{y}_t, y_t) - \sum_{t=1}^{T} \ell(u(x_t), y_t) = \text{Regret}(u)$$

where $u \in \mathcal{C}$

- We want the regret sublinear in T
  - e.g. $\mathcal{O}(\sqrt{T})$

Our Case: Linear losses in Hilbert Spaces

- Let $\mathcal{H}$ be a Hilbert space with inner product $\langle \cdot, \cdot \rangle$.
  - For example, RKHS, space of matrices, etc.
- The associated norm is denoted by $\|x\| = \sqrt{\langle x, x \rangle}$.
- In each round
  - Learner picks $w_t \in \mathcal{H}$
  - Adversary chooses g_t, for simplicity $\|g_t\| \leq 1$
  - Learner gains $\langle w_t, g_t \rangle$
- Minimize the Regret:

$$\text{Regret}(u) \equiv \sum_{t=1}^{T} \langle g_t, w_t - u \rangle$$

where $u \in \mathcal{H}$

Our Results

- General conditions to be able to calculate in a closed formula the optimal strategy w_t, not just up to constants
- We recover and extend prior results
- We close the gap between upper and lower bound for unconstrained online learning
- Relations between potential functions, mirror descent, conditional value of the game, etc.
- Optimal regret bound for unconstrained online learning with unknown T

Our Results

- General conditions to be able to calculate in a closed formula the optimal strategy w_t, not just up to constants
- We recover and extend prior results
- We close the gap between upper and lower bound for unconstrained online learning
- Relations between potential functions, mirror descent, conditional value of the game, etc.
- Optimal regret bound for unconstrained online learning with unknown T

Outline

1. Online Learning
2. Minimax Analysis for Unconstrained Games
   - Conditional Value of Unconstrained Games
   - Orthogonal Adversary
   - Parallel Adversary
3. Applications
   - A Power Family of Minimax Algorithms
   - Tight Bounds for Unconstrained Learning

Minimax View of Online Learning

$$\sup_{\boldsymbol{u} \in \mathcal{W}} \sum_{t=1}^{T} \langle w_t - u, g_t \rangle$$

(Abernethy et al., 2008, 2007; Abernethy and Warmuth, 2010; Abernethy et al., 2008b; Streeter and McMahan, 2012).

Minimax View of Online Learning

$$\min_{w_1 \in \mathcal{H}} \max_{\|g_1\| \le 1} \cdots \min_{w_T \in \mathcal{H}} \max_{\|g_T\| \le 1} \left(\sup_{\boldsymbol{u} \in \mathcal{W}} \sum_{t=1}^{T} \langle w_t - u, g_t \rangle \right)$$

(Abernethy et al., 2008, 2007; Abernethy and Warmuth, 2010; Abernethy et al., 2008b; Streeter and McMahan, 2012).

Minimax View of Online Learning

Value of the game

$$V \equiv \min_{w_1 \in \mathcal{H}} \max_{\|g_1\| \leq 1} \cdots \min_{w_T \in \mathcal{H}} \max_{\|g_T\| \leq 1} \left(\sup_{\boldsymbol{u} \in \mathcal{W}} \sum_{t=1}^{T} \langle w_t - u, g_t \rangle \right)$$

(Abernethy et al., 2008, 2007; Abernethy and Warmuth, 2010; Abernethy et al., 2008b; Streeter and McMahan, 2012).

Minimax View of Online Learning

Value of the game

$$V \equiv \min_{w_1 \in \mathcal{H}} \max_{\|g_1\| \le 1} \cdots \min_{w_T \in \mathcal{H}} \max_{\|g_T\| \le 1} \left(\sup_{\boldsymbol{u} \in \mathcal{W}} \sum_{t=1}^{T} \langle w_t - u, g_t \rangle \right)$$

$$= \min_{w_1 \in \mathcal{H}} \max_{\|g_1\| \le 1} \cdots \min_{w_T \in \mathcal{H}} \max_{\|g_T\| \le 1} \left(\sum_{t=1}^{T} \langle w_t, g_t \rangle + \sup_{u \in \mathcal{W}} \langle u, \theta_T \rangle \right)$$

where $\theta_t = \sum_{i=1}^{t} g_i$.

(Abernethy et al., 2008, 2007; Abernethy and Warmuth, 2010; Abernethy et al., 2008b; Streeter and McMahan, 2012).

Minimax View in Unconstrained Settings

$$V \equiv \min_{w_1 \in \mathcal{H}} \max_{\|g_1\| \leq 1} \cdots \min_{w_T \in \mathcal{H}} \max_{\|g_T\| \leq 1} \left(\sum_{t=1}^{T} \langle w_t, g_t \rangle + \sup_{u \in \mathcal{W}} \langle u, \theta_T \rangle \right)$$

where $\theta_t = \sum_{i=1}^{t} g_i$.

Minimax View in Unconstrained Settings

$$V \equiv \min_{w_1 \in \mathcal{H}} \max_{\|g_1\| \le 1} \cdots \min_{w_T \in \mathcal{H}} \max_{\|g_T\| \le 1} \left(\sum_{t=1}^{T} \langle w_t, g_t \rangle + f(\|\theta_T\|) \right)$$

where $\theta_t = \sum_{i=1}^{t} g_i$ and the Benchmark Function f is increasing and convex.

(McMahan and Abernethy, 2013)

Minimax View in Unconstrained Settings

$$V \equiv \min_{w_1 \in \mathcal{H}} \max_{\|g_1\| \le 1} \cdots \min_{w_T \in \mathcal{H}} \max_{\|g_T\| \le 1} \left(\sum_{t=1}^{T} \langle w_t, g_t \rangle + f(\|\theta_T\|) \right)$$

where $\theta_t = \sum_{i=1}^{t} g_i$ and the Benchmark Function f is increasing and convex.

Theorem

$$\forall u \in \mathcal{H}, \quad \text{Regret}(u) \le f^*(\|u\|) + V$$

where f^* is the Fenchel conjugate of f

Conditional Value of the Game

Recursive definition of the conditional value of the game

$$V_t(\theta_t) = \min_{w \in \mathcal{H}} \max_{\|g\| \leq 1} \langle g, w \rangle + V_{t+1}(\theta_t - g)$$

Conditional Value of the Game

Recursive definition of the conditional value of the game

$$V_t(\theta_t) = \min_{w \in \mathcal{H}} \max_{\|g\| \leq 1} \langle g, w \rangle + V_{t+1}(\theta_t - g)$$

$$\left. \begin{aligned} V_T(\theta_T) &= f(\|\theta_T\|) \\ V &= V_0(\mathbf{0}) \end{aligned} \right\} \Rightarrow \forall u \in \mathcal{H}, \quad \text{Regret}(u) \leq f^*(\|u\|) + V_0(\mathbf{0})$$

Conditional Value of the Game

Recursive definition of the conditional value of the game

$$V_t(\theta_t) = \min_{w \in \mathcal{H}} \max_{\|g\| \leq 1} \langle g, w \rangle + V_{t+1}(\theta_t - g)$$

$$\left.\begin{aligned} V_T(\theta_T) &= f(\|\theta_T\|) \\ V &= V_0(\mathbf{0}) \end{aligned}\right\} \Rightarrow \forall u \in \mathcal{H}, \quad \text{Regret}(u) \leq f^*(\|u\|) + V_0(\mathbf{0})$$

Minimax-optimal strategy is

$$w_{t+1} = \arg\min_{w} \max_{\|g\| \leq 1} \langle g, w \rangle + V_{t+1}(\theta_t - g) .$$

Conditional Value of the Game

Recursive definition of the conditional value of the game

$$V_t(\theta_t) = \min_{w \in \mathcal{H}} \max_{\|g\| \leq 1} \langle g, w \rangle + V_{t+1}(\theta_t - g)$$

$$\left.\begin{aligned} V_T(\theta_T) &= f(\|\theta_T\|) \\ V &= V_0(\mathbf{0}) \end{aligned}\right\} \Rightarrow \forall u \in \mathcal{H}, \quad \text{Regret}(u) \leq f^*(\|u\|) + V_0(\mathbf{0})$$

Minimax-optimal strategy is

$$w_{t+1} = \arg\min_{w} \max_{\|g\| \leq 1} \langle g, w \rangle + V_{t+1}(\theta_t - g) \; .$$

How to calculate the conditional value of the game?

Orthogonal Adversary

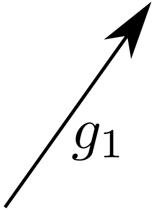

Orthogonal Adversary

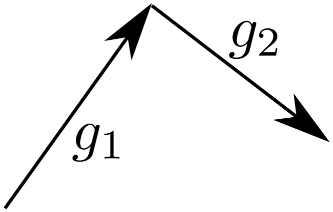

Orthogonal Adversary

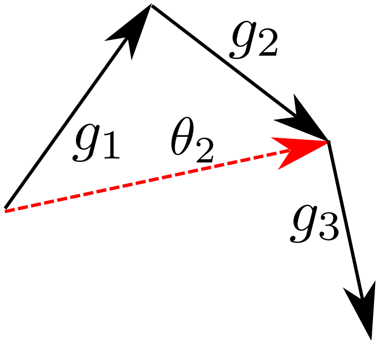

Orthogonal Adversary

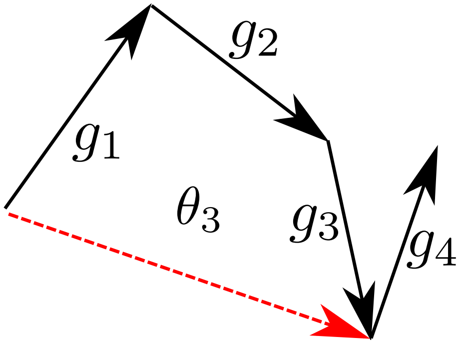

Orthogonal Adversary

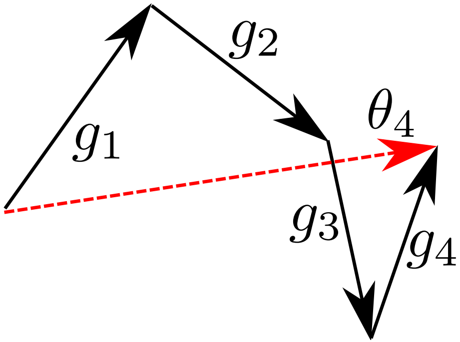

$$\|\theta_5\|^2 = \|\theta_4\|^2 + \|g_5\|^2 \text{ or } \|\theta_5\|^2 = \|\theta_4\|^2 - \|g_5\|^2$$

Is the Adversary Orthogonal?

Theorem

- $d > 1$, *f twice differentiable, and* $f''(x) \leq f'(x)/x, \forall x > 0$.

Is the Adversary Orthogonal?

Theorem

- $d > 1$, f *twice differentiable, and* $f''(x) \leq f'(x)/x, \forall x > 0$.

Then

- *Optimal* g_t *orthogonal to* θ_{t-1}

Is the Adversary Orthogonal?

Theorem

- $d > 1$, f *twice differentiable, and* $f''(x) \leq f'(x)/x, \forall x > 0$.

Then

- *Optimal* g_t *orthogonal to* θ_{t-1}
- *Conditional value of the game* $V_t(\theta_t) = f(\sqrt{\|\theta_t\|^2 + T - t})$

Is the Adversary Orthogonal?

Theorem

- $d > 1$, f *twice differentiable, and* $f''(x) \leq f'(x)/x, \forall x > 0$.

Then

- *Optimal* g_t *orthogonal to* θ_{t-1}
- *Conditional value of the game* $V_t(\theta_t) = f(\sqrt{\|\theta_t\|^2 + T - t})$
- *Value of the game is* $V = V_0(\mathbf{0}) = f(\sqrt{T})$

Hence

Is the Adversary Orthogonal?

Theorem

- $d > 1$, f *twice differentiable, and* $f''(x) \leq f'(x)/x, \forall x > 0$.

Then

- *Optimal* g_t *orthogonal to* θ_{t-1}
- *Conditional value of the game* $V_t(\theta_t) = f(\sqrt{\|\theta_t\|^2 + T - t})$
- *Value of the game is* $V = V_0(\mathbf{0}) = f(\sqrt{T})$

Hence

- *The minimax optimal strategy:* $w_{t+1} = \nabla V_t(\theta_t)$

Is the Adversary Orthogonal?

Theorem

- $d > 1$, f *twice differentiable, and* $f''(x) \leq f'(x)/x, \forall x > 0$.

Then

- *Optimal* g_t *orthogonal to* θ_{t-1}
- *Conditional value of the game* $V_t(\theta_t) = f(\sqrt{\|\theta_t\|^2 + T - t})$
- *Value of the game is* $V = V_0(\mathbf{0}) = f(\sqrt{T})$

Hence

- *The minimax optimal strategy:* $w_{t+1} = \nabla V_t(\theta_t)$
- *Regret:* $\text{Regret}(u) \leq f^*(\|u\|) + f(\sqrt{T})$

Is the Adversary Orthogonal?

Theorem

- *$d > 1$, f twice differentiable, and $f''(x) \le f'(x)/x, \forall x > 0$.*

Then

- *Optimal g_t orthogonal to θ_{t-1}*
- *Conditional value of the game $V_t(\theta_t) = f(\sqrt{\|\theta_t\|^2 + T - t})$*
- *Value of the game is $V = V_0(\mathbf{0}) = f(\sqrt{T})$*

Hence

- *The minimax optimal strategy: $w_{t+1} = \nabla V_t(\theta_t)$*
- *Regret:* $\text{Regret}(u) \le f^*(\|u\|) + f(\sqrt{T})$

Parallel Adversary

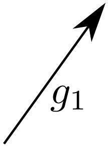

Parallel Adversary

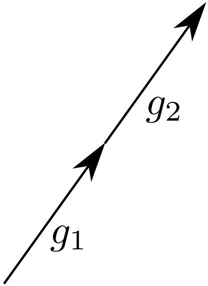

Parallel Adversary

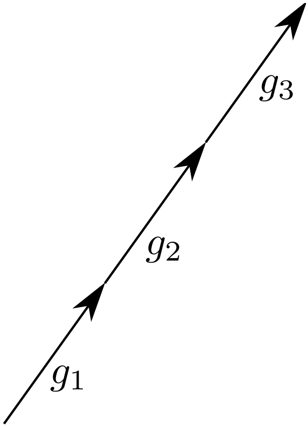

Parallel Adversary

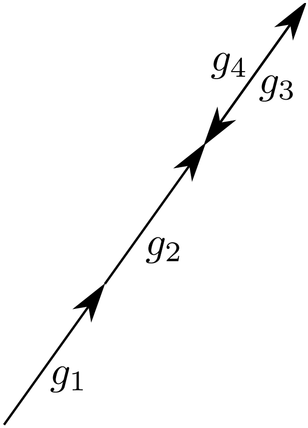

Parallel Adversary

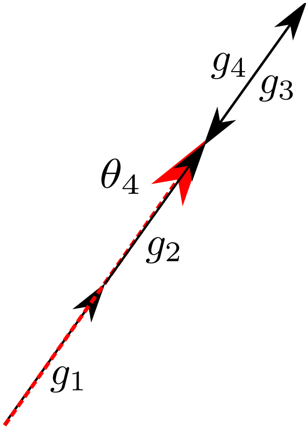

At any moment only two possible choices of g_t are possible

Is the Adversary Parallel? (Exact Case)

Theorem

- *Denote by* $r_t \sim \{-1, 1\}^t$ *the sum of* $T - t$ *IID Rademacher random variables and* $f_t(x) = \mathbb{E}_{r_t}\left[f\left(|x + r_t|\right)\right]$
- $d = 1$, *or* $d > 1$, f_t *twice differentiable,* $f_t''(x) \geq f_t'(x)/x, \forall x > 0$

Is the Adversary Parallel? (Exact Case)

Theorem

- *Denote by* $r_t \sim \{-1, 1\}^t$ *the sum of* $T - t$ *IID Rademacher random variables and* $f_t(x) = \mathbb{E}_{r_t}\left[f\left(|x + r_t|\right)\right]$
- $d = 1$, *or* $d > 1$, f_t *twice differentiable,* $f_t''(x) \geq f_t'(x)/x, \forall x > 0$

Then

- *Optimal* g_t *parallel to* θ_{t-1}

Is the Adversary Parallel? (Exact Case)

Theorem

- *Denote by* $r_t \sim \{-1, 1\}^t$ *the sum of* $T - t$ *IID Rademacher random variables and* $f_t(x) = \mathbb{E}_{r_t}\left[f\left(|x + r_t|\right)\right]$
- $d = 1$, *or* $d > 1$, f_t *twice differentiable,* $f_t''(x) \geq f_t'(x)/x, \forall x > 0$

Then

- *Optimal* g_t *parallel to* θ_{t-1}
- *Conditional value of the game* $V_t(\theta_t) = f_t(\|\theta_t\|)$

Is the Adversary Parallel? (Exact Case)

Theorem

- *Denote by* $r_t \sim \{-1, 1\}^t$ *the sum of* $T - t$ *IID Rademacher random variables and* $f_t(x) = \mathbb{E}_{r_t}[f(|x + r_t|)]$
- $d = 1$, *or* $d > 1$, f_t *twice differentiable,* $f_t''(x) \geq f_t'(x)/x, \forall x > 0$

Then

- *Optimal* g_t *parallel to* θ_{t-1}
- *Conditional value of the game* $V_t(\theta_t) = f_t(\|\theta_t\|)$
- *Value of the game* $V = V_0(\mathbf{0}) = \mathbb{E}_{r_0}[f(|r_0|)]$

Is the Adversary Parallel? (Exact Case)

Theorem

- *Denote by* $r_t \sim \{-1, 1\}^t$ *the sum of* $T - t$ *IID Rademacher random variables and* $f_t(x) = \mathbb{E}_{r_t}[f(|x + r_t|)]$
- $d = 1$, *or* $d > 1$, f_t *twice differentiable,* $f_t''(x) \geq f_t'(x)/x, \forall x > 0$

Then

- *Optimal* g_t *parallel to* θ_{t-1}
- *Conditional value of the game* $V_t(\theta_t) = f_t(\|\theta_t\|)$
- *Value of the game* $V = V_0(\mathbf{0}) = \mathbb{E}_{r_0}[f(|r_0|)]$

Hence

Is the Adversary Parallel? (Exact Case)

Theorem

- *Denote by* $r_t \sim \{-1,1\}^t$ *the sum of* $T - t$ *IID Rademacher random variables and* $f_t(x) = \mathbb{E}_{r_t}\left[f\left(|x + r_t|\right)\right]$
- $d = 1$, *or* $d > 1$, f_t *twice differentiable,* $f_t''(x) \geq f_t'(x)/x, \forall x > 0$

Then

- *Optimal* g_t *parallel to* θ_{t-1}
- *Conditional value of the game* $V_t(\theta_t) = f_t(\|\theta_t\|)$
- *Value of the game* $V = V_0(\mathbf{0}) = \mathbb{E}_{r_0}\left[f\left(|r_0|\right)\right]$

Hence

- *Minimax strategy:* $w_{t+1} = 0.5\ \theta_t \left(f_{t+1}\left(\|\theta_t\| + 1\right) - f_{t+1}\left(\|\theta_t\| - 1\right)\right)$

Is the Adversary Parallel? (Exact Case)

Theorem

- *Denote by* $r_t \sim \{-1, 1\}^t$ *the sum of* $T - t$ *IID Rademacher random variables and* $f_t(x) = \mathbb{E}_{r_t}[f(|x + r_t|)]$
- $d = 1$, *or* $d > 1$, f_t *twice differentiable,* $f_t''(x) \geq f_t'(x)/x, \forall x > 0$

Then

- *Optimal* g_t *parallel to* θ_{t-1}
- *Conditional value of the game* $V_t(\theta_t) = f_t(\|\theta_t\|)$
- *Value of the game* $V = V_0(\mathbf{0}) = \mathbb{E}_{r_0}[f(|r_0|)]$

Hence

- *Minimax strategy:* $w_{t+1} = 0.5\, \theta_t (f_{t+1}(\|\theta_t\| + 1) - f_{t+1}(\|\theta_t\| - 1))$
- *Regret:* $\text{Regret}(u) \leq f^*(\|u\|) + \mathbb{E}_{r_0}[f(|r_0|)]$

Is the Adversary Parallel? (Exact Case)

Theorem

- *Denote by* $r_t \sim \{-1,1\}^t$ *the sum of* $T-t$ *IID Rademacher random variables and* $f_t(x) = \mathbb{E}_{r_t}\left[f\left(|x+r_t|\right)\right]$
- $d=1$, *or* $d>1$, f_t *twice differentiable,* $f_t''(x) \geq f_t'(x)/x, \forall x>0$

Then

- *Optimal* g_t *parallel to* θ_{t-1}
- *Conditional value of the game* $V_t(\theta_t) = f_t(\|\theta_t\|)$
- *Value of the game* $V = V_0(\mathbf{0}) = \mathbb{E}_{r_0}\left[f\left(|r_0|\right)\right]$

Hence

- *Minimax strategy:* $w_{t+1} = 0.5\ \theta_t \left(f_{t+1}\left(\|\theta_t\|+1\right) - f_{t+1}\left(\|\theta_t\|-1\right)\right)$
- *Regret:* $\text{Regret}(u) \leq f^*(\|u\|) + \mathbb{E}_{r_0}\left[f\left(|r_0|\right)\right]$

Is the Adversary Parallel? (Exact Case)

Theorem

- *Denote by $r_t \sim \{-1, 1\}^t$ the sum of $T - t$ IID Rademacher random variables and $f_t(x) = \mathbb{E}_{r_t}[f(|x + r_t|)]$*
- *$d = 1$, or $d > 1$, f_t twice differentiable, $f_t''(x) \geq f_t'(x)/x, \forall x > 0$*

Then

- *Optimal g_t parallel to θ_{t-1}*
- *Conditional value of the game $V_t(\theta_t) = f_t(\|\theta_t\|)$*
- *Value of the game $V = V_0(\mathbf{0}) = \mathbb{E}_{r_0}[f(|r_0|)]$*

Hence

- *Minimax strategy: $w_{t+1} = 0.5\, \theta_t\, (f_{t+1}(\|\theta_t\| + 1) - f_{t+1}(\|\theta_t\| - 1))$*
- *Regret:* $\text{Regret}(u) \leq f^*(\|u\|) + \mathbb{E}_{r_0}[f(|r_0|)]$

Is the Adversary Parallel? (Normal Approximation)

Theorem

- *Denote by* $\phi_t \sim N(0, (T-t)\frac{\pi}{2})$ *the sum of* $T-t$ *and* $\hat{f}_t(x) = \mathbb{E}_{\phi_t}[f(|x+\phi_t|)]$
- $d=1$, *or* $d>1$, $\hat{f}_t$ *twice differentiable,* $\hat{f}''_t(x) \geq \hat{f}'_t(x)/x, \forall x>0$

Then

- *Optimal* g_t *parallel to* θ_{t-1}
- $f_t(\|\theta_t\|)$ *is an admissible relaxation to* $V_t(\theta_t)$
- *Value of the game* V *is upper bounded by* $\mathbb{E}_{\phi_0}[f(|\phi_0|)]$

Hence

- *Strategy:* $w_{t+1} = 0.5\ \theta_t \left(\hat{f}_{t+1}(\|\theta_t\|+1) - \hat{f}_{t+1}(\|\theta_t\|-1)\right)$
- *Regret:* $\text{Regret}(u) \leq f^*(\|u\|) + \mathbb{E}_{\phi_0}[f(|\phi_0|)]$

Is the Adversary Parallel? (Normal Approximation)

Theorem

- *Denote by* $\phi_t \sim N(0,(T-t)\frac{\pi}{2})$ *the sum of* $T-t$ *and* $\hat{f}_t(x) = \mathbb{E}_{\phi_t}\left[f\left(|x+\phi_t|\right)\right]$
- $d=1$, *or* $d>1$, $\hat{f}_t$ *twice differentiable,* $\hat{f}_t''(x) \geq \hat{f}_t'(x)/x, \forall x > 0$

Then

- *Optimal* g_t *parallel to* θ_{t-1}
- $f_t(\|\theta_t\|)$ *is an admissible relaxation to* $V_t(\theta_t)$
- *Value of the game* V *is upper bounded by* $\mathbb{E}_{\phi_0}\left[f\left(|\phi_0|\right)\right]$

Hence

- *Strategy:* $w_{t+1} = 0.5\ \theta_t \left(\hat{f}_{t+1}\left(\|\theta_t\|+1\right) - \hat{f}_{t+1}\left(\|\theta_t\|-1\right)\right)$
- *Regret:* $\text{Regret}(u) \leq f^*(\|u\|) + \mathbb{E}_{\phi_0}\left[f\left(|\phi_0|\right)\right]$

Is the Adversary Parallel? (Normal Approximation)

Theorem

- *Denote by* $\phi_t \sim N(0, (T-t)\frac{\pi}{2})$ *the sum of* $T - t$ *and* $\hat{f}_t(x) = \mathbb{E}_{\phi_t}[f(|x + \phi_t|)]$
- $d = 1$, *or* $d > 1$, $\hat{f}_t$ *twice differentiable,* $\hat{f}_t''(x) \geq \hat{f}_t'(x)/x, \forall x > 0$

Then

- *Optimal* g_t *parallel to* θ_{t-1}
- $f_t(\|\theta_t\|)$ *is an admissible relaxation to* $V_t(\theta_t)$
- *Value of the game* V *is upper bounded by* $\mathbb{E}_{\phi_0}[f(|\phi_0|)]$

Hence

- *Strategy:* $w_{t+1} = 0.5\ \theta_t \left(\hat{f}_{t+1}(\|\theta_t\| + 1) - \hat{f}_{t+1}(\|\theta_t\| - 1)\right)$
- *Regret:* $\text{Regret}(u) \leq f^*(\|u\|) + \mathbb{E}_{\phi_0}[f(|\phi_0|)]$

Is the Adversary Parallel? (Normal Approximation)

Theorem

- *Denote by* $\phi_t \sim N(0, (T-t)\frac{\pi}{2})$ *the sum of* $T-t$ *and* $\hat{f}_t(x) = \mathbb{E}_{\phi_t}[f(|x+\phi_t|)]$
- $d = 1$, *or* $d > 1$, $\hat{f}_t$ *twice differentiable,* $\hat{f}_t''(x) \geq \hat{f}_t'(x)/x, \forall x > 0$

Then

- *Optimal* g_t *parallel to* θ_{t-1}
- $f_t(\|\theta_t\|)$ *is an admissible relaxation to* $V_t(\theta_t)$
- *Value of the game* V *is upper bounded by* $\mathbb{E}_{\phi_0}[f(|\phi_0|)]$

Hence

- *Strategy:* $w_{t+1} = 0.5\, \theta_t \left(\hat{f}_{t+1}(\|\theta_t\|+1) - \hat{f}_{t+1}(\|\theta_t\|-1)\right)$
- *Regret:* $\text{Regret}(u) \leq f^*(\|u\|) + \mathbb{E}_{\phi_0}[f(|\phi_0|)]$

Outline

1. Online Learning
2. Minimax Analysis for Unconstrained Games
   - Conditional Value of Unconstrained Games
   - Orthogonal Adversary
   - Parallel Adversary
3. Applications
   - A Power Family of Minimax Algorithms
   - Tight Bounds for Unconstrained Learning

A Power Family of Minimax Algorithms

Corollary

- $f(x) = \frac{W}{p}|x|^p$ *for parameters* $W > 0$ *and* $p \in [1, 2]$, $d > 1$

A Power Family of Minimax Algorithms

Corollary

- $f(x) = \frac{W}{p}|x|^p$ *for parameters* $W > 0$ *and* $p \in [1, 2]$, $d > 1$

Then

- *We are in the case of the Orthogonal Adversary*

A Power Family of Minimax Algorithms

Corollary

- $f(x) = \frac{W}{p}|x|^p$ *for parameters* $W > 0$ *and* $p \in [1,2]$, $d > 1$

Then

- *We are in the case of the Orthogonal Adversary*
- *Conditional Value of the Game* $V_t(\theta_t) = \frac{W}{p}\left(\|\theta_t\|^2 + T - t\right)^{\frac{p}{2}}$

A Power Family of Minimax Algorithms

Corollary

- $f(x) = \frac{W}{p}|x|^p$ *for parameters* $W > 0$ *and* $p \in [1, 2]$, $d > 1$

Then

- *We are in the case of the Orthogonal Adversary*
- *Conditional Value of the Game* $V_t(\theta_t) = \frac{W}{p}\left(\|\theta_t\|^2 + T - t\right)^{\frac{p}{2}}$
-
$$\text{Regret}(u) \leq f^*(\|u\|) + V$$

A Power Family of Minimax Algorithms

Corollary

- *$f(x) = \frac{W}{p}|x|^p$ for parameters $W > 0$ and $p \in [1,2]$, $d > 1$*

Then

- *We are in the case of the Orthogonal Adversary*
- *Conditional Value of the Game $V_t(\theta_t) = \frac{W}{p}\left(\|\theta_t\|^2 + T - t\right)^{\frac{p}{2}}$*
- $$\text{Regret}(u) \leq \frac{1}{W^{q-1}q}\|u\|^q + \frac{W}{p}(\sqrt{T})^p$$
  - *For $p = 1$, we have*

$$\forall u : \|u\| \leq W, \quad \text{Regret}(u) \leq W\sqrt{T}$$

It extends (Abernethy et al., 2008) to $d = 2$

Tight Bounds for Unconstrained Learning

Corollary

- *Set* $f(x) = \epsilon \exp\left(\frac{x^2}{2aT}\right)$, $\phi_t \sim N(0, (T-t)\frac{\pi}{2})$

  $\hat{f}_t(x) = \mathbb{E}_{\phi_t}\left[f\left(|x + \phi_t|\right)\right]$

Tight Bounds for Unconstrained Learning

Corollary

- *Set* $f(x) = \epsilon \exp\left(\frac{x^2}{2aT}\right)$, $\phi_t \sim N(0, (T-t)\frac{\pi}{2})$

  $\hat{f}_t(x) = \mathbb{E}_{\phi_t}\left[f\left(|x+\phi_t|\right)\right] = \epsilon\left(1 - \frac{\pi(T-t)}{2aT}\right)^{-\frac{1}{2}} \exp\left(\frac{x^2}{2aT - \pi(T-t)}\right)$

Tight Bounds for Unconstrained Learning

Corollary

- *Set* $f(x) = \epsilon \exp\left(\frac{x^2}{2aT}\right)$, $\phi_t \sim N(0, (T-t)\frac{\pi}{2})$

  $\hat{f}_t(x) = \mathbb{E}_{\phi_t}\left[f\left(|x+\phi_t|\right)\right] = \epsilon\left(1 - \frac{\pi(T-t)}{2aT}\right)^{-\frac{1}{2}} \exp\left(\frac{x^2}{2aT - \pi(T-t)}\right)$

Then

- *We are in the case of the Parallel Adversary*

Tight Bounds for Unconstrained Learning

Corollary

- *Set* $f(x) = \epsilon \exp\left(\frac{x^2}{2aT}\right)$, $\phi_t \sim N(0, (T-t)\frac{\pi}{2})$

  $\hat{f}_t(x) = \mathbb{E}_{\phi_t}[f(|x+\phi_t|)] = \epsilon\left(1 - \frac{\pi(T-t)}{2aT}\right)^{-\frac{1}{2}} \exp\left(\frac{x^2}{2aT - \pi(T-t)}\right)$

Then

- *We are in the case of the Parallel Adversary*
- $w_{t+1} = \epsilon\theta_t \frac{\exp\left(\frac{(\|\theta_t\|+1)^2}{2aT-\pi(T-t-1)}\right) - \exp\left(\frac{(\|\theta_t\|-1)^2}{2aT-\pi(T-t-1)}\right)}{2\sqrt{1-\frac{\pi(T-t-1)}{2aT}}}$
- $\text{Regret}(u) \leq f^*(\|u\|) + \mathbb{E}_{\phi_0}[f(|\phi_0|)]$

Tight Bounds for Unconstrained Learning

Corollary

- *Set* $f(x) = \epsilon \exp\left(\frac{x^2}{2aT}\right)$, $\phi_t \sim N(0, (T-t)\frac{\pi}{2})$

  $\hat{f}_t(x) = \mathbb{E}_{\phi_t}\left[f\left(|x+\phi_t|\right)\right] = \epsilon\left(1 - \frac{\pi(T-t)}{2aT}\right)^{-\frac{1}{2}} \exp\left(\frac{x^2}{2aT - \pi(T-t)}\right)$

Then

- *We are in the case of the Parallel Adversary*
- $w_{t+1} = \epsilon\theta_t \frac{\exp\left(\frac{(\|\theta_t\|+1)^2}{2aT-\pi(T-t-1)}\right) - \exp\left(\frac{(\|\theta_t\|-1)^2}{2aT-\pi(T-t-1)}\right)}{2\sqrt{1-\frac{\pi(T-t-1)}{2aT}}}$
- $\text{Regret}(u) \le \|u\| \sqrt{2aT \log\left(\frac{\sqrt{aT}\|u\|}{\epsilon} + 1\right)} + \epsilon\left(1 - \frac{\pi}{2a}\right)^{-\frac{1}{2}}$
  - *Optimal up to constants, it improves over (Orabona, 2013)*

Conclusions

General conditions to calculate minimax optimal strategy for unconstrained games.

- Future Work:
  - Connection between optimal regret in unconstrained setting and optimal rates in statistical setting (check ArXiv in the next days...)
  - First expert algorithm that does not require communication
  - Generalization to arbitrary norms

Thanks for your attention

`http://francesco.orabona.com`