# Online Learning
# Lecture 5

Nicolò Cesa-Bianchi

Università degli Studi di Milano

# Exploiting curvature of the losses

# Exploiting curvature of the losses

- Convex and $G$-Lipschitz losses: FTRL with $\psi = \frac{1}{2}\left\|\cdot\right\|_2^2$ achieves $R_T = \mathcal{O}(GD\sqrt{T})$

# Exploiting curvature of the losses

- Convex and $G$-Lipschitz losses: FTRL with $\psi = \frac{1}{2}\left\|\cdot\right\|_2^2$ achieves $R_T = \mathcal{O}(GD\sqrt{T})$
- Strongly convex and $G$-Lipschitz losses: FTL achieves $R_T = \mathcal{O}(G^2 \ln T)$

# Exploiting curvature of the losses

- Convex and $G$-Lipschitz losses: FTRL with $\psi = \frac{1}{2}\|\cdot\|_2^2$ achieves $R_T = \mathcal{O}(GD\sqrt{T})$
- Strongly convex and $G$-Lipschitz losses: FTL achieves $R_T = \mathcal{O}(G^2 \ln T)$

## Strong convexity in the direction of the gradient (exp-concavity)

$$\ell_t(\boldsymbol{u}) \geq \ell_t(\boldsymbol{w}) + \boldsymbol{g}_t^\top(\boldsymbol{u}-\boldsymbol{w}) + \frac{\lambda}{2}\|\boldsymbol{u}-\boldsymbol{w}\|^2_{\boldsymbol{g}_t\boldsymbol{g}_t^\top} \qquad \boldsymbol{u},\boldsymbol{w} \in \mathbb{V}$$

where $\boldsymbol{g}_t = \nabla \ell_t(\boldsymbol{w})$ and $\|\boldsymbol{w}\|_M^2 = \boldsymbol{w}^\top M \boldsymbol{w}$

# Exploiting curvature of the losses

- Convex and $G$-Lipschitz losses: FTRL with $\psi = \frac{1}{2}\|\cdot\|_2^2$ achieves $R_T = \mathcal{O}(GD\sqrt{T})$
- Strongly convex and $G$-Lipschitz losses: FTL achieves $R_T = \mathcal{O}(G^2 \ln T)$

## Strong convexity in the direction of the gradient (exp-concavity)

$$\ell_t(\boldsymbol{u}) \geq \ell_t(\boldsymbol{w}) + \boldsymbol{g}_t^\top(\boldsymbol{u}-\boldsymbol{w}) + \frac{\lambda}{2}\|\boldsymbol{u}-\boldsymbol{w}\|^2_{\boldsymbol{g}_t\boldsymbol{g}_t^\top} \qquad \boldsymbol{u},\boldsymbol{w} \in \mathbb{V}$$

where $\boldsymbol{g}_t = \nabla\ell_t(\boldsymbol{w})$ and $\|\boldsymbol{w}\|_M^2 = \boldsymbol{w}^\top M \boldsymbol{w}$

## Some losses satisfying the condition (in a bounded domain)

- Square loss $\ell(\boldsymbol{w}) = \frac{1}{2}(\boldsymbol{w}^\top\boldsymbol{x} - y)^2$

# Exploiting curvature of the losses

- Convex and $G$-Lipschitz losses: FTRL with $\psi = \frac{1}{2}\|\cdot\|_2^2$ achieves $R_T = \mathcal{O}(GD\sqrt{T})$
- Strongly convex and $G$-Lipschitz losses: FTL achieves $R_T = \mathcal{O}(G^2 \ln T)$

## Strong convexity in the direction of the gradient (exp-concavity)

$$\ell_t(\boldsymbol{u}) \geq \ell_t(\boldsymbol{w}) + \boldsymbol{g}_t^\top(\boldsymbol{u}-\boldsymbol{w}) + \frac{\lambda}{2}\|\boldsymbol{u}-\boldsymbol{w}\|^2_{\boldsymbol{g}_t\boldsymbol{g}_t^\top} \qquad \boldsymbol{u},\boldsymbol{w} \in \mathbb{V}$$

where $\boldsymbol{g}_t = \nabla \ell_t(\boldsymbol{w})$ and $\|\boldsymbol{w}\|_M^2 = \boldsymbol{w}^\top M \boldsymbol{w}$

## Some losses satisfying the condition (in a bounded domain)

- Square loss $\ell(\boldsymbol{w}) = \frac{1}{2}(\boldsymbol{w}^\top\boldsymbol{x} - y)^2$
- Logistic loss $\ell(\boldsymbol{w}) = \ln\left(1 + \exp(-y\boldsymbol{w}^\top\boldsymbol{x})\right)$

# Online Newton Step for exp-concave losses

Choose the model minimizing a second-order approximation of the true loss:

$$\boldsymbol{w}_{t+1} = \operatorname*{argmin}_{\boldsymbol{w} \in \mathbb{V}} \sum_{s=1}^{t} \widehat{\ell}_s(\boldsymbol{w}) \qquad \text{(FTL on a surrogate loss)}$$

$$\widehat{\ell}_t(\boldsymbol{w}) = \ell_t(\boldsymbol{w}_t) + \boldsymbol{g}_t^\top (\boldsymbol{w} - \boldsymbol{w}_t) + \frac{\lambda}{2} \|\boldsymbol{w} - \boldsymbol{w}_t\|^2_{\boldsymbol{g}_t \boldsymbol{g}_t^\top} \qquad \boldsymbol{g}_t = \nabla \ell_t(\boldsymbol{w}_t)$$

# Online Newton Step for exp-concave losses

Choose the model minimizing a second-order approximation of the true loss:

$$\boldsymbol{w}_{t+1} = \operatorname*{argmin}_{\boldsymbol{w}\in\mathbb{V}} \sum_{s=1}^{t} \widehat{\ell}_s(\boldsymbol{w}) \qquad \text{(FTL on a surrogate loss)}$$

$$\widehat{\ell}_t(\boldsymbol{w}) = \ell_t(\boldsymbol{w}_t) + \boldsymbol{g}_t^\top(\boldsymbol{w}-\boldsymbol{w}_t) + \frac{\lambda}{2}\|\boldsymbol{w}-\boldsymbol{w}_t\|^2_{\boldsymbol{g}_t\boldsymbol{g}_t^\top} \qquad \boldsymbol{g}_t = \nabla\ell_t(\boldsymbol{w}_t)$$

Properties:

# Online Newton Step for exp-concave losses

Choose the model minimizing a second-order approximation of the true loss:

$$\boldsymbol{w}_{t+1} = \operatorname*{argmin}_{\boldsymbol{w}\in\mathbb{V}} \sum_{s=1}^{t} \widehat{\ell}_s(\boldsymbol{w}) \qquad \text{(FTL on a surrogate loss)}$$

$$\widehat{\ell}_t(\boldsymbol{w}) = \ell_t(\boldsymbol{w}_t) + \boldsymbol{g}_t^\top(\boldsymbol{w}-\boldsymbol{w}_t) + \frac{\lambda}{2}\|\boldsymbol{w}-\boldsymbol{w}_t\|^2_{\boldsymbol{g}_t\boldsymbol{g}_t^\top} \qquad \boldsymbol{g}_t = \nabla\ell_t(\boldsymbol{w}_t)$$

Properties:

- $\widehat{\ell}_t(\boldsymbol{u}) \le \ell_t(\boldsymbol{u})$ for all $\boldsymbol{u}\in\mathbb{V}$

# Online Newton Step for exp-concave losses

Choose the model minimizing a second-order approximation of the true loss:

$$\boldsymbol{w}_{t+1} = \underset{\boldsymbol{w}\in\mathbb{V}}{\operatorname{argmin}} \sum_{s=1}^{t} \widehat{\ell}_s(\boldsymbol{w}) \qquad \text{(FTL on a surrogate loss)}$$

$$\widehat{\ell}_t(\boldsymbol{w}) = \ell_t(\boldsymbol{w}_t) + \boldsymbol{g}_t^\top(\boldsymbol{w}-\boldsymbol{w}_t) + \frac{\lambda}{2}\left\|\boldsymbol{w}-\boldsymbol{w}_t\right\|^2_{\boldsymbol{g}_t\boldsymbol{g}_t^\top} \qquad \boldsymbol{g}_t = \nabla\ell_t(\boldsymbol{w}_t)$$

## Properties:

- $\widehat{\ell}_t(\boldsymbol{u}) \le \ell_t(\boldsymbol{u})$ for all $\boldsymbol{u} \in \mathbb{V}$
- $\widehat{\ell}_t(\boldsymbol{w}_t) = \ell_t(\boldsymbol{w}_t)$

# Online Newton Step for exp-concave losses

Choose the model minimizing a second-order approximation of the true loss:

$$\boldsymbol{w}_{t+1} = \operatorname*{argmin}_{\boldsymbol{w}\in\mathbb{V}} \sum_{s=1}^{t} \widehat{\ell}_s(\boldsymbol{w}) \qquad \text{(FTL on a surrogate loss)}$$

$$\widehat{\ell}_t(\boldsymbol{w}) = \ell_t(\boldsymbol{w}_t) + \boldsymbol{g}_t^\top(\boldsymbol{w} - \boldsymbol{w}_t) + \frac{\lambda}{2}\|\boldsymbol{w} - \boldsymbol{w}_t\|^2_{\boldsymbol{g}_t\boldsymbol{g}_t^\top} \qquad \boldsymbol{g}_t = \nabla\ell_t(\boldsymbol{w}_t)$$

Properties:

- $\widehat{\ell}_t(\boldsymbol{u}) \le \ell_t(\boldsymbol{u})$ for all $\boldsymbol{u} \in \mathbb{V}$
- $\widehat{\ell}_t(\boldsymbol{w}_t) = \ell_t(\boldsymbol{w}_t)$
- Regret bound: $R_T(\boldsymbol{u}) \le \sum_{t=1}^{T} \widehat{\ell}_t(\boldsymbol{w}_t) - \sum_{t=1}^{T} \widehat{\ell}_t(\boldsymbol{u}) = \mathcal{O}(GDd\ln T)$

# Online Newton Step for exp-concave losses

Choose the model minimizing a second-order approximation of the true loss:

$$\boldsymbol{w}_{t+1} = \operatorname*{argmin}_{\boldsymbol{w}\in\mathbb{V}} \sum_{s=1}^{t} \widehat{\ell}_s(\boldsymbol{w}) \qquad \text{(FTL on a surrogate loss)}$$

$$\widehat{\ell}_t(\boldsymbol{w}) = \ell_t(\boldsymbol{w}_t) + \boldsymbol{g}_t^\top(\boldsymbol{w}-\boldsymbol{w}_t) + \frac{\lambda}{2}\|\boldsymbol{w}-\boldsymbol{w}_t\|^2_{\boldsymbol{g}_t\boldsymbol{g}_t^\top} \qquad \boldsymbol{g}_t = \nabla\ell_t(\boldsymbol{w}_t)$$

Properties:

- $\widehat{\ell}_t(\boldsymbol{u}) \le \ell_t(\boldsymbol{u})$ for all $\boldsymbol{u} \in \mathbb{V}$
- $\widehat{\ell}_t(\boldsymbol{w}_t) = \ell_t(\boldsymbol{w}_t)$
- Regret bound: $R_T(\boldsymbol{u}) \le \sum_{t=1}^{T} \widehat{\ell}_t(\boldsymbol{w}_t) - \sum_{t=1}^{T} \widehat{\ell}_t(\boldsymbol{u}) = \mathcal{O}(GDd\ln T)$
- This matches the $\mathcal{O}(\ln T)$ bound for strongly convex losses

# Unconstrained online convex optimization

- Assume $\ell_t$ is $1$-Lipschitz for $t \geq 1$

# Unconstrained online convex optimization

- Assume $\ell_t$ is $1$-Lipschitz for $t \geq 1$
- Run FTRL with Euclidean regularizer $\psi = \frac{1}{2}\left\|\cdot\right\|_2^2$, no projection, and learning rate $\eta = \alpha/\sqrt{T}$ for $\alpha > 0$

# Unconstrained online convex optimization

- Assume $\ell_t$ is $1$-Lipschitz for $t \geq 1$
- Run FTRL with Euclidean regularizer $\psi = \frac{1}{2}\left\|\cdot\right\|_2^2$, no projection, and learning rate $\eta = \alpha/\sqrt{T}$ for $\alpha > 0$
- $R_T(\boldsymbol{u}) \leq \dfrac{\psi(\boldsymbol{u}) - \psi(\boldsymbol{w}_1)}{\eta} + \eta T = \dfrac{1}{2}\left(\dfrac{\|\boldsymbol{u}\|_2^2}{\alpha} + \alpha\right)\sqrt{T} \qquad \forall \boldsymbol{u} \in \mathbb{R}^d$

# Unconstrained online convex optimization

- Assume $\ell_t$ is $1$-Lipschitz for $t \geq 1$
- Run FTRL with Euclidean regularizer $\psi = \frac{1}{2}\left\|\cdot\right\|_2^2$, no projection, and learning rate $\eta = \alpha/\sqrt{T}$ for $\alpha > 0$
- $R_T(\boldsymbol{u}) \leq \dfrac{\psi(\boldsymbol{u}) - \psi(\boldsymbol{w}_1)}{\eta} + \eta T = \dfrac{1}{2}\left(\dfrac{\|\boldsymbol{u}\|_2^2}{\alpha} + \alpha\right)\sqrt{T} \qquad \forall \boldsymbol{u} \in \mathbb{R}^d$
- $R_T(\boldsymbol{u}) \leq \|\boldsymbol{u}\|_2 \sqrt{T}$ for $\alpha = \|\boldsymbol{u}\|_2$

# Unconstrained online convex optimization

- Assume $\ell_t$ is $1$-Lipschitz for $t \geq 1$
- Run FTRL with Euclidean regularizer $\psi = \frac{1}{2}\left\|\cdot\right\|_2^2$, no projection, and learning rate $\eta = \alpha/\sqrt{T}$ for $\alpha > 0$
- $R_T(\boldsymbol{u}) \leq \dfrac{\psi(\boldsymbol{u}) - \psi(\boldsymbol{w}_1)}{\eta} + \eta T = \dfrac{1}{2}\left(\dfrac{\|\boldsymbol{u}\|_2^2}{\alpha} + \alpha\right)\sqrt{T} \qquad \forall \boldsymbol{u} \in \mathbb{R}^d$
- $R_T(\boldsymbol{u}) \leq \|\boldsymbol{u}\|_2 \sqrt{T}$ for $\alpha = \|\boldsymbol{u}\|_2$
- This bound cannot be simultaneously achieved for all $\boldsymbol{u}$!

# Main idea

- Control $R_T(\boldsymbol{u})$ by learning length $w = \|\boldsymbol{u}\|_2$ and direction $\boldsymbol{v} = \boldsymbol{u} / \|\boldsymbol{u}\|_2$ separately

# Main idea

- Control $R_T(\boldsymbol{u})$ by learning length $w = \|\boldsymbol{u}\|_2$ and direction $\boldsymbol{v} = \boldsymbol{u}/\|\boldsymbol{u}\|_2$ separately
- The direction can be learned using FTRL with projection onto the unit Euclidean ball

# Main idea

- Control $R_T(\boldsymbol{u})$ by learning length $w = \|\boldsymbol{u}\|_2$ and direction $\boldsymbol{v} = \boldsymbol{u}/\|\boldsymbol{u}\|_2$ separately
- The direction can be learned using FTRL with projection onto the unit Euclidean ball
- The length is learned using a parameterless $1$-dimensional online learning algorithm

# Main idea

- Control $R_T(\boldsymbol{u})$ by learning length $w = \|\boldsymbol{u}\|_2$ and direction $\boldsymbol{v} = \boldsymbol{u} / \|\boldsymbol{u}\|_2$ separately
- The direction can be learned using FTRL with projection onto the unit Euclidean ball
- The length is learned using a parameterless $1$-dimensional online learning algorithm
- We predict with $w\boldsymbol{v}$

# Analysis

# Analysis

$$R_T(\boldsymbol{u}) = \sum_{t=1}^{T} \ell_t(w_t \boldsymbol{v}_t) - \sum_{t=1}^{T} \ell_t(\boldsymbol{u})$$

# Analysis

$$R_T(\boldsymbol{u}) = \sum_{t=1}^{T} \ell_t(w_t \boldsymbol{v}_t) - \sum_{t=1}^{T} \ell_t(\boldsymbol{u})$$

$$\leq \sum_{t=1}^{T} \boldsymbol{g}_t^\top (w_t \boldsymbol{v}_t - \boldsymbol{u}) \qquad \text{(linearized regret)}$$

# Analysis

$$
\begin{aligned}
R_T(\boldsymbol{u}) &= \sum_{t=1}^{T} \ell_t(w_t \boldsymbol{v}_t) - \sum_{t=1}^{T} \ell_t(\boldsymbol{u}) \\
&\le \sum_{t=1}^{T} \boldsymbol{g}_t^\top (w_t \boldsymbol{v}_t - \boldsymbol{u}) \qquad \text{(linearized regret)} \\
&= \sum_{t=1}^{T} \Big( w_t \, \boldsymbol{g}_t^\top \boldsymbol{v}_t - \|\boldsymbol{u}\|_2 \, \boldsymbol{g}_t^\top \boldsymbol{v}_t \Big) + \|\boldsymbol{u}\|_2 \sum_{t=1}^{T} \left( \boldsymbol{g}_t^\top \boldsymbol{v}_t - \boldsymbol{g}_t^\top \frac{\boldsymbol{u}}{\|\boldsymbol{u}\|_2} \right)
\end{aligned}
$$

# Analysis

$$
\begin{aligned}
R_T(\boldsymbol{u}) &= \sum_{t=1}^{T} \ell_t(w_t \boldsymbol{v}_t) - \sum_{t=1}^{T} \ell_t(\boldsymbol{u}) \\
&\le \sum_{t=1}^{T} \boldsymbol{g}_t^\top (w_t \boldsymbol{v}_t - \boldsymbol{u}) \qquad \text{(linearized regret)} \\
&= \sum_{t=1}^{T} \underbrace{\left(w_t\, \ell_t'(w_t) - \|\boldsymbol{u}\|_2\, \ell_t'(w_t)\right)}_{\text{parameterless}} + \|\boldsymbol{u}\|_2 \underbrace{\sum_{t=1}^{T} \left( \boldsymbol{g}_t^\top \boldsymbol{v}_t - \boldsymbol{g}_t^\top \frac{\boldsymbol{u}}{\|\boldsymbol{u}\|_2} \right)}_{\text{FTRL}}
\end{aligned}
$$

# Learning and investing

$1$-dimensional parameterless online algorithms extracted from investment strategies

# Learning and investing

1-dimensional parameterless online algorithms extracted from investment strategies

## The betting game

# Learning and investing

$1$-dimensional parameterless online algorithms extracted from investment strategies

## The betting game

- The bettor starts out with an initial wealth of $C_0 = 1$

# Learning and investing

$1$-dimensional parameterless online algorithms extracted from investment strategies

## The betting game

- The bettor starts out with an initial wealth of $C_0 = 1$
- In each round $t = 1, 2, \ldots$ of the game

# Learning and investing

$1$-dimensional parameterless online algorithms extracted from investment strategies

## The betting game

- The bettor starts out with an initial wealth of $C_0 = 1$
- In each round $t = 1, 2, \ldots$ of the game
    1. The bettor bets $\alpha_t \in [-1, 1]$

# Learning and investing

$1$-dimensional parameterless online algorithms extracted from investment strategies

## The betting game

- The bettor starts out with an initial wealth of $C_0 = 1$
- In each round $t = 1, 2, \dots$ of the game
  1. The bettor bets $\alpha_t \in [-1, 1]$
  2. The market reveals $x_t \in [-1, 1]$

# Learning and investing

$1$-dimensional parameterless online algorithms extracted from investment strategies

## The betting game

- The bettor starts out with an initial wealth of $C_0 = 1$
- In each round $t = 1, 2, \dots$ of the game
  1. The bettor bets $\alpha_t \in [-1, 1]$
  2. The market reveals $x_t \in [-1, 1]$
  3. The bettor's wealth is $C_t = (1 + \alpha_t x_t) C_{t-1}$

# Learning and investing

$1$-dimensional parameterless online algorithms extracted from investment strategies

## The betting game

- The bettor starts out with an initial wealth of $C_0 = 1$
- In each round $t = 1, 2, \dots$ of the game
  1. The bettor bets $\alpha_t \in [-1, 1]$
  2. The market reveals $x_t \in [-1, 1]$
  3. The bettor's wealth is $C_t = (1 + \alpha_t x_t)C_{t-1}$

## A reduction from prediction to investment

# Learning and investing

$1$-dimensional parameterless online algorithms extracted from investment strategies

## The betting game

- The bettor starts out with an initial wealth of $C_0 = 1$
- In each round $t = 1, 2, \ldots$ of the game
  1. The bettor bets $\alpha_t \in [-1, 1]$
  2. The market reveals $x_t \in [-1, 1]$
  3. The bettor's wealth is $C_t = (1 + \alpha_t x_t) C_{t-1}$

## A reduction from prediction to investment

- Predict using $w_t = \alpha_t \, C_{t-1}$ implying $C_t = C_{t-1} + w_t x_t$

# Learning and investing

$1$-dimensional parameterless online algorithms extracted from investment strategies

## The betting game

- The bettor starts out with an initial wealth of $C_0 = 1$
- In each round $t = 1, 2, \ldots$ of the game
  1. The bettor bets $\alpha_t \in [-1, 1]$
  2. The market reveals $x_t \in [-1, 1]$
  3. The bettor's wealth is $C_t = (1 + \alpha_t x_t) C_{t-1}$

## A reduction from prediction to investment

- Predict using $w_t = \alpha_t \, C_{t-1}$ implying $C_t = C_{t-1} + w_t x_t$
- Provide feedback $x_t = -\ell_t'(w_t) = -\boldsymbol{g}_t^\top \boldsymbol{v}_t$

# Learning and investing

$1$-dimensional parameterless online algorithms extracted from investment strategies

## The betting game

- The bettor starts out with an initial wealth of $C_0 = 1$
- In each round $t = 1, 2, \ldots$ of the game
  1. The bettor bets $\alpha_t \in [-1, 1]$
  2. The market reveals $x_t \in [-1, 1]$
  3. The bettor's wealth is $C_t = (1 + \alpha_t x_t) C_{t-1}$

## A reduction from prediction to investment

- Predict using $w_t = \alpha_t \, C_{t-1}$ implying $C_t = C_{t-1} + w_t x_t$
- Provide feedback $x_t = -\ell_t'(w_t) = -\boldsymbol{g}_t^\top \boldsymbol{v}_t$
- $C_T = \prod_{t=1}^T (1 + \alpha_t x_t) = 1 + \sum_{t=1}^T w_t x_t = 1 - \sum_{t=1}^T w_t \, \ell_t'(w_t)$

# Connecting wealth and regret

For a convex $\phi$ assume a betting strategy achieves $C_T \geq \phi\left(\sum_{t=1}^{T} x_t\right) = \phi\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right)$

# Connecting wealth and regret

For a convex $\phi$ assume a betting strategy achieves $C_T \ge \phi\left(\sum_{t=1}^{T} x_t\right) = \phi\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right)$

$$R_T(u) \le \sum_{t=1}^{T} (w_t - u)\ell_t'(w_t) \qquad \text{(any } u \in \mathbb{R}\text{)}$$

# Connecting wealth and regret

For a convex $\phi$ assume a betting strategy achieves $C_T \geq \phi\left(\sum_{t=1}^{T} x_t\right) = \phi\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right)$

$$R_T(u) \leq \sum_{t=1}^{T} (w_t - u)\ell_t'(w_t) \qquad \text{(any } u \in \mathbb{R}\text{)}$$

$$= -u \sum_{t=1}^{T} \ell_t'(w_t) - \left(1 - \sum_{t=1}^{T} w_t \ell_t'(w_t)\right) + 1$$

# Connecting wealth and regret

For a convex $\phi$ assume a betting strategy achieves $C_T \geq \phi\left(\sum_{t=1}^{T} x_t\right) = \phi\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right)$

$$
\begin{aligned}
R_T(u) &\leq \sum_{t=1}^{T} (w_t - u)\ell_t'(w_t) && \text{(any } u \in \mathbb{R}\text{)} \\
&= -u \sum_{t=1}^{T} \ell_t'(w_t) - \left(1 - \sum_{t=1}^{T} w_t \ell_t'(w_t)\right) + 1 \\
&= -u \sum_{t=1}^{T} \ell_t'(w_t) - C_T + 1 && \left(\text{using } C_T = 1 - \textstyle\sum_{t=1}^{T} w_t\, \ell_t'(w_t)\right)
\end{aligned}
$$

# Connecting wealth and regret

For a convex $\phi$ assume a betting strategy achieves $C_T \geq \phi\left(\sum_{t=1}^{T} x_t\right) = \phi\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right)$

$$
\begin{aligned}
R_T(u) &\leq \sum_{t=1}^{T} (w_t - u)\ell_t'(w_t) && \text{(any } u \in \mathbb{R}\text{)} \\
&= -u \sum_{t=1}^{T} \ell_t'(w_t) - \left(1 - \sum_{t=1}^{T} w_t \ell_t'(w_t)\right) + 1 \\
&= -u \sum_{t=1}^{T} \ell_t'(w_t) - C_T + 1 && \left(\text{using } C_T = 1 - \textstyle\sum_{t=1}^{T} w_t\, \ell_t'(w_t)\right) \\
&\leq -u \sum_{t=1}^{T} \ell_t'(w_t) - \phi\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right) + 1
\end{aligned}
$$

## Connecting wealth and regret

For a convex $\phi$ assume a betting strategy achieves $C_T \geq \phi\left(\sum_{t=1}^{T} x_t\right) = \phi\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right)$

$$
\begin{aligned}
R_T(u) &\leq \sum_{t=1}^{T} (w_t - u)\ell_t'(w_t) && \text{(any } u \in \mathbb{R}\text{)} \\
&= -u \sum_{t=1}^{T} \ell_t'(w_t) - \left(1 - \sum_{t=1}^{T} w_t \ell_t'(w_t)\right) + 1 \\
&= -u \sum_{t=1}^{T} \ell_t'(w_t) - C_T + 1 && \text{(using } C_T = 1 - \textstyle\sum_{t=1}^{T} w_t\, \ell_t'(w_t)\text{)} \\
&\leq -u \sum_{t=1}^{T} \ell_t'(w_t) - \phi\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right) + 1 \\
&\leq \sup_{\theta \in \mathbb{R}} u\theta - \phi(\theta) + 1 && (\theta = -\ell_1'(w_1) - \cdots - \ell_T(w_T))
\end{aligned}
$$

# Connecting wealth and regret

For a convex $\phi$ assume a betting strategy achieves $C_T \geq \phi\left(\sum_{t=1}^{T} x_t\right) = \phi\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right)$

$$
\begin{aligned}
R_T(u) &\leq \sum_{t=1}^{T} (w_t - u)\ell_t'(w_t) && \text{(any } u \in \mathbb{R}\text{)} \\
&= -u \sum_{t=1}^{T} \ell_t'(w_t) - \left(1 - \sum_{t=1}^{T} w_t \ell_t'(w_t)\right) + 1 \\
&= -u \sum_{t=1}^{T} \ell_t'(w_t) - C_T + 1 && \text{(using } C_T = 1 - \textstyle\sum_{t=1}^{T} w_t\, \ell_t'(w_t)\text{)} \\
&\leq -u \sum_{t=1}^{T} \ell_t'(w_t) - \phi\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right) + 1 \\
&\leq \sup_{\theta \in \mathbb{R}} u\theta - \phi(\theta) + 1 && (\theta = -\ell_1'(w_1) - \cdots - \ell_T(w_T)) \\
&= \phi^*(\theta) + 1
\end{aligned}
$$

# Regret bound

- Betting strategy: $\alpha_1 = 0$ and $\alpha_t = (x_1 + \cdots + x_{t-1})/t$ for $t \geq 1$ (Krichevsky-Trofimov estimator)

# Regret bound

- Betting strategy: $\alpha_1 = 0$ and $\alpha_t = (x_1 + \cdots + x_{t-1})/t$ for $t \geq 1$ (Krichevsky-Trofimov estimator)
- Achieved wealth: $C_T \geq \frac{1}{\sqrt{T}} \exp\left(\frac{1}{2T}\left(\sum_{t=1}^{T} x_t\right)^2\right) = \phi\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right)$

# Regret bound

- Betting strategy: $\alpha_1 = 0$ and $\alpha_t = (x_1 + \cdots + x_{t-1})/t$ for $t \geq 1$ (Krichevsky-Trofimov estimator)
- Achieved wealth: $C_T \geq \frac{1}{\sqrt{T}} \exp\left(\frac{1}{2T}\left(\sum_{t=1}^{T} x_t\right)^2\right) = \phi\left(-\sum_{t=1}^{T} \ell'_t(w_t)\right)$
- Resulting regret: $R_T(u) = \phi^*\left(-\sum_{t=1}^{T} \ell'_t(w_t)\right) = \mathcal{O}\left(|u|\sqrt{T\ln(u^2T+1)} + 1\right)$ for any $u \in \mathbb{R}$

# Regret bound

- Betting strategy: $\alpha_1 = 0$ and $\alpha_t = (x_1 + \cdots + x_{t-1})/t$ for $t \geq 1$ (Krichevsky-Trofimov estimator)
- Achieved wealth: $C_T \geq \frac{1}{\sqrt{T}} \exp\left(\frac{1}{2T}\left(\sum_{t=1}^{T} x_t\right)^2\right) = \phi\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right)$
- Resulting regret: $R_T(u) = \phi^*\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right) = \mathcal{O}\left(|u|\sqrt{T\ln(u^2T+1)}+1\right)$ for any $u \in \mathbb{R}$

$R_T(\boldsymbol{u}) \leq \sum_{t=1}^{T}\Big(w_t\,\ell_t'(w_t) - \|\boldsymbol{u}\|_2\,\ell_t'(w_t)\Big) + \|\boldsymbol{u}\|_2 \sum_{t=1}^{T}\left(\boldsymbol{g}_t^\top \boldsymbol{v}_t - \boldsymbol{g}_t^\top \frac{\boldsymbol{u}}{\|\boldsymbol{u}\|_2}\right)$ (for any $\boldsymbol{u} \in \mathbb{R}^d$)

# Regret bound

- Betting strategy: $\alpha_1 = 0$ and $\alpha_t = (x_1 + \cdots + x_{t-1})/t$ for $t \geq 1$ (Krichevsky-Trofimov estimator)
- Achieved wealth: $C_T \geq \frac{1}{\sqrt{T}} \exp\left(\frac{1}{2T}\left(\sum_{t=1}^{T} x_t\right)^2\right) = \phi\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right)$
- Resulting regret: $R_T(u) = \phi^*\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right) = \mathcal{O}\left(|u|\sqrt{T \ln(u^2 T + 1)} + 1\right)$ for any $u \in \mathbb{R}$

$$R_T(\boldsymbol{u}) \leq \sum_{t=1}^{T}\Big(w_t\,\ell_t'(w_t) - \|\boldsymbol{u}\|_2\,\ell_t'(w_t)\Big) + \|\boldsymbol{u}\|_2 \sum_{t=1}^{T}\left(\boldsymbol{g}_t^\top \boldsymbol{v}_t - \boldsymbol{g}_t^\top \frac{\boldsymbol{u}}{\|\boldsymbol{u}\|_2}\right) \quad \text{(for any } \boldsymbol{u} \in \mathbb{R}^d\text{)}$$

$$= \mathcal{O}\Big(\|\boldsymbol{u}\|_2 \sqrt{T \ln\left(\|\boldsymbol{u}\|_2^2\, T + 1\right)} + 1\Big) + \mathcal{O}\big(\|\boldsymbol{u}\|_2 \sqrt{T}\big)$$

# Regret bound

- Betting strategy: $\alpha_1 = 0$ and $\alpha_t = (x_1 + \cdots + x_{t-1})/t$ for $t \geq 1$ (Krichevsky-Trofimov estimator)
- Achieved wealth: $C_T \geq \frac{1}{\sqrt{T}} \exp\left(\frac{1}{2T}\left(\sum_{t=1}^{T} x_t\right)^2\right) = \phi\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right)$
- Resulting regret: $R_T(u) = \phi^*\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right) = \mathcal{O}\left(|u|\sqrt{T\ln(u^2T+1)}+1\right)$ for any $u \in \mathbb{R}$

$$R_T(\boldsymbol{u}) \leq \sum_{t=1}^{T}\Big(w_t\,\ell_t'(w_t) - \|\boldsymbol{u}\|_2\,\ell_t'(w_t)\Big) + \|\boldsymbol{u}\|_2 \sum_{t=1}^{T}\left(\boldsymbol{g}_t^\top \boldsymbol{v}_t - \boldsymbol{g}_t^\top \frac{\boldsymbol{u}}{\|\boldsymbol{u}\|_2}\right) \quad \text{(for any } \boldsymbol{u} \in \mathbb{R}^d\text{)}$$

$$= \mathcal{O}\Big(\|\boldsymbol{u}\|_2\sqrt{T\ln\left(\|\boldsymbol{u}\|_2^2\,T+1\right)}+1\Big) + \mathcal{O}\big(\|\boldsymbol{u}\|_2\sqrt{T}\big)$$

$$= \mathcal{O}\Big(\|\boldsymbol{u}\|_2\sqrt{T\ln\left(\|\boldsymbol{u}\|_2^2\,T+1\right)}+1\Big)$$

# Regret bound

- Betting strategy: $\alpha_1 = 0$ and $\alpha_t = (x_1 + \cdots + x_{t-1})/t$ for $t \geq 1$ (Krichevsky-Trofimov estimator)
- Achieved wealth: $C_T \geq \frac{1}{\sqrt{T}} \exp\left(\frac{1}{2T}\left(\sum_{t=1}^{T} x_t\right)^2\right) = \phi\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right)$
- Resulting regret: $R_T(u) = \phi^*\left(-\sum_{t=1}^{T} \ell_t'(w_t)\right) = \mathcal{O}\left(|u|\sqrt{T\ln(u^2T+1)}+1\right)$ for any $u \in \mathbb{R}$

$$R_T(\boldsymbol{u}) \leq \sum_{t=1}^{T}\Big(w_t\,\ell_t'(w_t) - \|\boldsymbol{u}\|_2\,\ell_t'(w_t)\Big) + \|\boldsymbol{u}\|_2 \sum_{t=1}^{T}\left(\boldsymbol{g}_t^\top \boldsymbol{v}_t - \boldsymbol{g}_t^\top \frac{\boldsymbol{u}}{\|\boldsymbol{u}\|_2}\right) \qquad \text{(for any } \boldsymbol{u} \in \mathbb{R}^d\text{)}$$

$$= \mathcal{O}\Big(\|\boldsymbol{u}\|_2\sqrt{T\ln\big(\|\boldsymbol{u}\|_2^2\,T+1\big)}+1\Big) + \mathcal{O}\big(\|\boldsymbol{u}\|_2\sqrt{T}\big)$$

$$= \mathcal{O}\Big(\|\boldsymbol{u}\|_2\sqrt{T\ln\big(\|\boldsymbol{u}\|_2^2\,T+1\big)}+1\Big)$$

Result matches the $\|\boldsymbol{u}\|_2\sqrt{T}$ bound up to log factors that are unavoidable if $\|\boldsymbol{u}\|_2$ is unknown

# Other notions of regret

# Other notions of regret

- If the loss sequence $\ell_1, \ell_2, \dots$ is such that no $\boldsymbol{u} \in \mathbb{V}$ achieves a small cumulative loss $\ell_1(\boldsymbol{u}) + \ell_2(\boldsymbol{u}) + \cdots$, then regret bounds are meaningless

# Other notions of regret

- If the loss sequence $\ell_1, \ell_2, \dots$ is such that no $\boldsymbol{u} \in \mathbb{V}$ achieves a small cumulative loss $\ell_1(\boldsymbol{u}) + \ell_2(\boldsymbol{u}) + \cdots$, then regret bounds are meaningless
- Lack of a single good minimizer in $\mathbb{V}$ caused by a highly nonstationary data sequence

# Other notions of regret

- If the loss sequence $\ell_1, \ell_2, \dots$ is such that no $\boldsymbol{u} \in \mathbb{V}$ achieves a small cumulative loss $\ell_1(\boldsymbol{u}) + \ell_2(\boldsymbol{u}) + \cdots$, then regret bounds are meaningless
- Lack of a single good minimizer in $\mathbb{V}$ caused by a highly nonstationary data sequence
- In this case, the regret should be replaced by more robust measures

# Other notions of regret

- If the loss sequence $\ell_1, \ell_2, \dots$ is such that no $\boldsymbol{u} \in \mathbb{V}$ achieves a small cumulative loss $\ell_1(\boldsymbol{u}) + \ell_2(\boldsymbol{u}) + \cdots$, then regret bounds are meaningless
- Lack of a single good minimizer in $\mathbb{V}$ caused by a highly nonstationary data sequence
- In this case, the regret should be replaced by more robust measures
- Dynamic regret $\quad R_T^{\mathrm{dyn}}(\boldsymbol{u}_1, \dots, \boldsymbol{u}_T) = \sum_{t=1}^{T} \ell_t(\boldsymbol{w}_t) - \sum_{t=1}^{T} \ell_t(\boldsymbol{u}_t) \quad$ where $\boldsymbol{u}_1, \dots, \boldsymbol{u}_T \in \mathbb{V}$

# Other notions of regret

- If the loss sequence $\ell_1, \ell_2, \dots$ is such that no $\boldsymbol{u} \in \mathbb{V}$ achieves a small cumulative loss $\ell_1(\boldsymbol{u}) + \ell_2(\boldsymbol{u}) + \cdots$, then regret bounds are meaningless
- Lack of a single good minimizer in $\mathbb{V}$ caused by a highly nonstationary data sequence
- In this case, the regret should be replaced by more robust measures
- Dynamic regret $\quad R_T^{\mathrm{dyn}}(\boldsymbol{u}_1, \dots, \boldsymbol{u}_T) = \sum_{t=1}^{T} \ell_t(\boldsymbol{w}_t) - \sum_{t=1}^{T} \ell_t(\boldsymbol{u}_t) \quad$ where $\boldsymbol{u}_1, \dots, \boldsymbol{u}_T \in \mathbb{V}$
- Complexity parameter: $\Pi_T = \sum_{t=1}^{T-1} \|\boldsymbol{u}_{t+1} - \boldsymbol{u}_t\|$

# Other notions of regret

- If the loss sequence $\ell_1, \ell_2, \dots$ is such that no $\boldsymbol{u} \in \mathbb{V}$ achieves a small cumulative loss $\ell_1(\boldsymbol{u}) + \ell_2(\boldsymbol{u}) + \cdots$, then regret bounds are meaningless
- Lack of a single good minimizer in $\mathbb{V}$ caused by a highly nonstationary data sequence
- In this case, the regret should be replaced by more robust measures
- Dynamic regret $\quad R_T^{\text{dyn}}(\boldsymbol{u}_1, \dots, \boldsymbol{u}_T) = \sum_{t=1}^{T} \ell_t(\boldsymbol{w}_t) - \sum_{t=1}^{T} \ell_t(\boldsymbol{u}_t) \quad$ where $\boldsymbol{u}_1, \dots, \boldsymbol{u}_T \in \mathbb{V}$
- Complexity parameter: $\Pi_T = \sum_{t=1}^{T-1} \|\boldsymbol{u}_{t+1} - \boldsymbol{u}_t\|$
- Lower bound: $\Omega\big(G\sqrt{(D + \Pi_T)DT}\big)$

# Other notions of regret

- If the loss sequence $\ell_1, \ell_2, \dots$ is such that no $\boldsymbol{u} \in \mathbb{V}$ achieves a small cumulative loss $\ell_1(\boldsymbol{u}) + \ell_2(\boldsymbol{u}) + \cdots$, then regret bounds are meaningless
- Lack of a single good minimizer in $\mathbb{V}$ caused by a highly nonstationary data sequence
- In this case, the regret should be replaced by more robust measures
- Dynamic regret $\quad R_T^{\text{dyn}}(\boldsymbol{u}_1, \dots, \boldsymbol{u}_T) = \sum_{t=1}^{T} \ell_t(\boldsymbol{w}_t) - \sum_{t=1}^{T} \ell_t(\boldsymbol{u}_t) \quad$ where $\boldsymbol{u}_1, \dots, \boldsymbol{u}_T \in \mathbb{V}$
- Complexity parameter: $\Pi_T = \sum_{t=1}^{T-1} \|\boldsymbol{u}_{t+1} - \boldsymbol{u}_t\|$
- Lower bound: $\Omega\big(G\sqrt{(D + \Pi_T)DT}\big)$
- When $\Pi_T = 0$ this reduces to the standard lower bound $\Omega(GD\sqrt{T})$

# Other notions of regret

- If the loss sequence $\ell_1, \ell_2, \dots$ is such that no $\boldsymbol{u} \in \mathbb{V}$ achieves a small cumulative loss $\ell_1(\boldsymbol{u}) + \ell_2(\boldsymbol{u}) + \cdots$, then regret bounds are meaningless
- Lack of a single good minimizer in $\mathbb{V}$ caused by a highly nonstationary data sequence
- In this case, the regret should be replaced by more robust measures
- Dynamic regret $\quad R_T^{\mathrm{dyn}}(\boldsymbol{u}_1, \dots, \boldsymbol{u}_T) = \displaystyle\sum_{t=1}^{T} \ell_t(\boldsymbol{w}_t) - \sum_{t=1}^{T} \ell_t(\boldsymbol{u}_t) \quad$ where $\boldsymbol{u}_1, \dots, \boldsymbol{u}_T \in \mathbb{V}$
- Complexity parameter: $\Pi_T = \displaystyle\sum_{t=1}^{T-1} \|\boldsymbol{u}_{t+1} - \boldsymbol{u}_t\|$
- Lower bound: $\Omega\big(G\sqrt{(D + \Pi_T)DT}\big)$
- When $\Pi_T = 0$ this reduces to the standard lower bound $\Omega(GD\sqrt{T})$
- Matching upper bound obtained by using Hedge to aggregate $\mathcal{O}(\ln T)$ instances of FTRL each tuned to a different value of $\Pi_T$

# Adaptive regret

# Adaptive regret

- Evaluate the performance of the online algorithm against that of the best fixed comparator in any interval of time

# Adaptive regret

- Evaluate the performance of the online algorithm against that of the best fixed comparator in any interval of time
- $R^{\text{ada}}_{\tau,T} = \max_{s=1,\dots,T-\tau+1} \left( \sum_{t=s}^{s+\tau-1} \ell_t(\boldsymbol{w}_t) - \min_{\boldsymbol{u}\in\mathbb{V}} \sum_{t=s}^{s+\tau-1} \ell_t(\boldsymbol{u}) \right)$ where $\tau \in \{1,\dots,T\}$

# Adaptive regret

- Evaluate the performance of the online algorithm against that of the best fixed comparator in any interval of time
- $R^{\text{ada}}_{\tau,T} = \max_{s=1,\dots,T-\tau+1} \left( \sum_{t=s}^{s+\tau-1} \ell_t(\boldsymbol{w}_t) - \min_{\boldsymbol{u}\in\mathbb{V}} \sum_{t=s}^{s+\tau-1} \ell_t(\boldsymbol{u}) \right)$ where $\tau \in \{1,\dots,T\}$
- Best known upper bound: $R^{\text{ada}}_{\tau,T}(\boldsymbol{u}) = \mathcal{O}\big(DG\sqrt{\tau} + \sqrt{(\ln T)\tau}\big)$

# Adaptive regret

- Evaluate the performance of the online algorithm against that of the best fixed comparator in any interval of time
- $R^{\text{ada}}_{\tau,T} = \max_{s=1,\dots,T-\tau+1} \left( \sum_{t=s}^{s+\tau-1} \ell_t(\boldsymbol{w}_t) - \min_{\boldsymbol{u}\in\mathbb{V}} \sum_{t=s}^{s+\tau-1} \ell_t(\boldsymbol{u}) \right)$ where $\tau \in \{1,\dots,T\}$
- Best known upper bound: $R^{\text{ada}}_{\tau,T}(\boldsymbol{u}) = \mathcal{O}\big(DG\sqrt{\tau} + \sqrt{(\ln T)\tau}\big)$
- Obtained by combining several instances of a standard online algorithm each run in a specific interval of time

# Adaptive regret

- Evaluate the performance of the online algorithm against that of the best fixed comparator in any interval of time
- $R^{\text{ada}}_{\tau,T} = \max_{s=1,\dots,T-\tau+1} \left( \sum_{t=s}^{s+\tau-1} \ell_t(\boldsymbol{w}_t) - \min_{\boldsymbol{u}\in\mathbb{V}} \sum_{t=s}^{s+\tau-1} \ell_t(\boldsymbol{u}) \right)$ where $\tau \in \{1,\dots,T\}$
- Best known upper bound: $R^{\text{ada}}_{\tau,T}(\boldsymbol{u}) = \mathcal{O}\big(DG\sqrt{\tau} + \sqrt{(\ln T)\tau}\big)$
- Obtained by combining several instances of a standard online algorithm each run in a specific interval of time
- The set of intervals is carefully designed so that the overall number of instances to be run is $\mathcal{O}(\ln T)$

# From sequential to statistical learning

- Statistical risk for a convex and bounded loss $\ell_{\mathcal{D}}(\boldsymbol{w}) = \mathbb{E}\Big[\ell(\boldsymbol{w}^\top \boldsymbol{X}, Y)\Big]$

# From sequential to statistical learning

- Statistical risk for a convex and bounded loss $\ell_{\mathcal{D}}(\boldsymbol{w}) = \mathbb{E}\Big[\ell(\boldsymbol{w}^\top \boldsymbol{X}, Y)\Big]$
- Let $\overline{\boldsymbol{w}} = \dfrac{1}{T}\displaystyle\sum_{t=1}^{T} \boldsymbol{w}_t$ where $\boldsymbol{w}_1, \dots, \boldsymbol{w}_T$ are generated by an online algorithm over $(\boldsymbol{X}_1, Y_1), (\boldsymbol{X}_2, Y_2), \dots$ drawn i.i.d. from an unknown distribution $\mathcal{D}$

# From sequential to statistical learning

- Statistical risk for a convex and bounded loss $\ell_{\mathcal{D}}(\boldsymbol{w}) = \mathbb{E}\Big[\ell(\boldsymbol{w}^\top \boldsymbol{X}, Y)\Big]$
- Let $\overline{\boldsymbol{w}} = \dfrac{1}{T}\sum_{t=1}^{T} \boldsymbol{w}_t$ where $\boldsymbol{w}_1, \ldots, \boldsymbol{w}_T$ are generated by an online algorithm over $(\boldsymbol{X}_1, Y_1), (\boldsymbol{X}_2, Y_2), \ldots$ drawn i.i.d. from an unknown distribution $\mathcal{D}$
- By Jensen's inequality

$$\ell_{\mathcal{D}}(\overline{\boldsymbol{w}}) = \mathbb{E}\Big[\ell(\overline{\boldsymbol{w}}^\top \boldsymbol{X}, Y)\Big] \le \mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T} \ell(\boldsymbol{w}_t^\top \boldsymbol{X}_t, Y_t)\right] = \frac{1}{T}\sum_{t=1}^{T} \ell_{\mathcal{D}}(\boldsymbol{w}_t)$$

# From sequential to statistical learning

- Statistical risk for a convex and bounded loss $\ell_{\mathcal{D}}(\boldsymbol{w}) = \mathbb{E}\Big[\ell(\boldsymbol{w}^\top \boldsymbol{X}, Y)\Big]$
- Let $\overline{\boldsymbol{w}} = \frac{1}{T}\sum_{t=1}^{T} \boldsymbol{w}_t$ where $\boldsymbol{w}_1, \dots, \boldsymbol{w}_T$ are generated by an online algorithm over $(\boldsymbol{X}_1, Y_1), (\boldsymbol{X}_2, Y_2), \dots$ drawn i.i.d. from an unknown distribution $\mathcal{D}$
- By Jensen's inequality

$$\ell_{\mathcal{D}}(\overline{\boldsymbol{w}}) = \mathbb{E}\Big[\ell(\overline{\boldsymbol{w}}^\top \boldsymbol{X}, Y)\Big] \le \mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T} \ell(\boldsymbol{w}_t^\top \boldsymbol{X}_t, Y_t)\right] = \frac{1}{T}\sum_{t=1}^{T} \ell_{\mathcal{D}}(\boldsymbol{w}_t)$$

- Note also that $\mathbb{E}\Big[\ell_{\mathcal{D}}(\boldsymbol{w}_t) - \ell(\boldsymbol{w}_t^\top \boldsymbol{X}_t, Y_t) \,\Big|\, (\boldsymbol{X}_1, Y_1), \dots, (\boldsymbol{X}_{t-1}, Y_{t-1})\Big] = 0$

# From sequential to statistical learning

- Statistical risk for a convex and bounded loss $\ell_{\mathcal{D}}(\boldsymbol{w}) = \mathbb{E}\Big[\ell(\boldsymbol{w}^\top \boldsymbol{X}, Y)\Big]$
- Let $\overline{\boldsymbol{w}} = \dfrac{1}{T}\displaystyle\sum_{t=1}^{T} \boldsymbol{w}_t$ where $\boldsymbol{w}_1, \dots, \boldsymbol{w}_T$ are generated by an online algorithm over $(\boldsymbol{X}_1, Y_1), (\boldsymbol{X}_2, Y_2), \dots$ drawn i.i.d. from an unknown distribution $\mathcal{D}$
- By Jensen's inequality

$$\ell_{\mathcal{D}}(\overline{\boldsymbol{w}}) = \mathbb{E}\Big[\ell(\overline{\boldsymbol{w}}^\top \boldsymbol{X}, Y)\Big] \le \mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T} \ell(\boldsymbol{w}_t^\top \boldsymbol{X}_t, Y_t)\right] = \frac{1}{T}\sum_{t=1}^{T} \ell_{\mathcal{D}}(\boldsymbol{w}_t)$$

- Note also that $\mathbb{E}\Big[\ell_{\mathcal{D}}(\boldsymbol{w}_t) - \ell(\boldsymbol{w}_t^\top \boldsymbol{X}_t, Y_t) \,\Big|\, (\boldsymbol{X}_1, Y_1), \dots, (\boldsymbol{X}_{t-1}, Y_{t-1})\Big] = 0$
- Using concentration inequalities for martingales (e.g., Hoeffding-Azuma inequality),

$$\frac{1}{T}\sum_{t=1}^{T} \ell_{\mathcal{D}}(\boldsymbol{w}_t) \le \frac{1}{T}\sum_{t=1}^{T} \ell(\boldsymbol{w}_t^\top \boldsymbol{X}_t, Y_t) + \mathcal{O}\left(\frac{1}{\sqrt{T}}\right) \quad \text{w.h.p.}$$

# Statistical risk bounds from regret bounds

Letting $\ell(\boldsymbol{w}^\top \boldsymbol{X}_t, Y_t) = \ell_t(\boldsymbol{w})$ we have $\ell_{\mathcal{D}}(\overline{\boldsymbol{w}}) \le \frac{1}{T}\sum_{t=1}^{T} \ell_t(\boldsymbol{w}_t) + \mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$ w.h.p.

# Statistical risk bounds from regret bounds

Letting $\ell(\boldsymbol{w}^\top \boldsymbol{X}_t, Y_t) = \ell_t(\boldsymbol{w})$ we have $\ell_{\mathcal{D}}(\overline{\boldsymbol{w}}) \le \frac{1}{T}\sum_{t=1}^{T} \ell_t(\boldsymbol{w}_t) + \mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$ w.h.p.

Let $\boldsymbol{u} = \operatorname*{argmin}_{\boldsymbol{w}\in\mathbb{V}} \ell_{\mathcal{D}}(\boldsymbol{w})$ and bound the estimation error $\ell_{\mathcal{D}}(\overline{\boldsymbol{w}}) - \ell_{\mathcal{D}}(\boldsymbol{u})$ w.h.p.

# Statistical risk bounds from regret bounds

Letting $\ell(\boldsymbol{w}^\top \boldsymbol{X}_t, Y_t) = \ell_t(\boldsymbol{w})$ we have $\ell_{\mathcal{D}}(\overline{\boldsymbol{w}}) \le \frac{1}{T}\sum_{t=1}^{T} \ell_t(\boldsymbol{w}_t) + \mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$ w.h.p.

Let $\boldsymbol{u} = \operatorname*{argmin}_{\boldsymbol{w}\in\mathbb{V}} \ell_{\mathcal{D}}(\boldsymbol{w})$ and bound the estimation error $\ell_{\mathcal{D}}(\overline{\boldsymbol{w}}) - \ell_{\mathcal{D}}(\boldsymbol{u})$ w.h.p.

$$\frac{1}{T}\sum_{t=1}^{T} \ell_t(\boldsymbol{w}_t) \le \frac{1}{T} \inf_{\boldsymbol{w}\in\mathbb{V}} \sum_{t=1}^{T} \ell_t(\boldsymbol{w}) + \frac{2GD}{\sqrt{T}} \qquad \text{(regret bound)}$$

# Statistical risk bounds from regret bounds

Letting $\ell(\boldsymbol{w}^\top \boldsymbol{X}_t, Y_t) = \ell_t(\boldsymbol{w})$ we have $\ell_{\mathcal{D}}(\overline{\boldsymbol{w}}) \le \frac{1}{T}\sum_{t=1}^{T} \ell_t(\boldsymbol{w}_t) + \mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$ w.h.p.

Let $\boldsymbol{u} = \operatorname*{argmin}_{\boldsymbol{w}\in\mathbb{V}} \ell_{\mathcal{D}}(\boldsymbol{w})$ and bound the estimation error $\ell_{\mathcal{D}}(\overline{\boldsymbol{w}}) - \ell_{\mathcal{D}}(\boldsymbol{u})$ w.h.p.

$$\frac{1}{T}\sum_{t=1}^{T} \ell_t(\boldsymbol{w}_t) \le \frac{1}{T} \inf_{\boldsymbol{w}\in\mathbb{V}} \sum_{t=1}^{T} \ell_t(\boldsymbol{w}) + \frac{2GD}{\sqrt{T}} \qquad \text{(regret bound)}$$

$$\le \frac{1}{T}\sum_{t=1}^{T} \ell_t(\boldsymbol{u}) + \frac{2GD}{\sqrt{T}}$$

# Statistical risk bounds from regret bounds

Letting $\ell(\boldsymbol{w}^\top \boldsymbol{X}_t, Y_t) = \ell_t(\boldsymbol{w})$ we have $\ell_{\mathcal{D}}(\overline{\boldsymbol{w}}) \le \frac{1}{T}\sum_{t=1}^{T} \ell_t(\boldsymbol{w}_t) + \mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$ w.h.p.

Let $\boldsymbol{u} = \operatorname*{argmin}_{\boldsymbol{w}\in\mathbb{V}} \ell_{\mathcal{D}}(\boldsymbol{w})$ and bound the estimation error $\ell_{\mathcal{D}}(\overline{\boldsymbol{w}}) - \ell_{\mathcal{D}}(\boldsymbol{u})$ w.h.p.

$$
\begin{aligned}
\frac{1}{T}\sum_{t=1}^{T} \ell_t(\boldsymbol{w}_t) &\le \frac{1}{T}\inf_{\boldsymbol{w}\in\mathbb{V}} \sum_{t=1}^{T} \ell_t(\boldsymbol{w}) + \frac{2GD}{\sqrt{T}} && \text{(regret bound)} \\
&\le \frac{1}{T}\sum_{t=1}^{T} \ell_t(\boldsymbol{u}) + \frac{2GD}{\sqrt{T}} \\
&\le \ell_{\mathcal{D}}(\boldsymbol{u}) + \frac{2GD}{\sqrt{T}} + \mathcal{O}\left(\frac{1}{\sqrt{T}}\right) && \text{w.h.p.}
\end{aligned}
$$

using concentration of $\ell_t(\boldsymbol{u})$ around $\ell_{\mathcal{D}}(\boldsymbol{u})$

# Final bound

If $\overline{\boldsymbol{w}} = \dfrac{1}{T}\displaystyle\sum_{t=1}^{T} \boldsymbol{w}_t$ where $\boldsymbol{w}_1, \dots, \boldsymbol{w}_T$ are generated by an online algorithm over $(\boldsymbol{X}_1, Y_1), (\boldsymbol{X}_2, Y_2), \dots$ drawn i.i.d. from an unknown distribution $\mathcal{D}$, then

$$\ell_{\mathcal{D}}(\overline{\boldsymbol{w}}) - \inf_{\boldsymbol{w} \in \mathbb{V}} \ell_{\mathcal{D}}(\boldsymbol{w}) \le \frac{2GD}{\sqrt{T}} \qquad \text{w.h.p.}$$