The Role of Convexity in Data-Driven Decision-Making

Peyman Mohajerin Esfahani

Delft University of Technology

Workshop on Insurance and Financial Mathematics

Hannover, November 2024

Outline

- Data-Driven Decision-Making

  - Three research questions

- Online Convex Optimization

  - Algorithms and regret bounds

Outline

- Data-Driven Decision-Making
  - Three research questions
- Online Convex Optimization
  - Algorithms and regret bounds

Data-Driven Decision-Making

Loss function: $\ell(x, \xi) \in \mathbb{R}$

$x \in \mathbb{X}$ decision

$\xi \in (\Xi, \mathbb{P})$ uncertainty

Data-Driven Decision-Making

Loss function: $\ell(x, \xi) \in \mathbb{R}$

$$\underbrace{x \in \mathbb{X}}_{\text{decision}} \qquad \underbrace{\xi \in (\Xi, \mathbb{P})}_{\text{uncertainty}}$$

$$\underbrace{(\widehat{\xi}_1, \ldots, \widehat{\xi}_N)}_{\text{past}} \rightsquigarrow \underbrace{\widehat{x}_N}_{\text{present}} \rightsquigarrow \underbrace{\ell(\widehat{x}_N, \xi)}_{\text{future}}$$

Data-Driven Decision-Making

$$\underbrace{(\widehat{\xi}_1, \ldots, \widehat{\xi}_N)}_{\text{past}} \rightsquigarrow \underbrace{\widehat{x}_N}_{\text{present}} \rightsquigarrow \underbrace{\ell(\widehat{x}_N, \xi)}_{\text{future}}$$

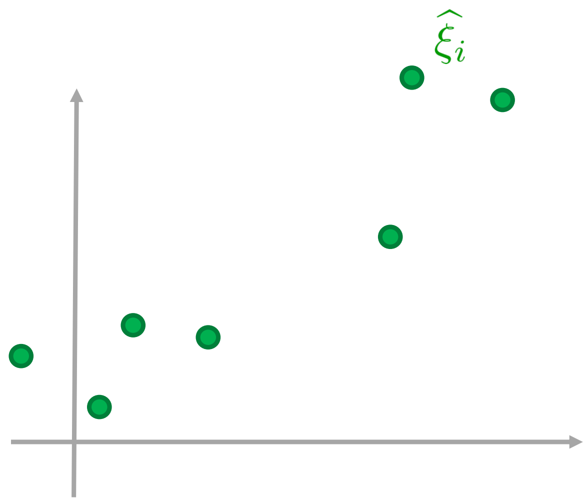

Data-Driven Decision-Making

$$\underbrace{(\widehat{\xi}_1, \ldots, \widehat{\xi}_N)}_{\text{past}} \rightsquigarrow \underbrace{\widehat{x}_N}_{\text{present}} \rightsquigarrow \underbrace{\ell(\widehat{x}_N, \xi)}_{\text{future}}$$

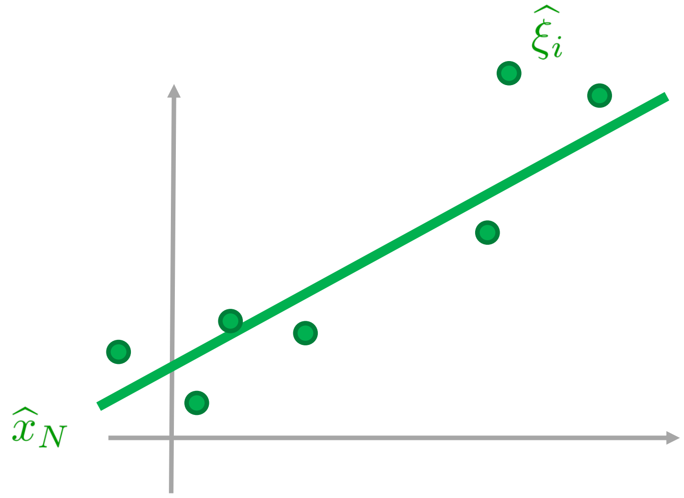

Data-Driven Decision-Making

$$\underbrace{(\widehat{\xi}_1, \ldots, \widehat{\xi}_N)}_{\text{past}} \rightsquigarrow \underbrace{\widehat{x}_N}_{\text{present}} \rightsquigarrow \underbrace{\ell(\widehat{x}_N, \xi)}_{\text{future}}$$

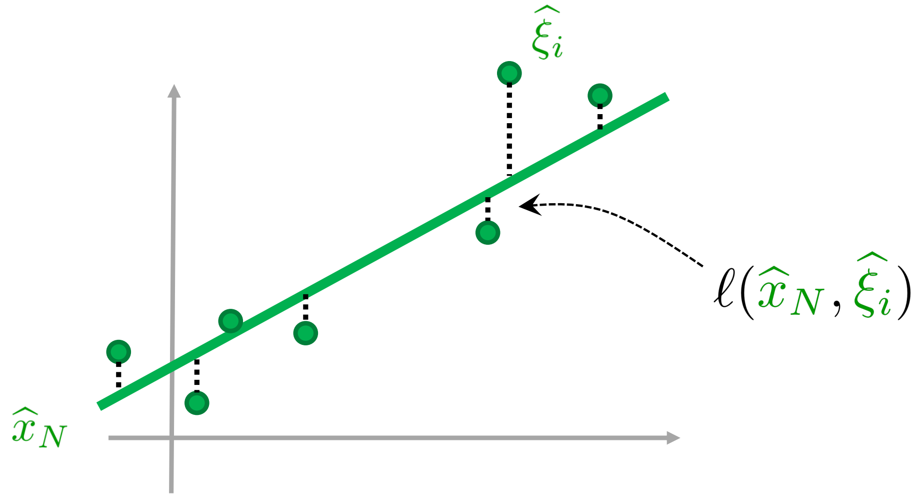

Data-Driven Decision-Making

$$\underbrace{(\widehat{\xi}_1, \ldots, \widehat{\xi}_N)}_{\text{past}} \rightsquigarrow \underbrace{\widehat{x}_N}_{\text{present}} \rightsquigarrow \underbrace{\ell(\widehat{x}_N, \xi)}_{\text{future}}$$

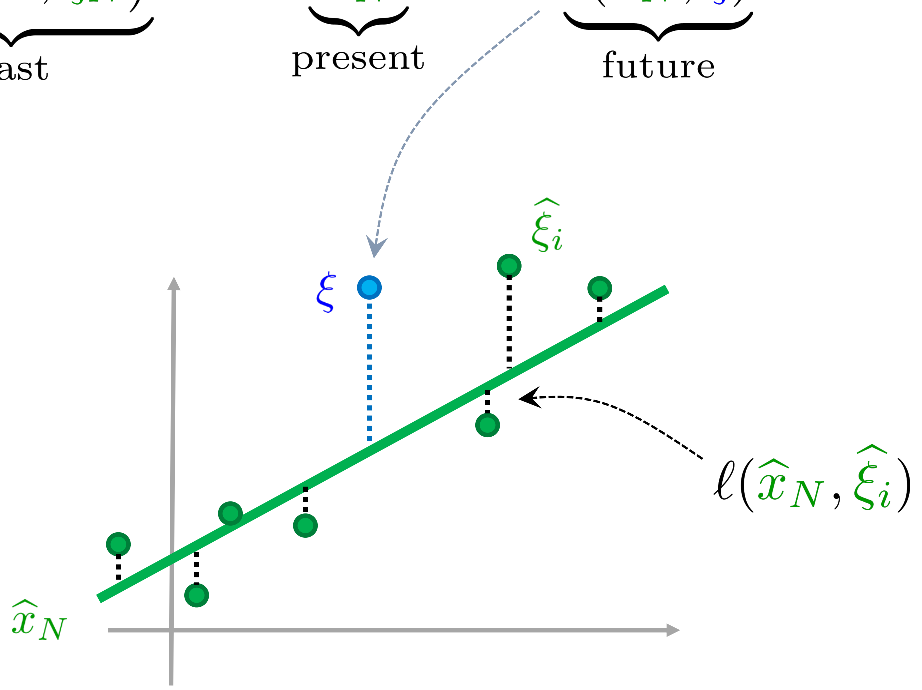

Data-Driven Decision-Making

$$\underbrace{(\widehat{\xi}_1,\ldots,\widehat{\xi}_N)}_{\text{past}} \overset{\text{(Prescriptive)}}{\rightsquigarrow} \underbrace{\widehat{x}_N}_{\text{present}} \overset{\text{(Predictive)}}{\rightsquigarrow} \underbrace{\ell(\widehat{x}_N,\xi)}_{\text{future}}$$

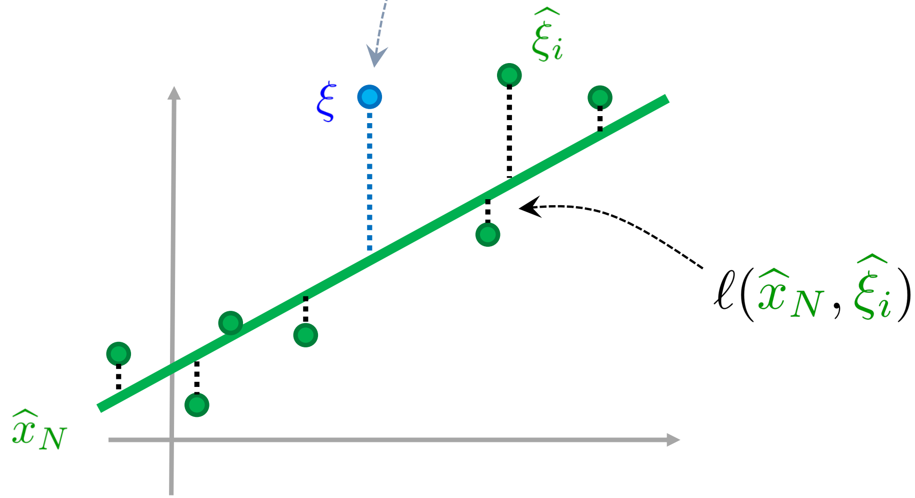

Data-Driven Decision-Making

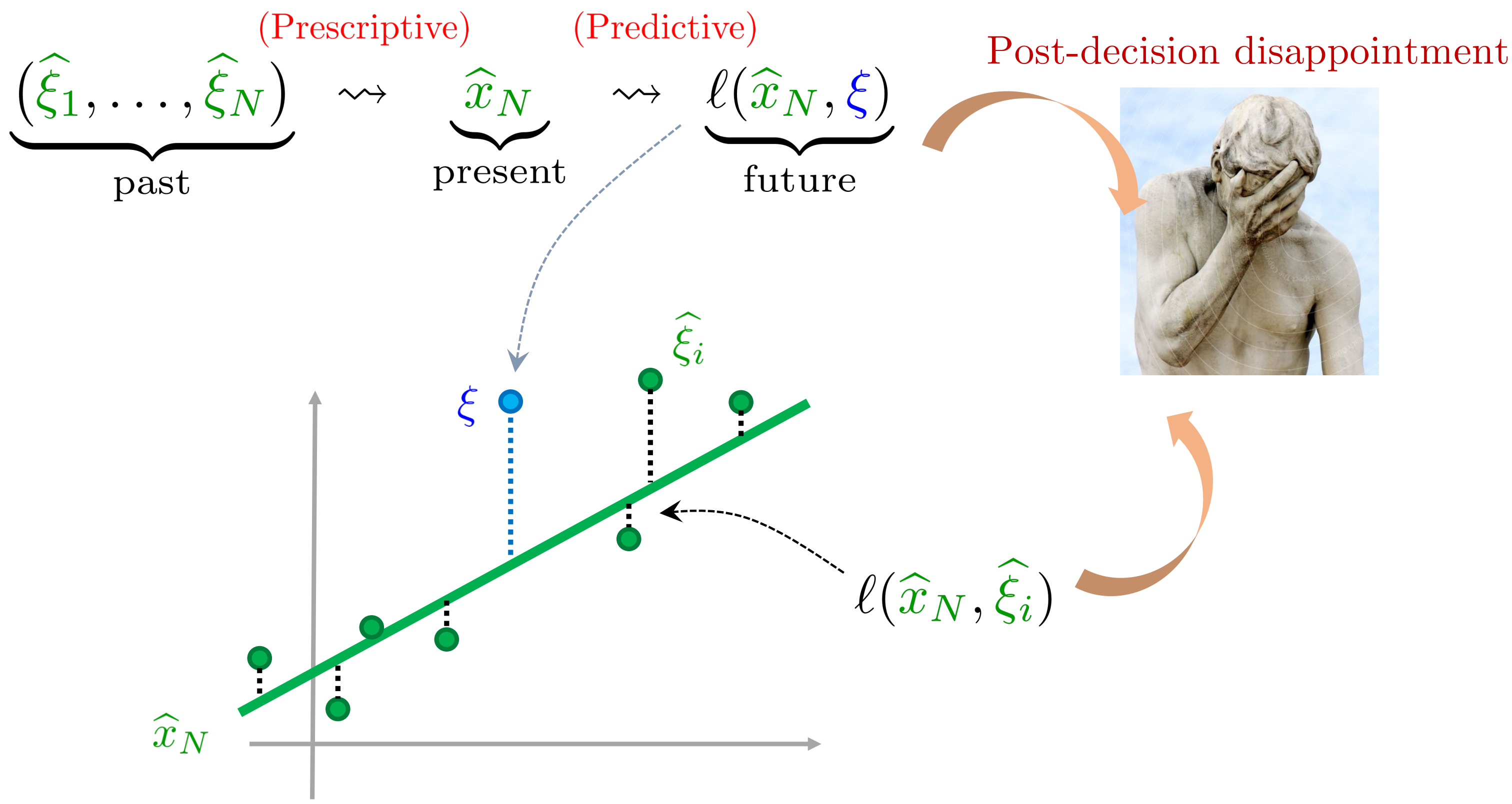

Data-Driven Decision-Making

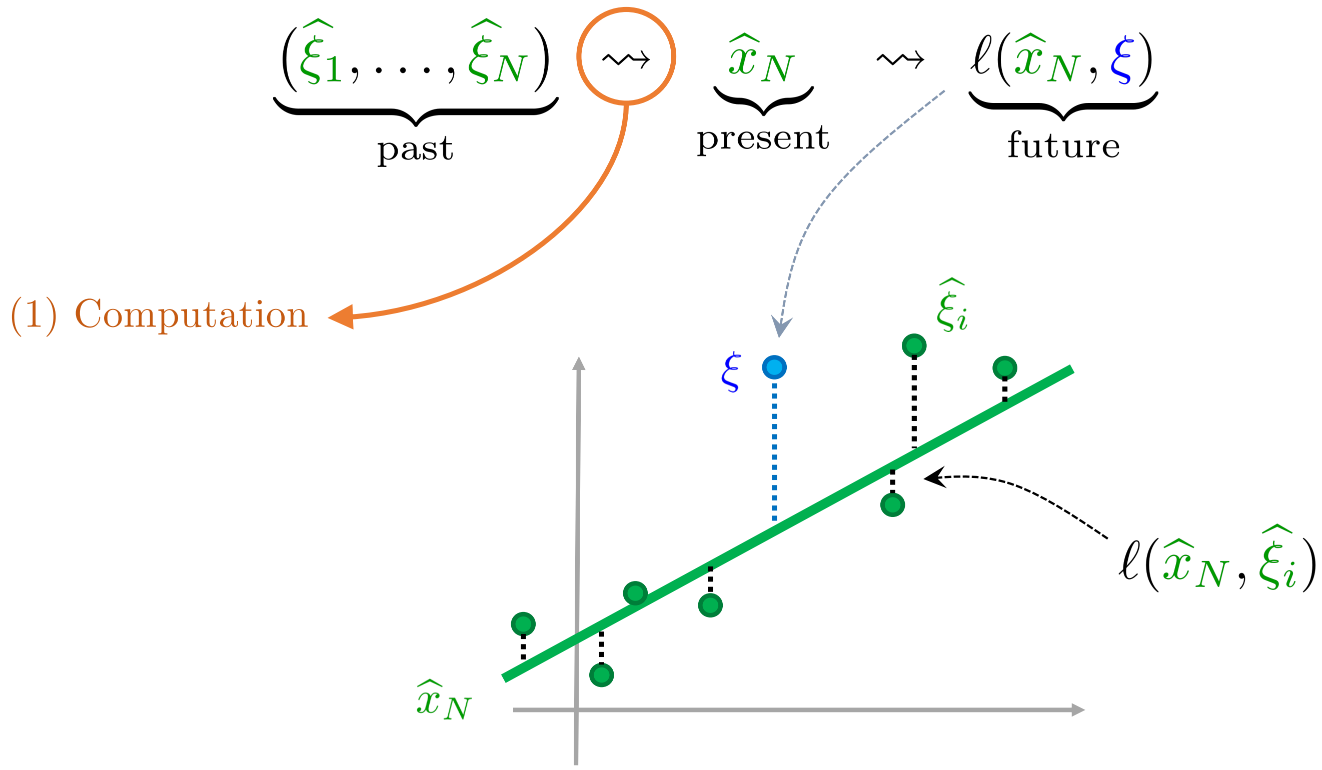

Data-Driven Decision-Making

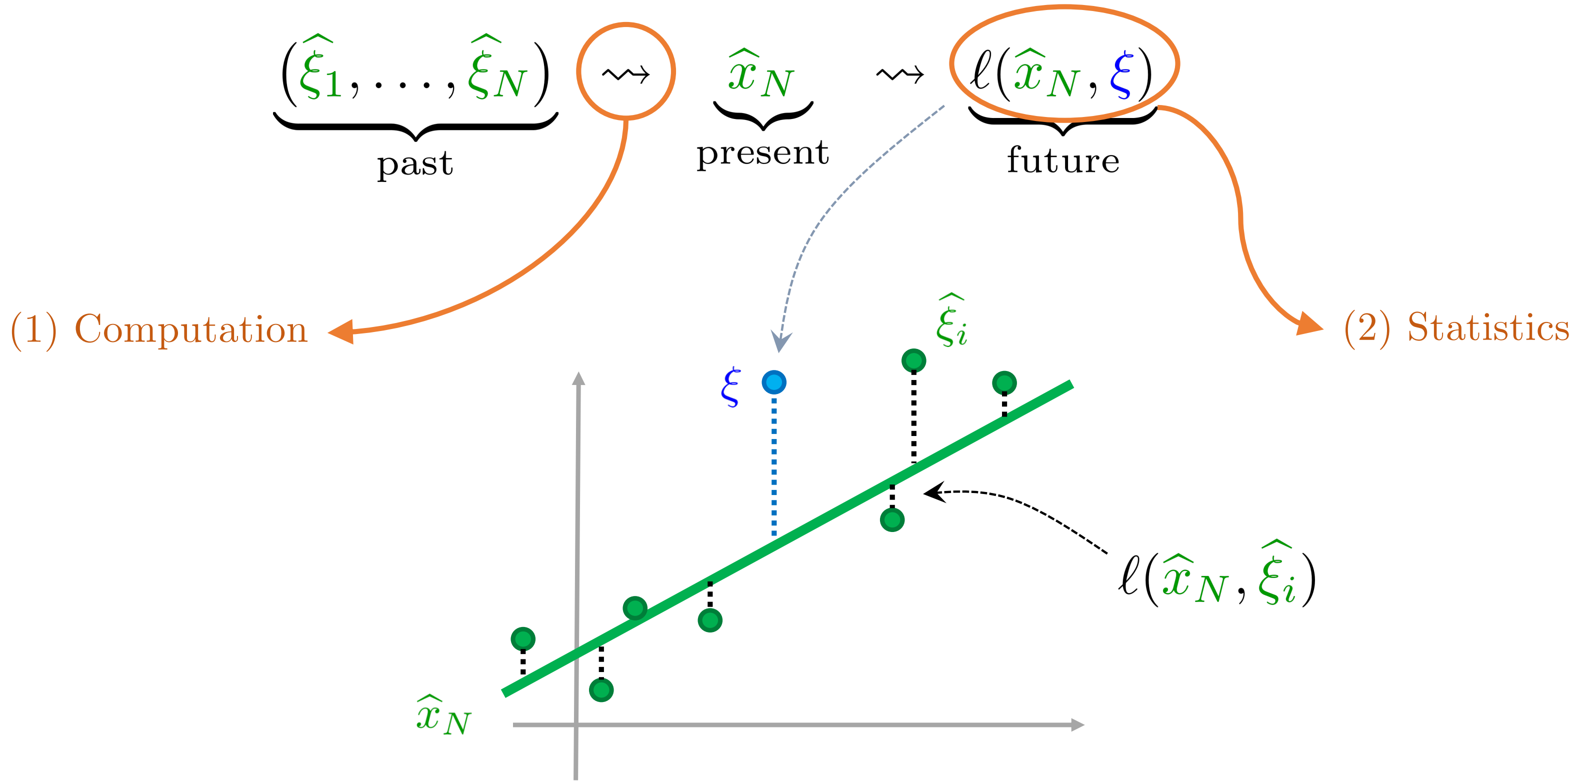

Data-Driven Decision-Making

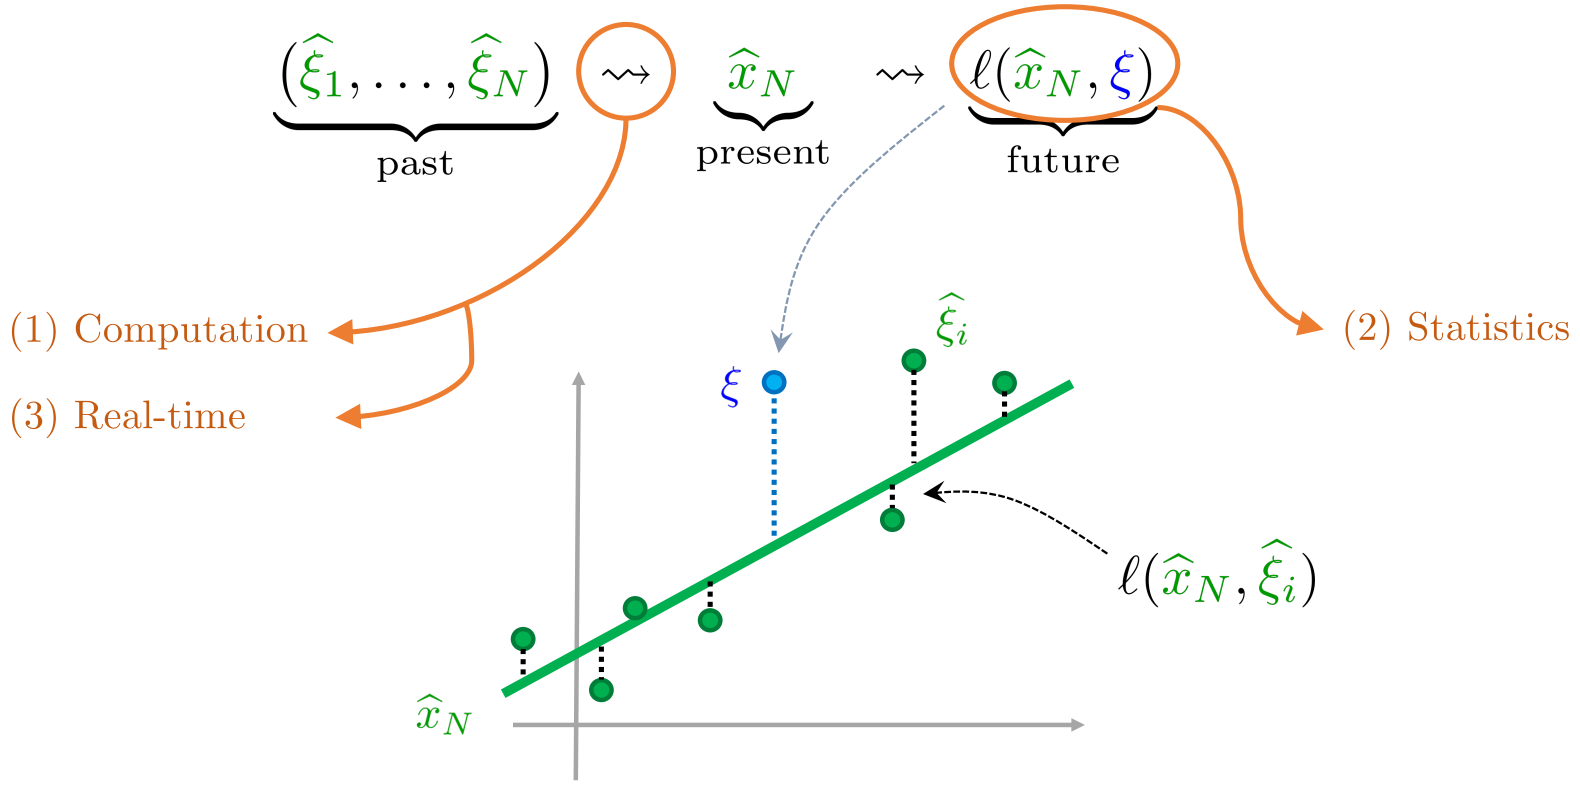

Data-Driven Decision-Making

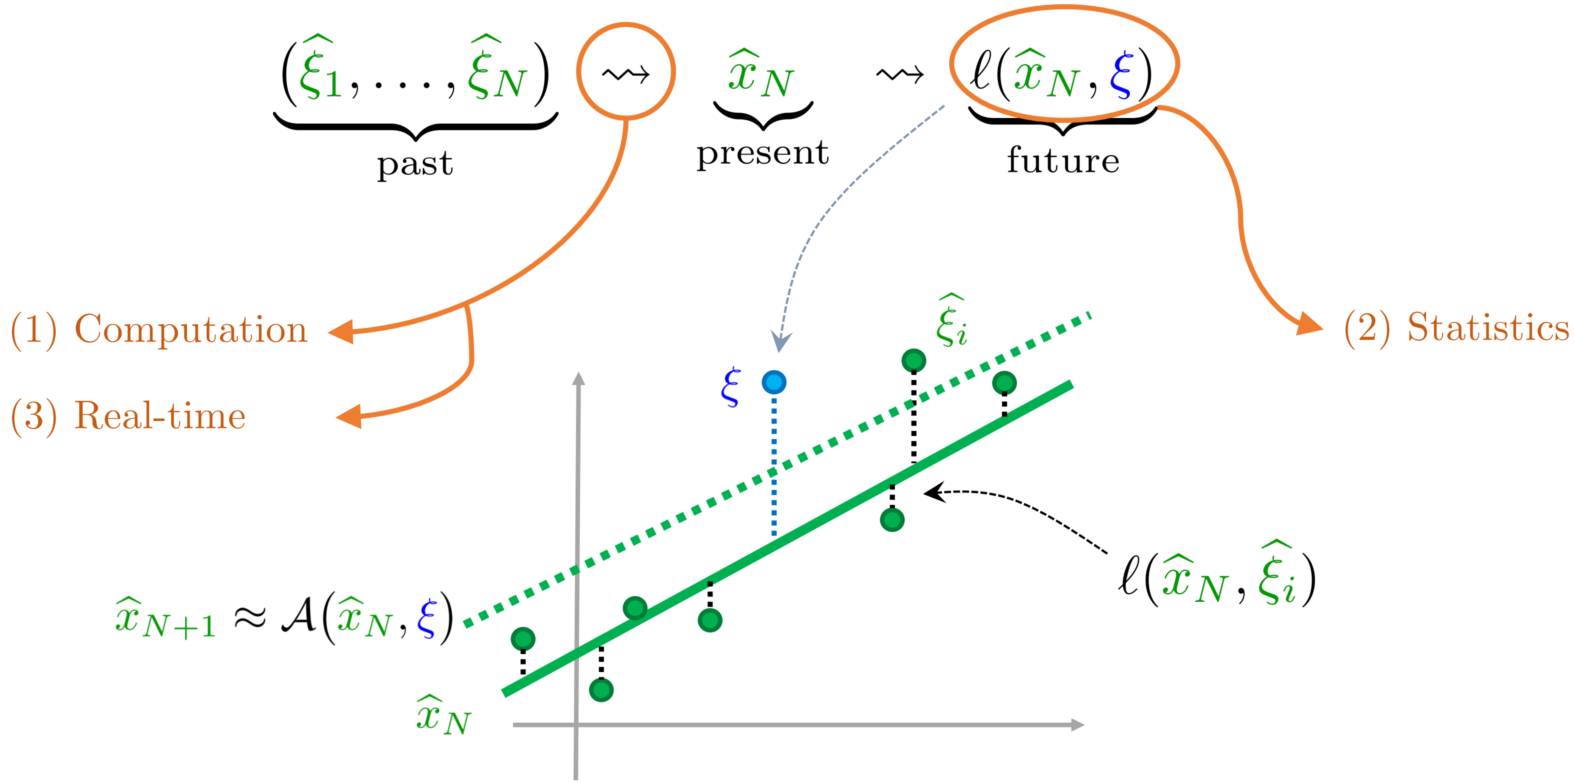

Data-Driven Decision-Making

The Role of Convexity

(1) Computational aspects

(2) Statistical aspects

(3) Real-time aspects

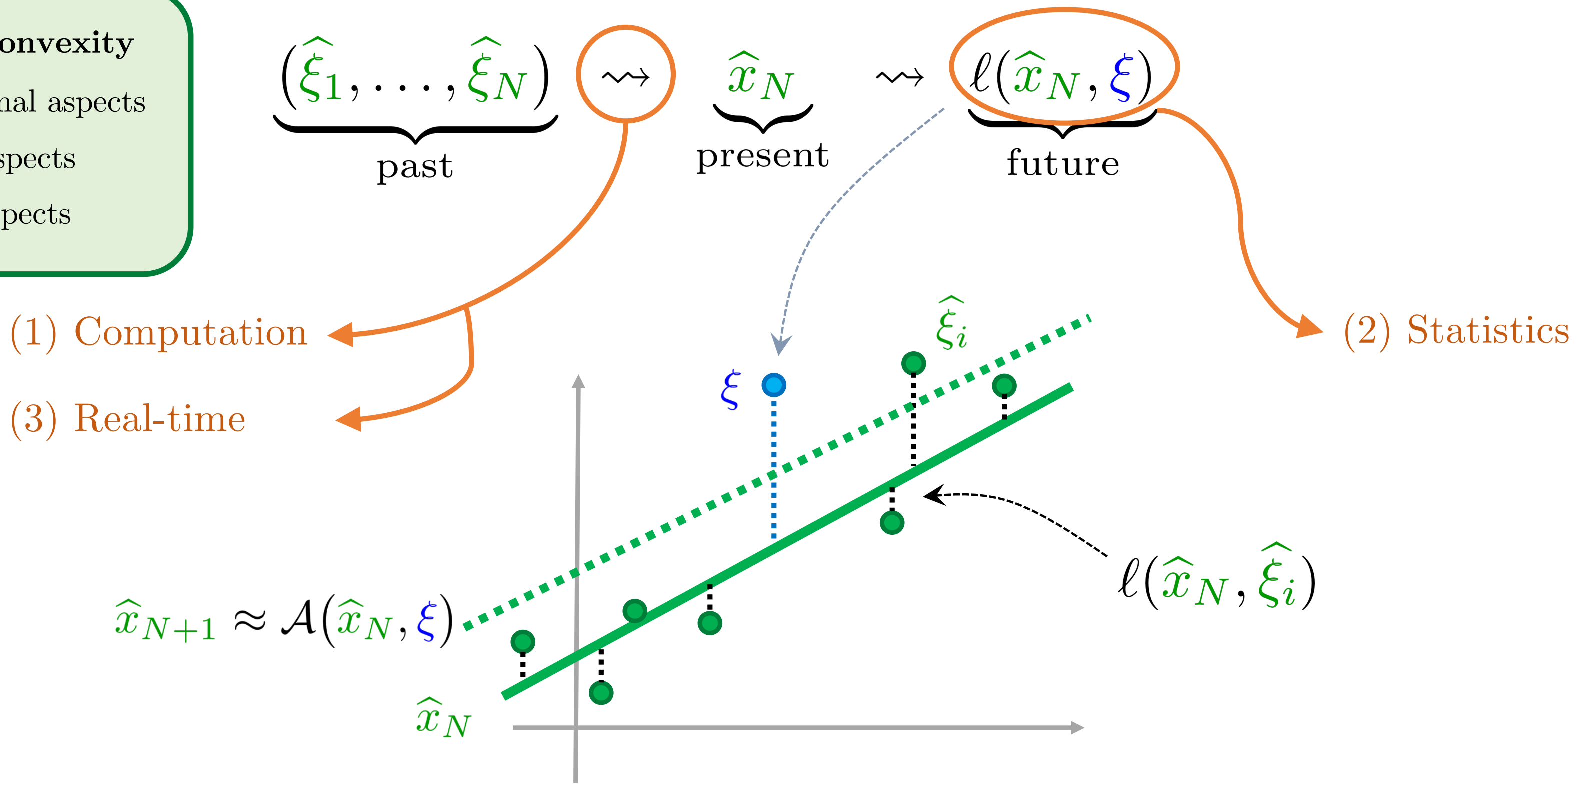

Data-Driven Decision-Making

The Role of Convexity

(1) Computational aspects

(2) Statistical aspects

(3) Real-time aspects

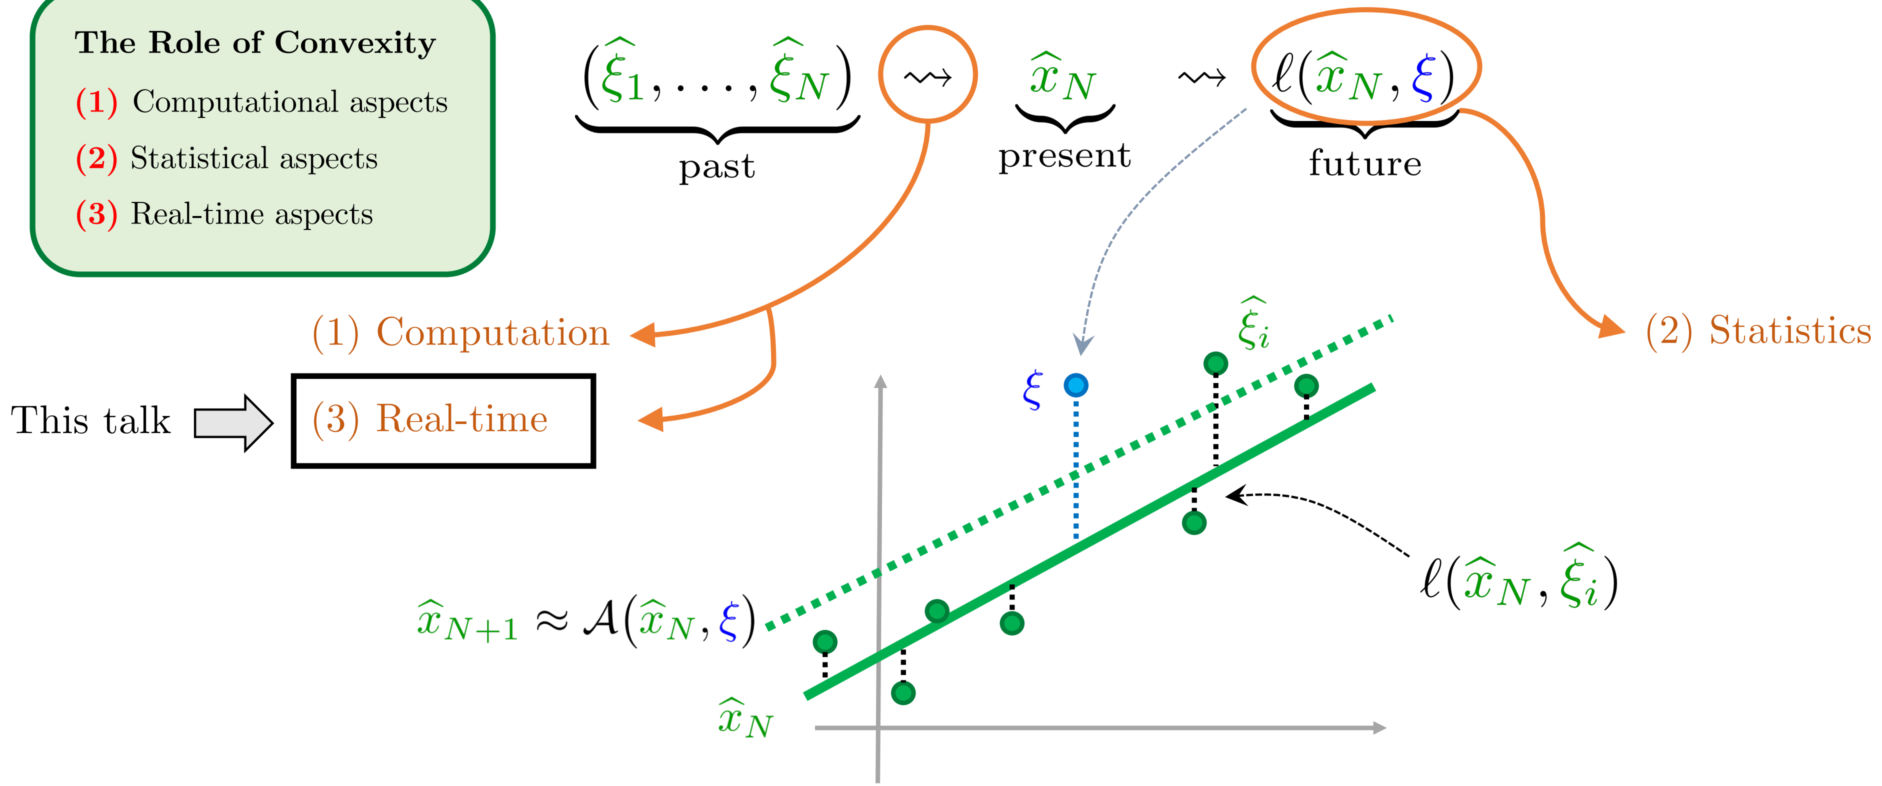

Outline

- Data-Driven Decision-Making
  - Three research questions
- Online Convex Optimization
  - Algorithms and regret bounds

Pedro Zattoni Scroccaro

Online Optimization: Setting

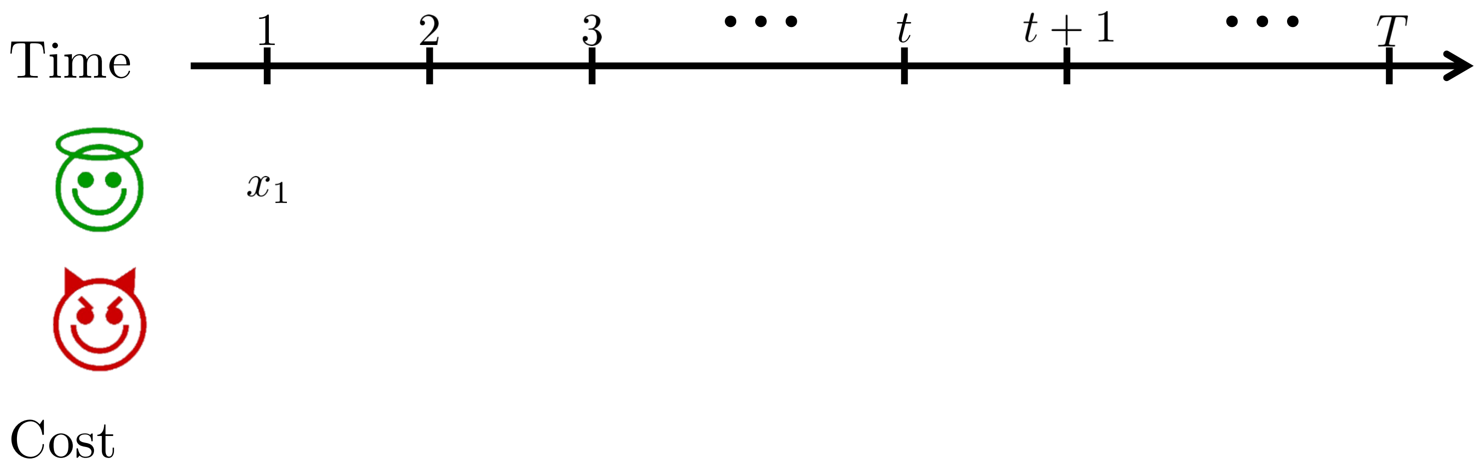

Online Optimization: Setting

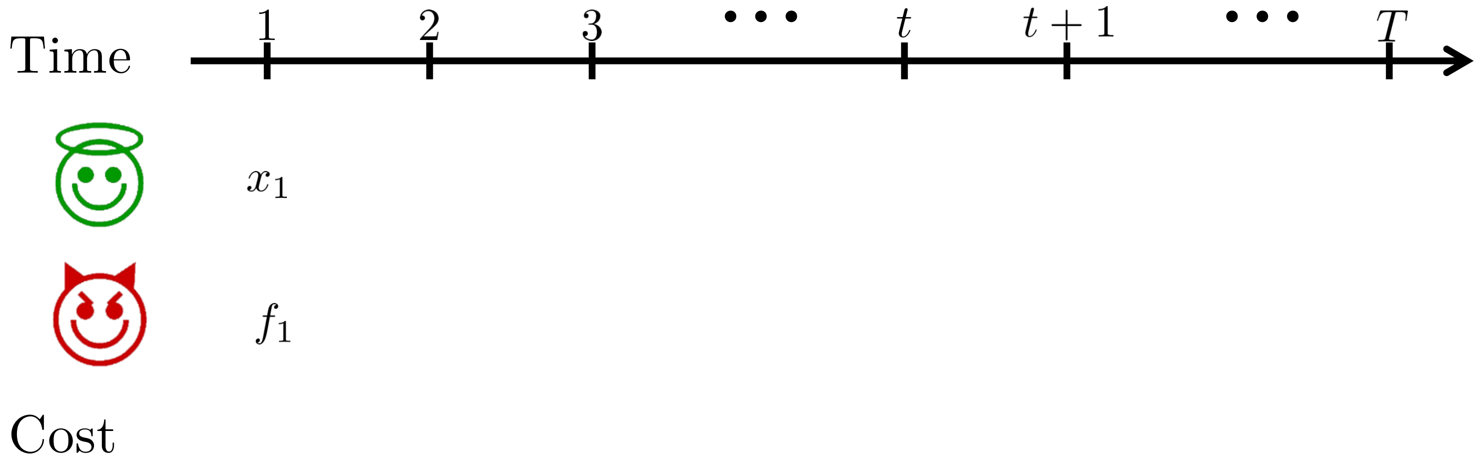

Online Optimization: Setting

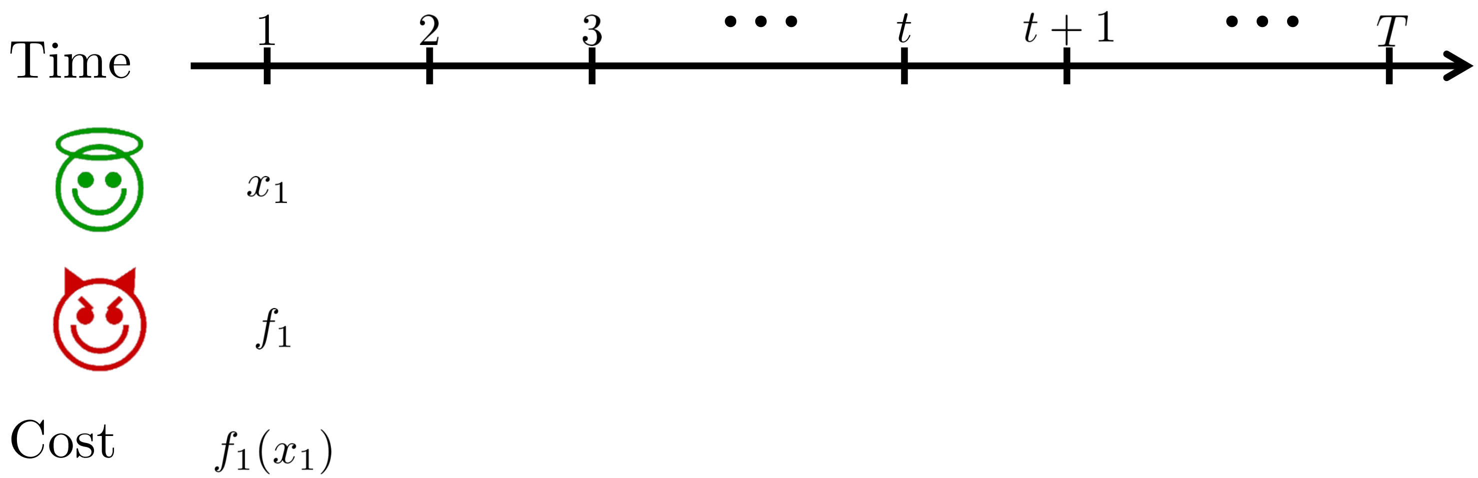

Online Optimization: Setting

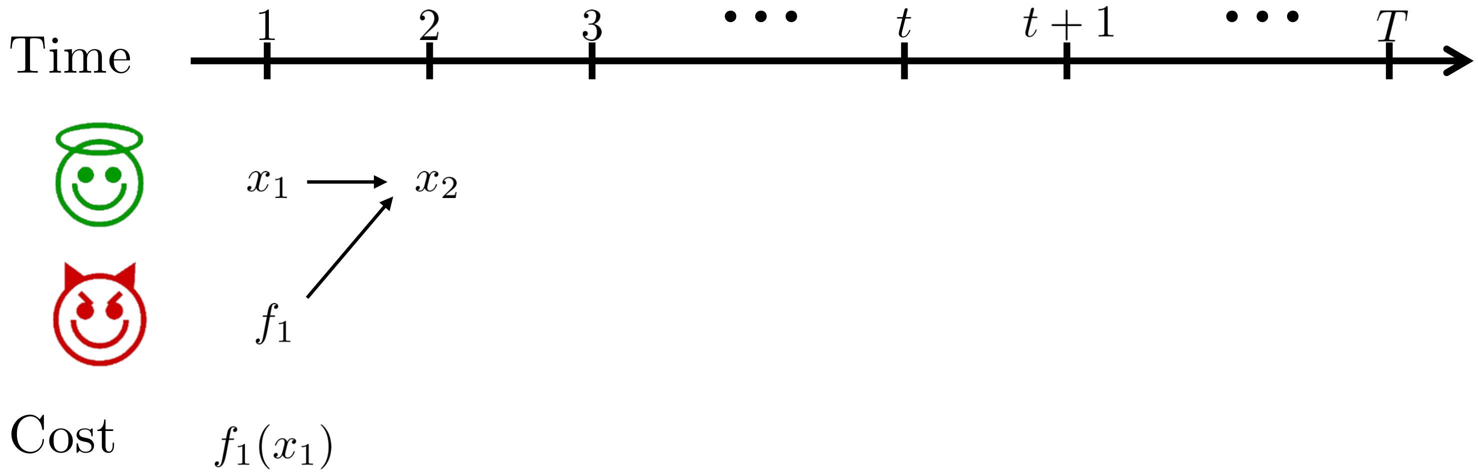

Online Optimization: Setting

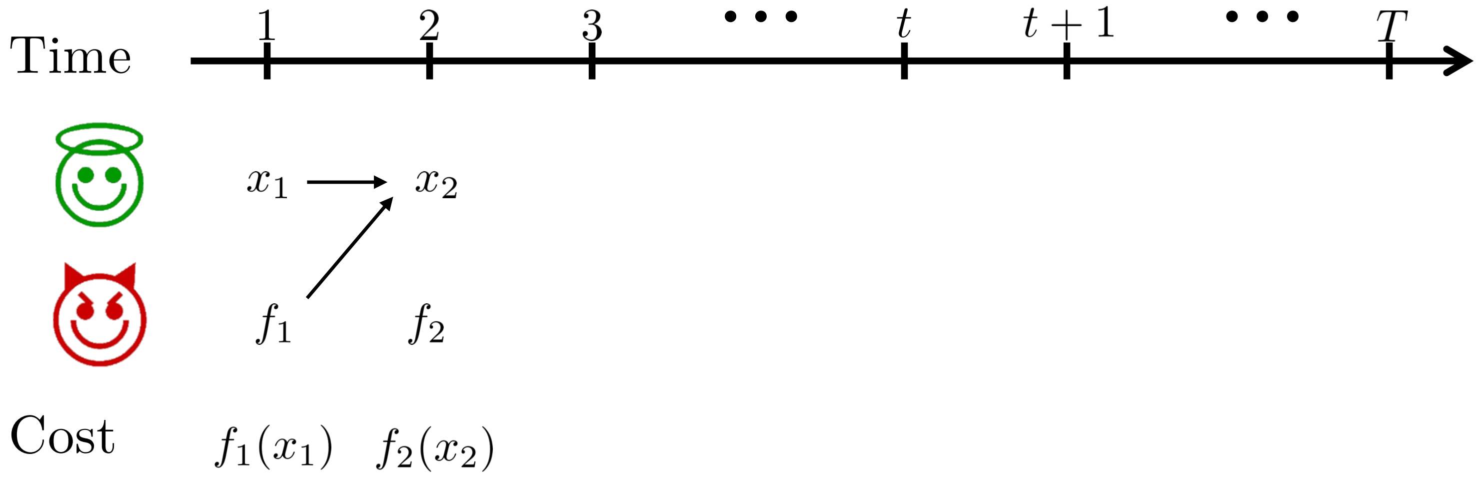

Online Optimization: Setting

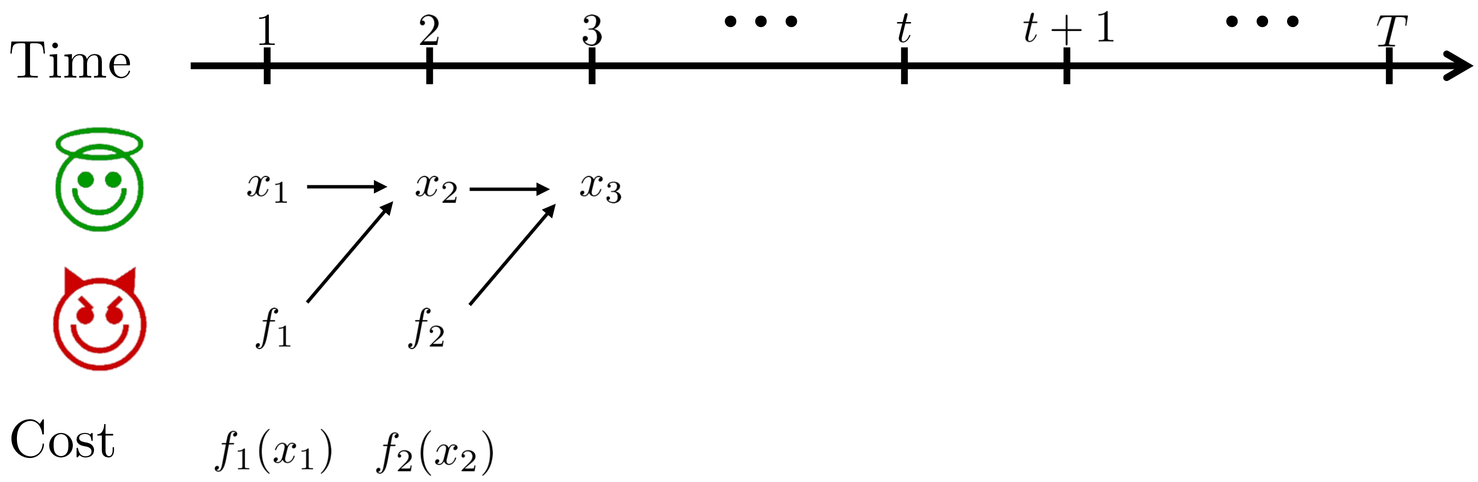

Online Optimization: Setting

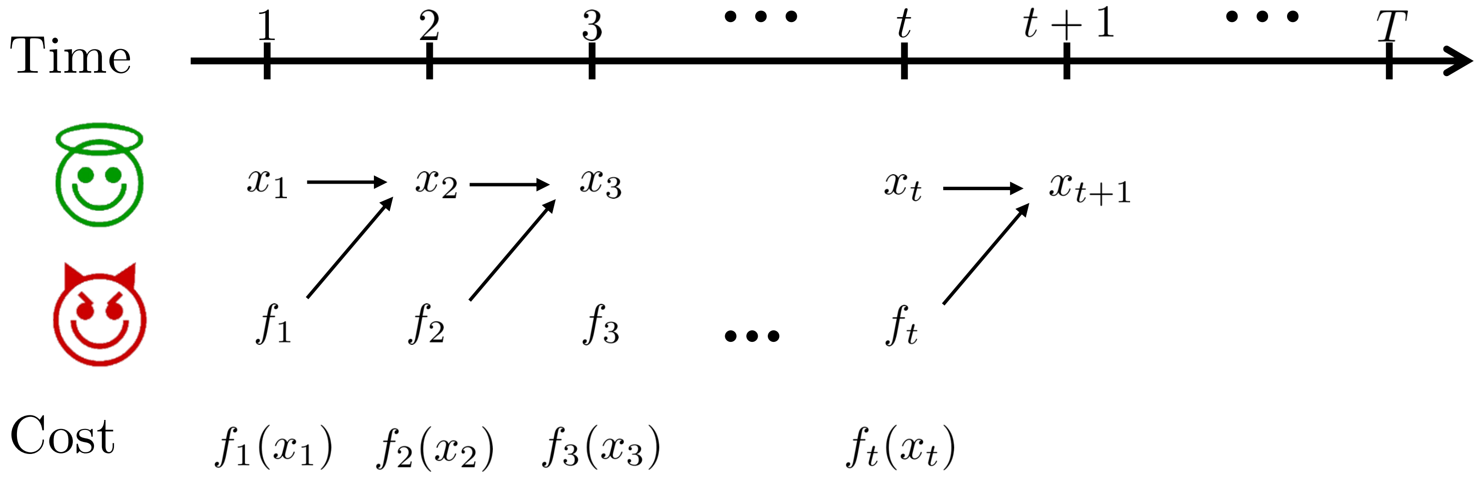

Online Optimization: Setting

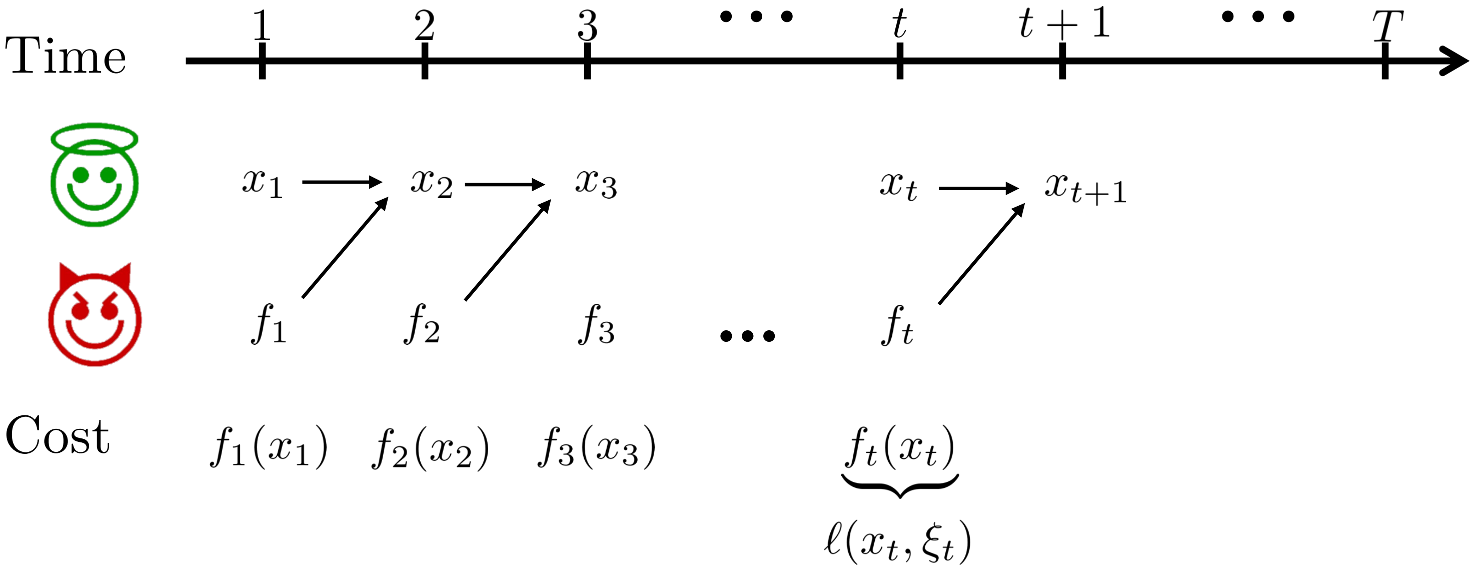

Online Optimization: Setting

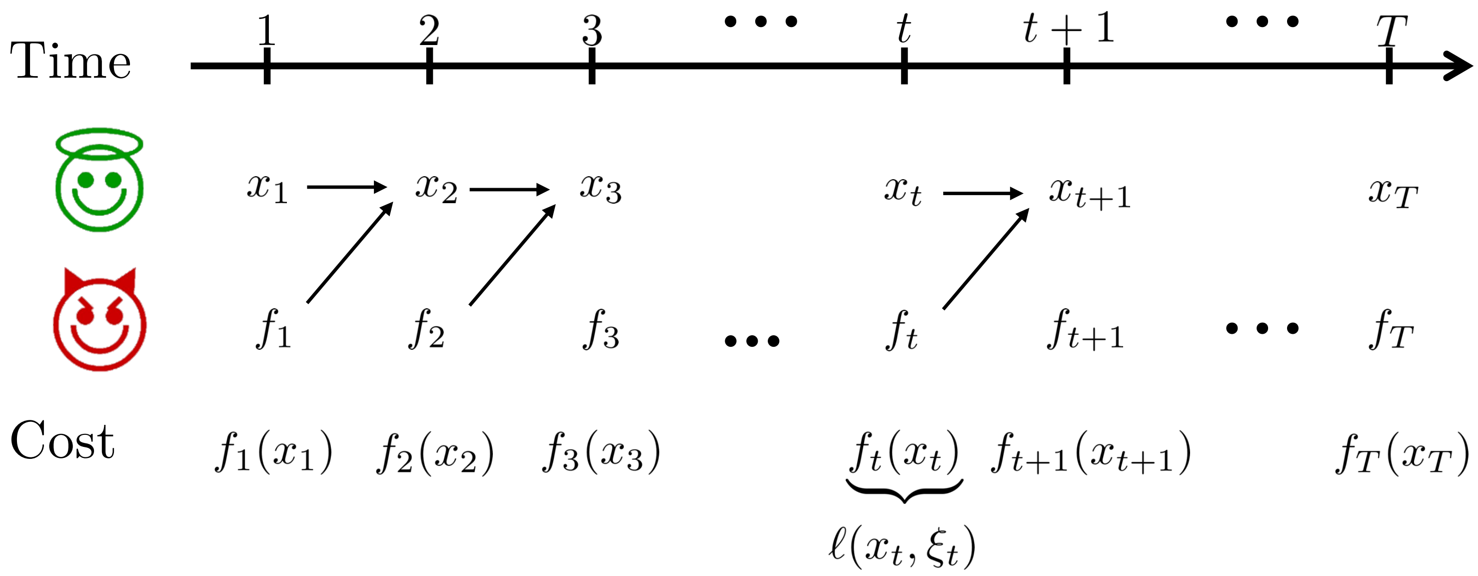

Online Optimization: Setting

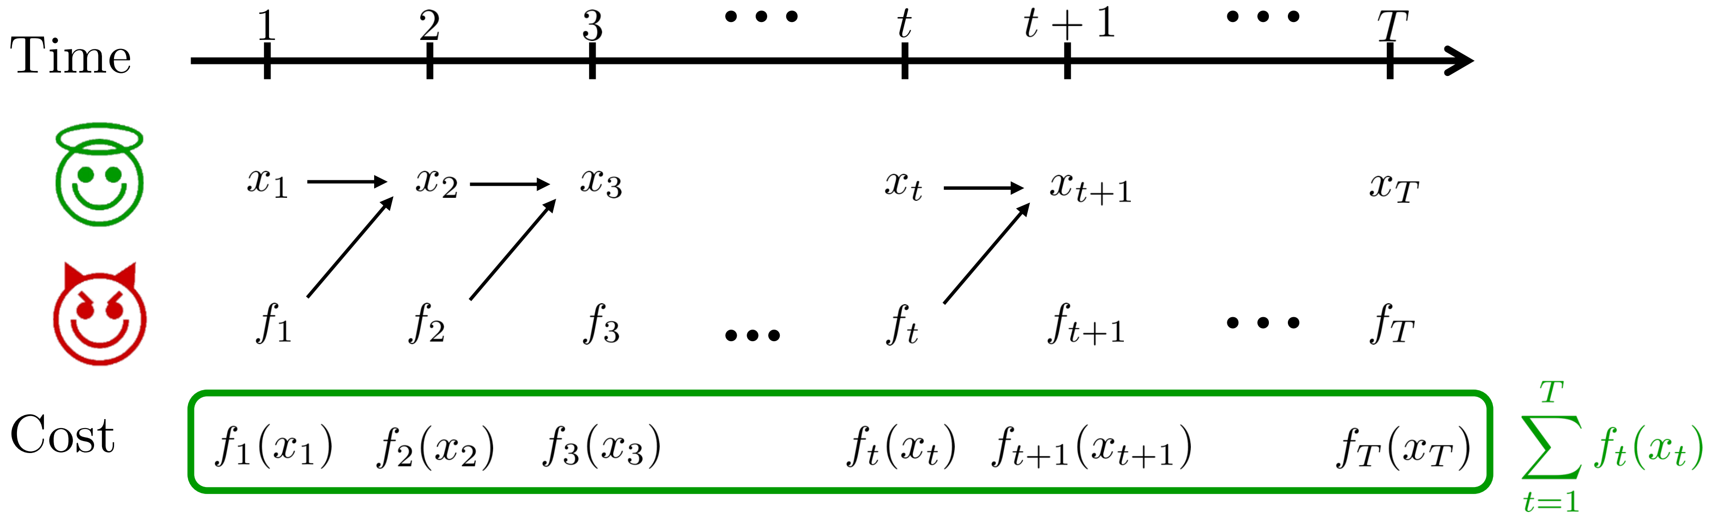

Online Optimization: Setting

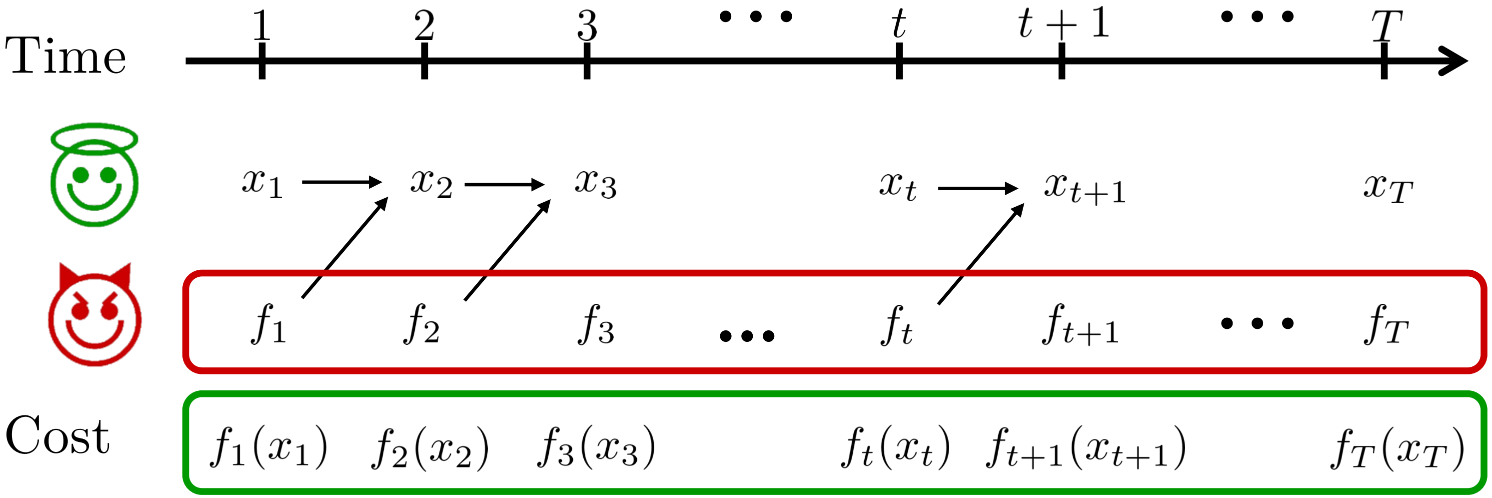

$$\min_{x \in \mathcal{X}} \sum_{t=1}^{T} f_t(x)$$

$$\sum_{t=1}^{T} f_t(x_t)$$

Online Optimization: Setting

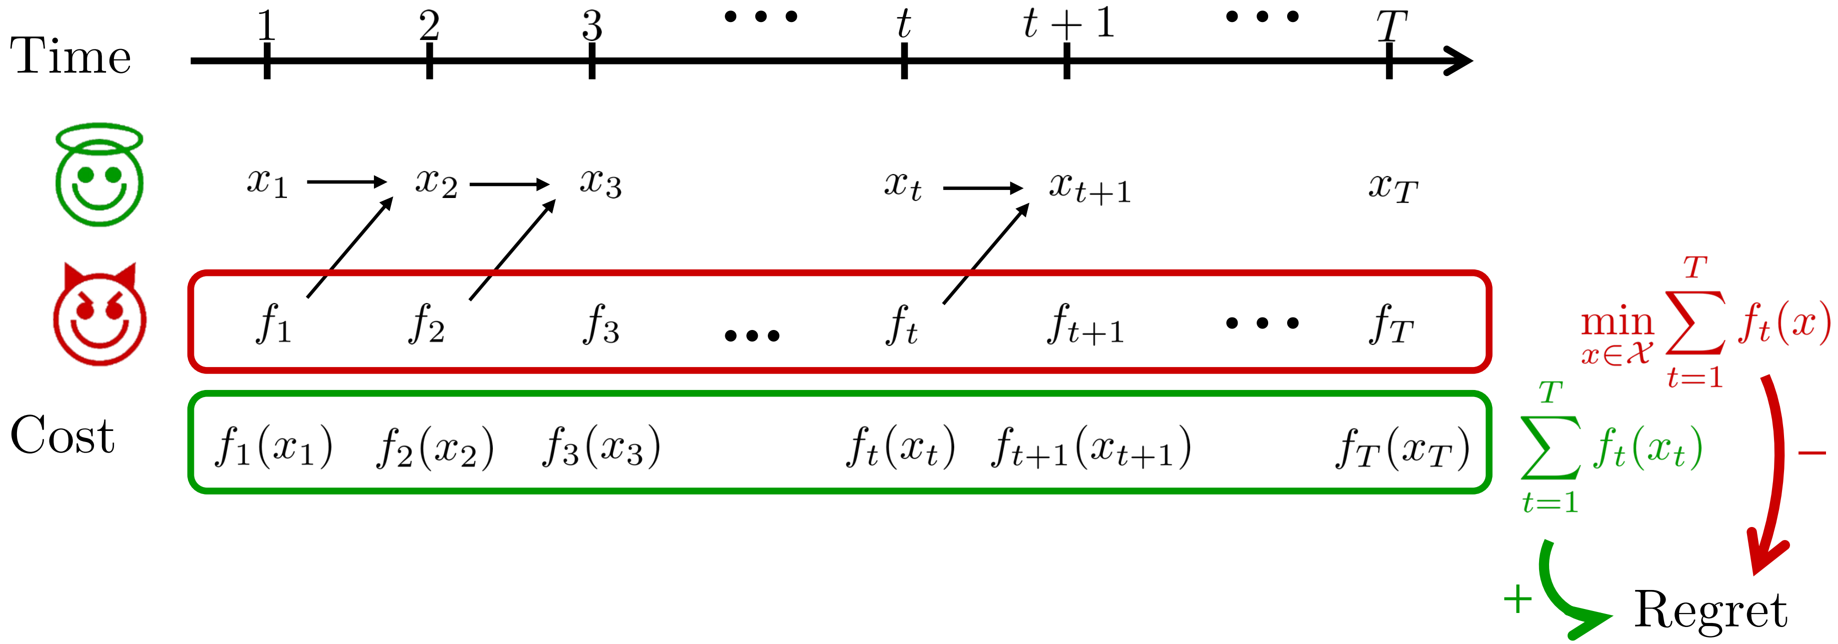

Online Optimization: Setting

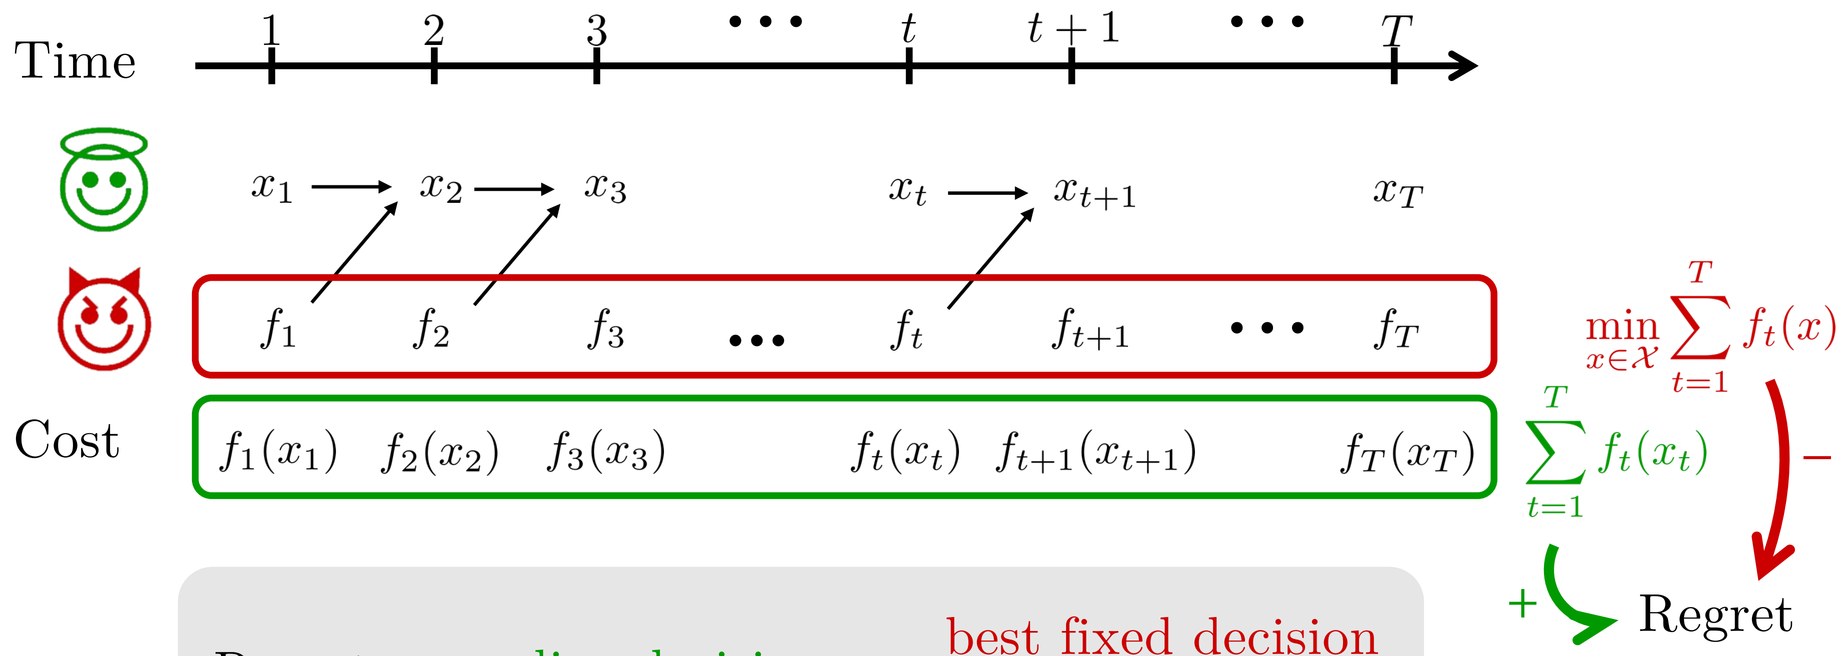

Regret = online decisions − best fixed decision in hindsight

$$\sum_{t=1}^{T} f_t(x_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^{T} f_t(x)$$

Online Optimization for Control

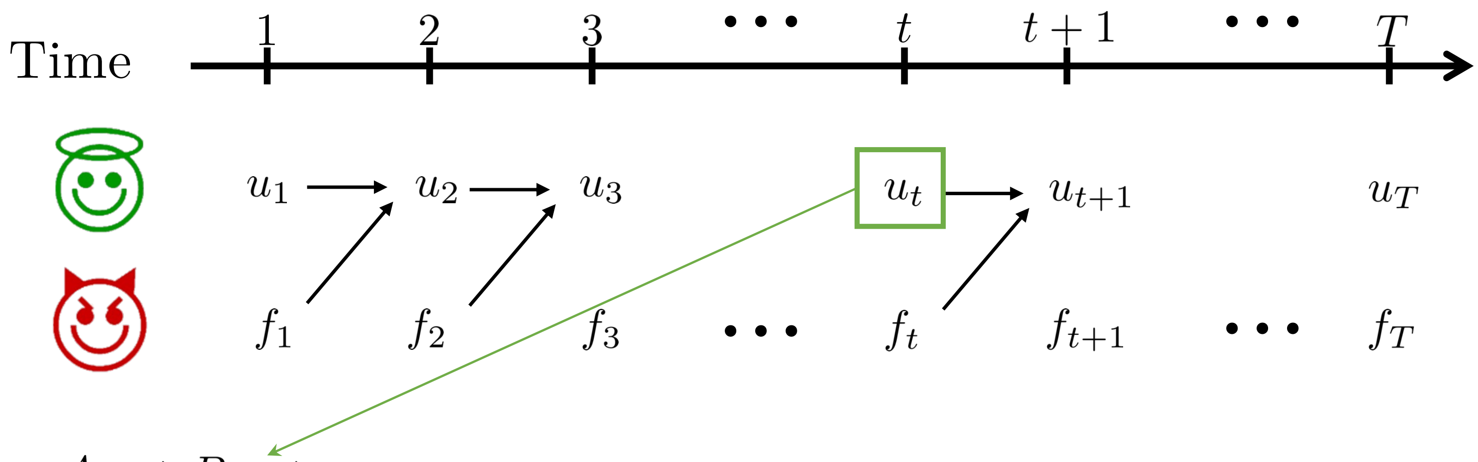

$$x_{t+1} = Ax_t + Bu_t + w_t$$

Online Optimization for Control

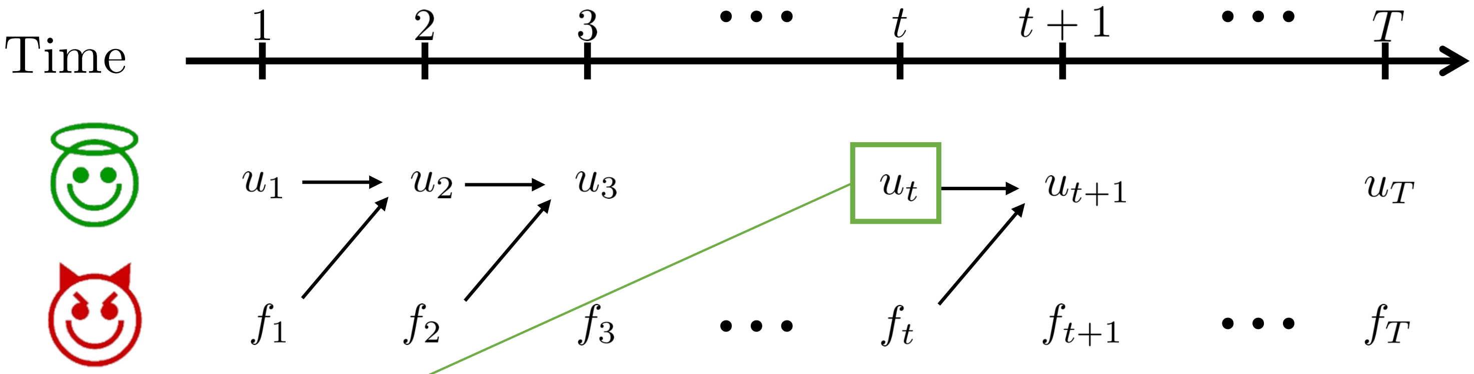

$$x_{t+1} = Ax_t + Bu_t + w_t$$

$$\text{Regret} = \sum_{t=1}^{T} f_t(u_t) - \min_{u \in \mathcal{U}} \sum_{t=1}^{T} f_t(u)$$

Online Optimization for Control

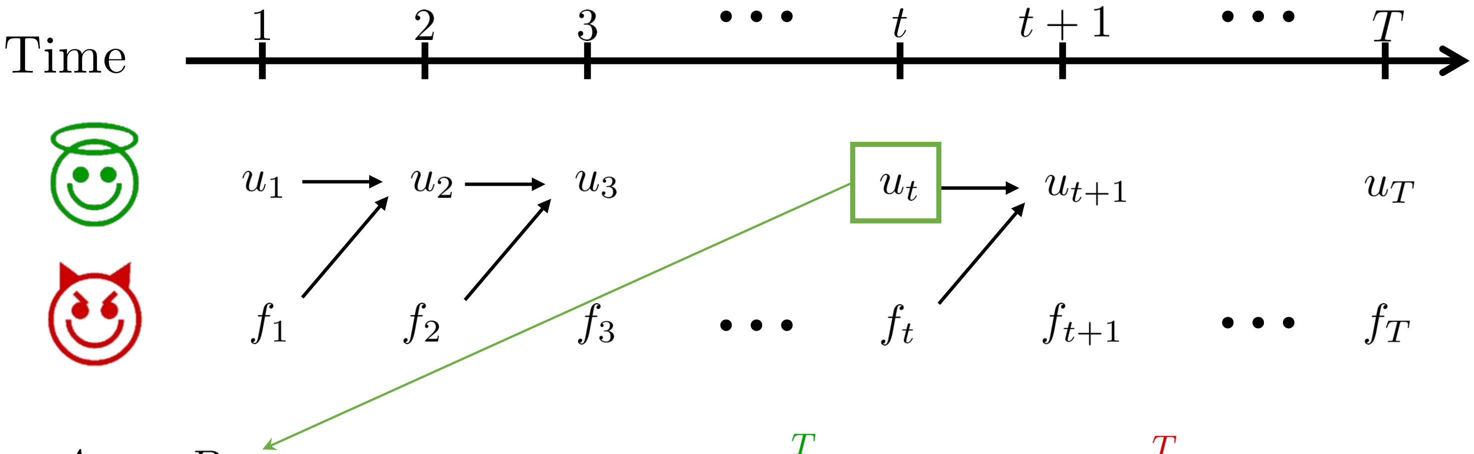

$$x_{t+1} = Ax_t + Bu_t + w_t$$

$$\text{Regret} = \sum_{t=1}^{T} f_t(u_t) - \min_{u \in \mathcal{U}} \sum_{t=1}^{T} f_t(u)$$

open-loop!

Online Optimization for Control

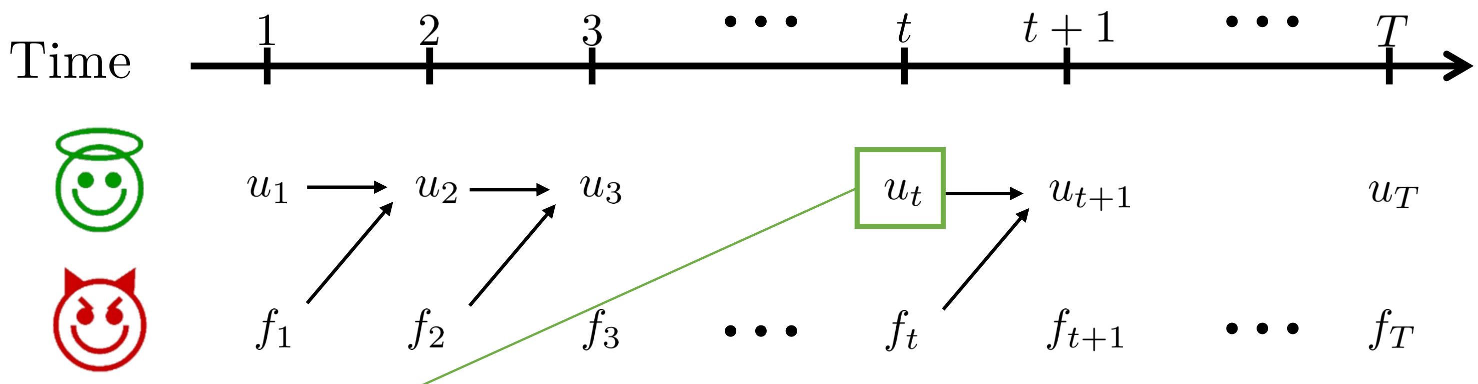

$$x_{t+1} = Ax_t + Bu_t + w_t$$

$$\text{Regret} = \sum_{t=1}^{T} f_t(u_t) - \min_{u \in \mathcal{U}} \sum_{t=1}^{T} f_t(u)$$

same costs,
meaningless regret!

Online Optimization for Control

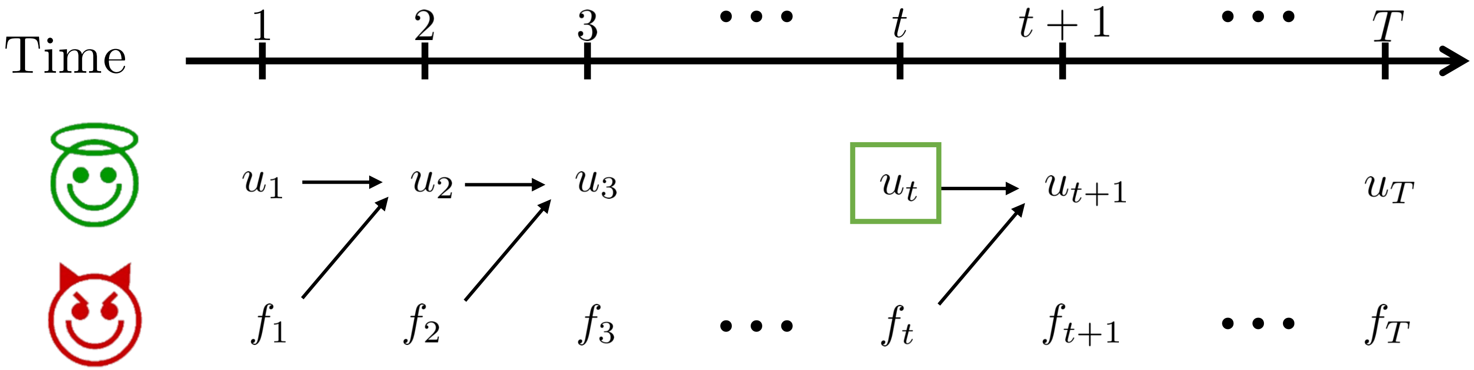

$$x_{t+1} = Ax_t + Bu_t + w_t$$

① $$u_t = \sum_{i=1}^{t-1} M_t^{[i]} w_i$$

disturbance feedback

$$M_t := \left(M_t^{[1]}, \ldots, M_t^{[t-1]}\right)$$

Online Optimization for Control

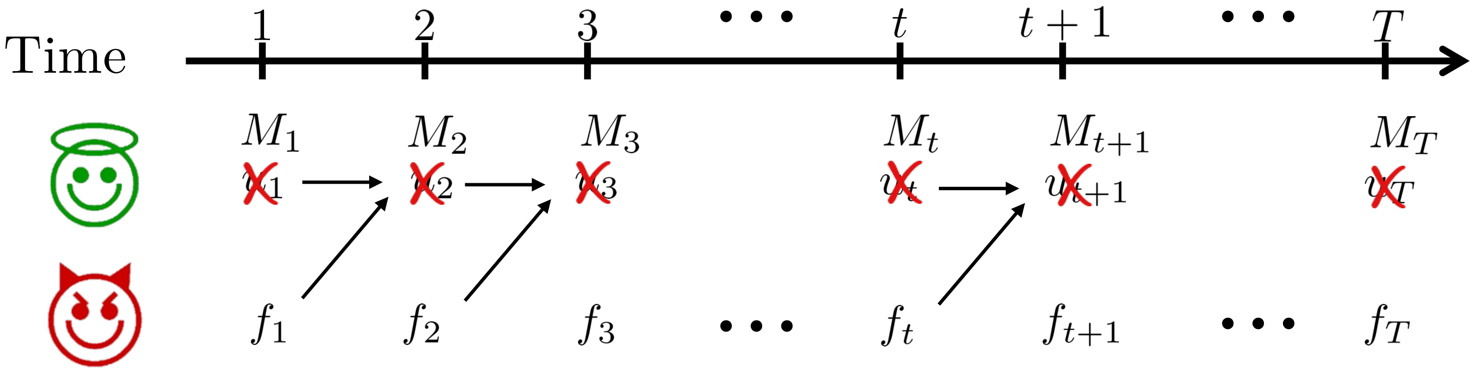

$$x_{t+1} = Ax_t + Bu_t + w_t$$

$$u_t \overset{①}{=} \sum_{i=1}^{t-1} M_t^{[i]} w_i$$

disturbance feedback

$$M_t := \left(M_t^{[1]}, \ldots, M_t^{[t-1]}\right)$$

Online Optimization for Control

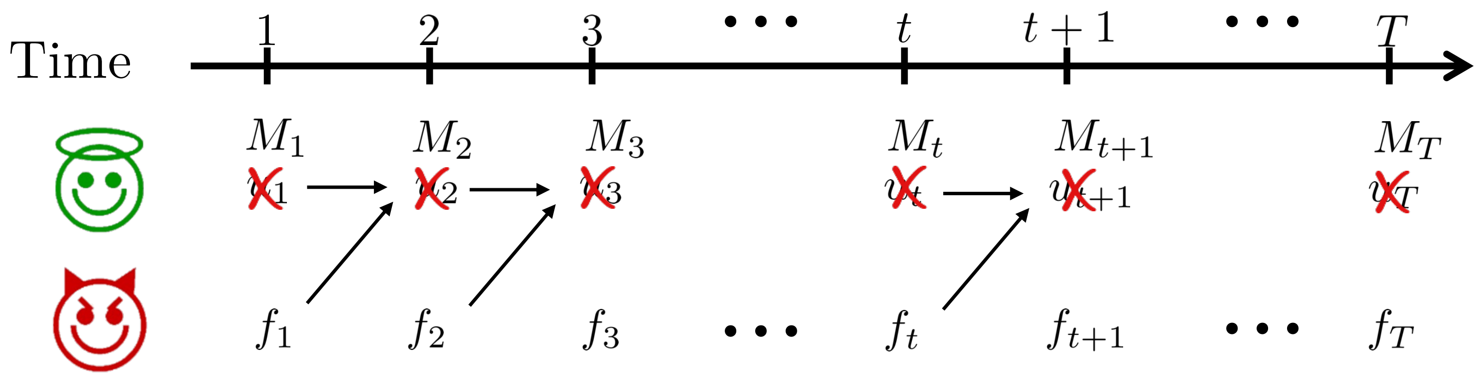

$$x_{t+1} = Ax_t + Bu_t + w_t$$

(1) (2)

$$u_t = \sum_{i=t-h}^{t-1} M_t^{[i]} w_i$$

disturbance feedback

(2)

$$M_t := \left(M_t^{[t-h]}, \dots, M_t^{[t-1]}\right)$$

finite memory

Online Optimization for Control

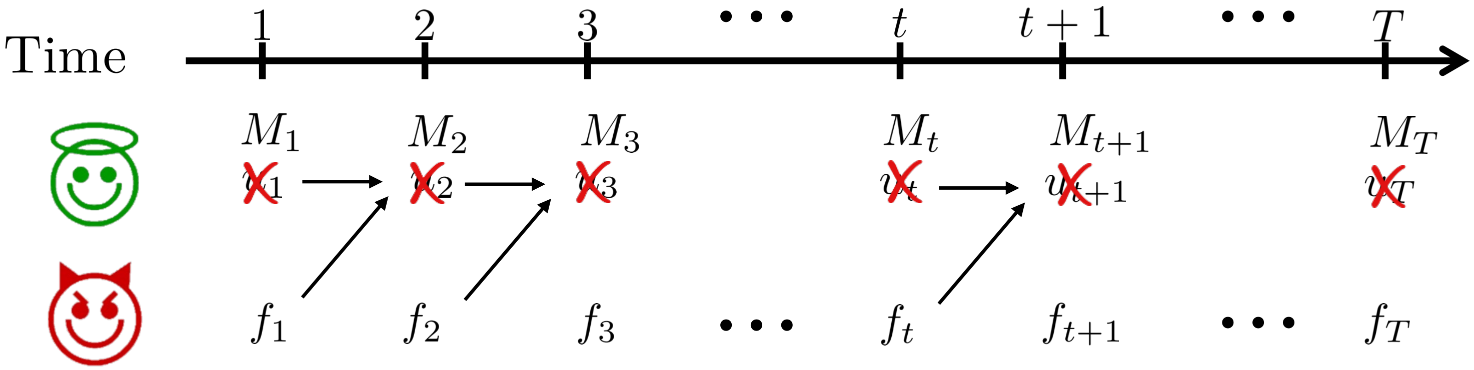

$$x_{t+1} = Ax_t + Bu_t + w_t$$

$$\textcircled{1}\ u_t = \sum_{i=t-h}^{t-1} M_t^{[i]} w_i \quad \textcircled{2}$$

disturbance feedback

$$M_t := \left(M_t^{[t-h]}, \dots, M_t^{[t-1]} \right) \quad \textcircled{2}$$

finite memory

$$\text{Regret} = \sum_{t=1}^{T} c(M_{t-h}, \dots, M_t \mid w_{1:t}) - \min_{M} \sum_{t=1}^{T} c(M, \dots, M \mid w_{1:t})$$

③ with memory ③

Online Optimization: Algorithms

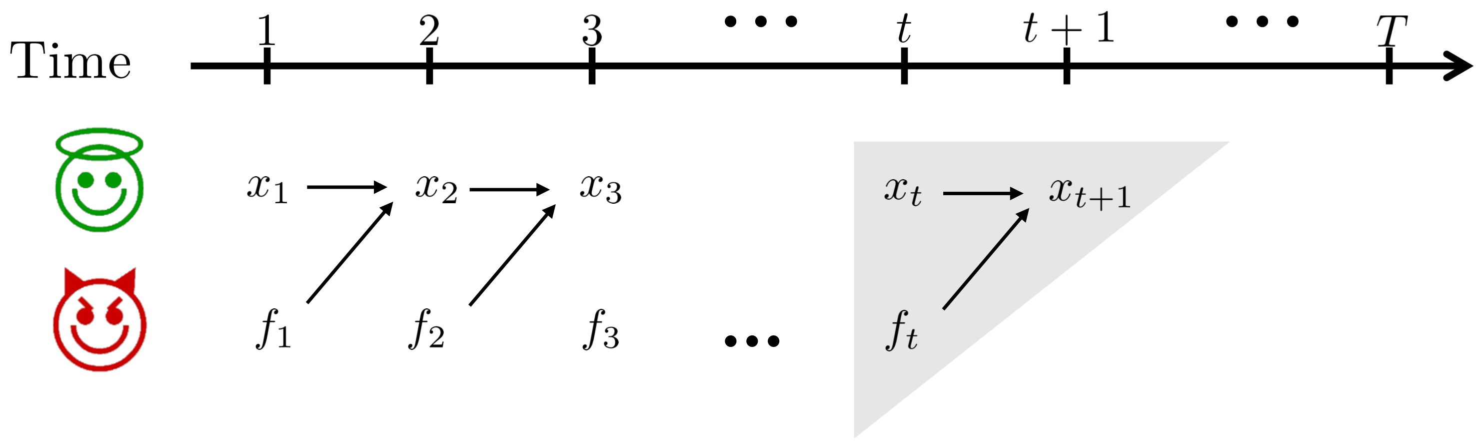

Standard Optimization Algorithms

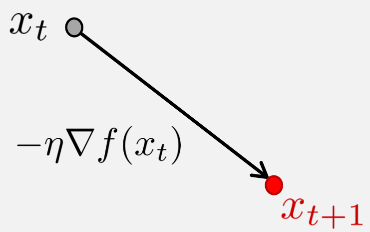

<u>Explicit</u> Gradient Descent

$$x_{t+1} = x_t - \eta \nabla f(x_t)$$

Standard Optimization Algorithms

Converges if
$\eta \leq \frac{1}{\beta}$

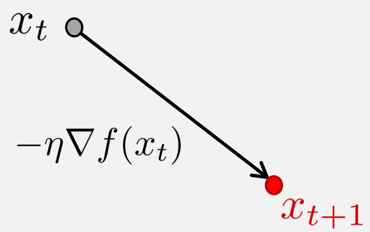

<u>Explicit</u> Gradient Descent

$$x_{t+1} = x_t - \eta \nabla f(x_t)$$

Standard Optimization Algorithms

Converges if
$\eta \leq \frac{1}{\beta}$

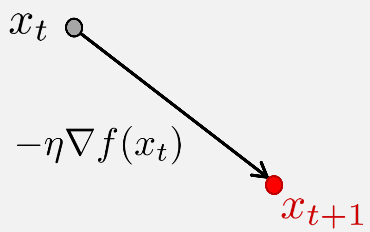

Explicit Gradient Descent

$$x_{t+1} = x_t - \eta \nabla f(x_t)$$

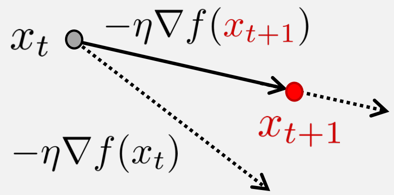

Implicit Gradient Descent

$$x_{t+1} = x_t - \eta \nabla f(x_{t+1})$$

Standard Optimization Algorithms

Converges if
$\eta \leq \frac{1}{\beta}$

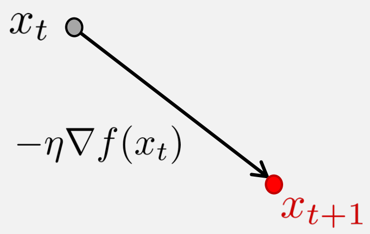

Explicit Gradient Descent

$$x_{t+1} = x_t - \eta\nabla f(x_t)$$

Converges
$\forall \eta > 0$

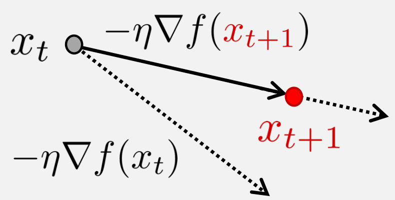

Implicit Gradient Descent

$$x_{t+1} = x_t - \eta\nabla f(x_{t+1})$$

Standard Optimization Algorithms

Converges if $\eta \leq \frac{1}{\beta}$

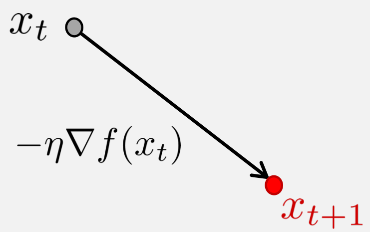

Explicit Gradient Descent

$$x_{t+1} = x_t - \eta\nabla f(x_t)$$

Converges $\forall\eta > 0$

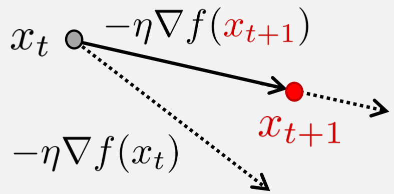

Implicit Gradient Descent

$$x_{t+1} = x_t - \eta\nabla f(x_{t+1})$$

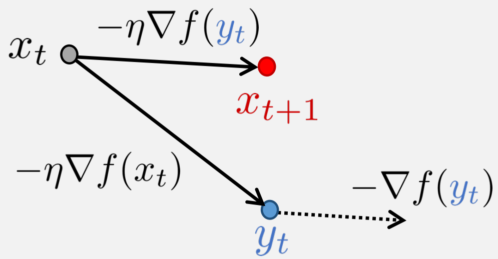

Extra Gradient Descent

$$y_t = x_t - \eta\nabla f(x_t)$$

$$x_{t+1} = x_t - \eta\nabla f(y_t)$$

Online Optimization: Algorithms

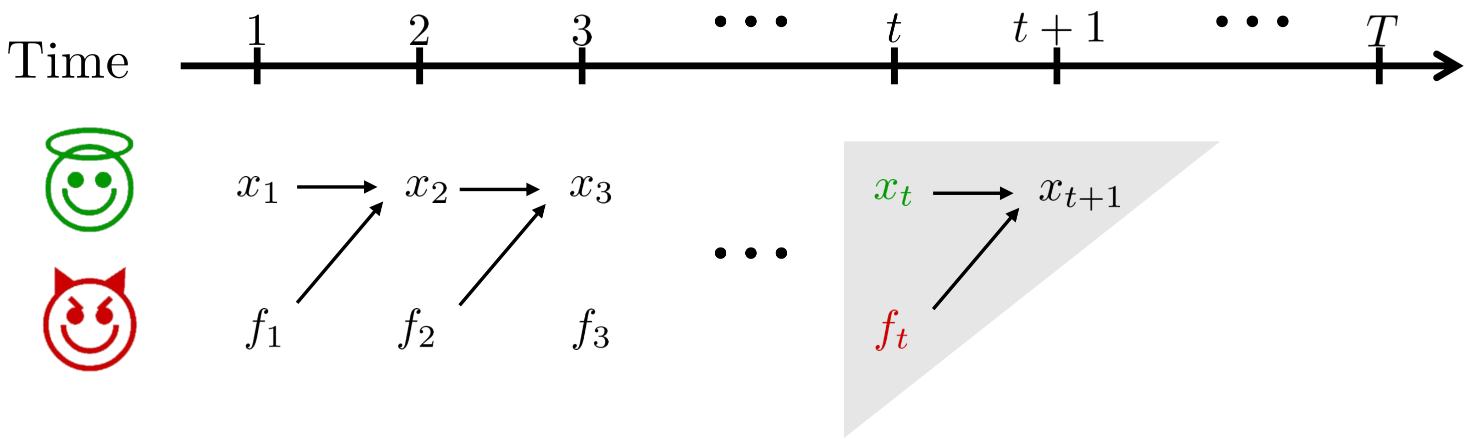

Online Optimization: Algorithms

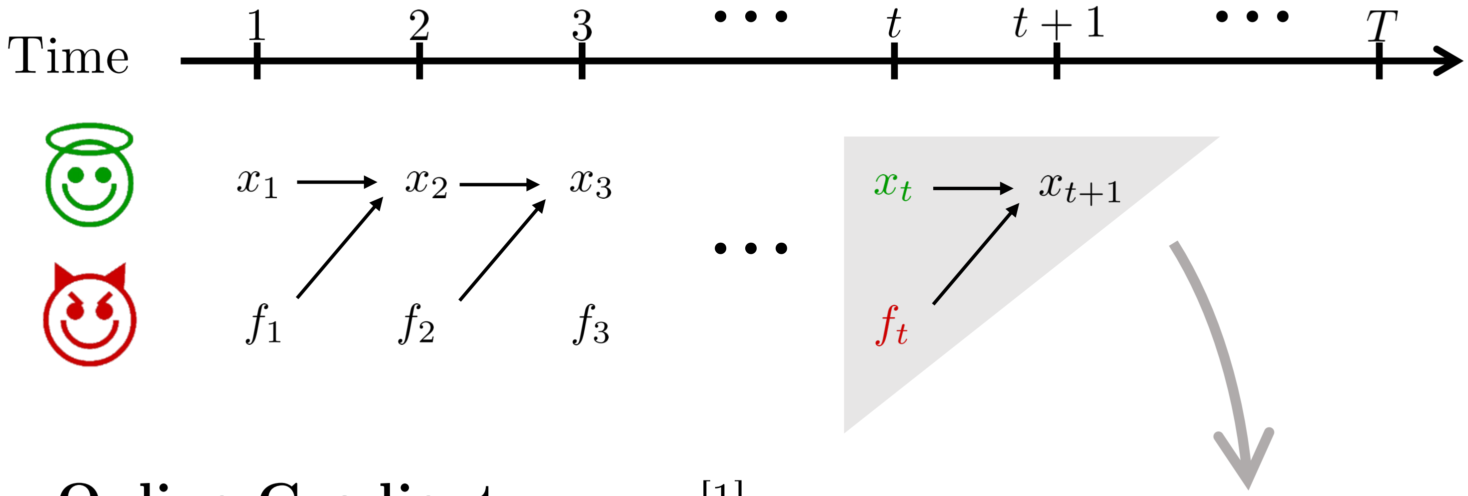

Online Gradient Descent [1]

$$x_{t+1} = \Pi_{\mathcal{X}}\big(x_t - \eta_t \nabla f_t(x_t)\big)$$

Euclidean projection onto $\mathcal{X}$

[1] Zinkevich (2003)

Online Optimization: Algorithms

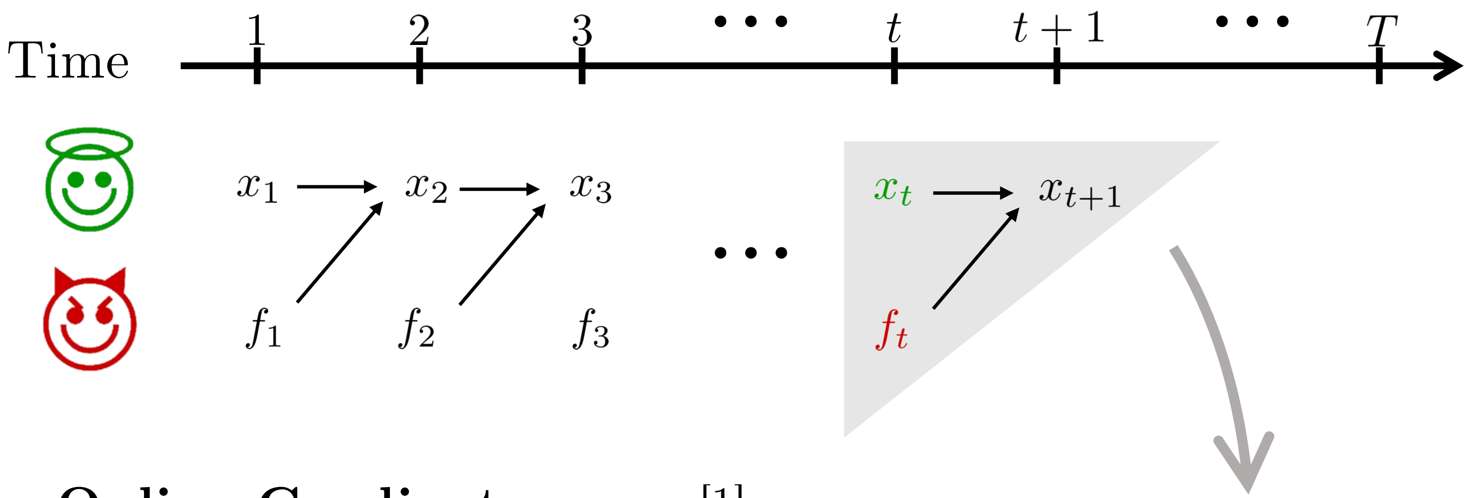

Online Gradient Descent [1]

$$x_{t+1} = \boxed{\Pi_{\mathcal{X}}}\big(x_t - \eta_t \nabla f_t(x_t)\big)$$

f_t convex

$\eta_t \propto 1/\sqrt{t}$

$\longrightarrow$ Regret $= O(\sqrt{T})$

Euclidean projection onto $\mathcal{X}$

[1] Zinkevich (2003)

Online Optimization with Prediction

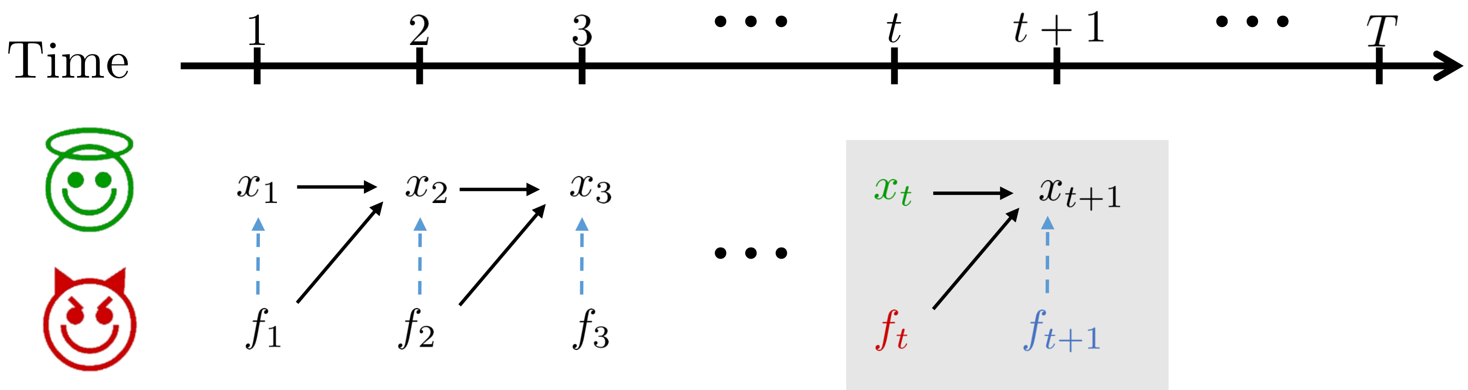

Online Optimization with Prediction

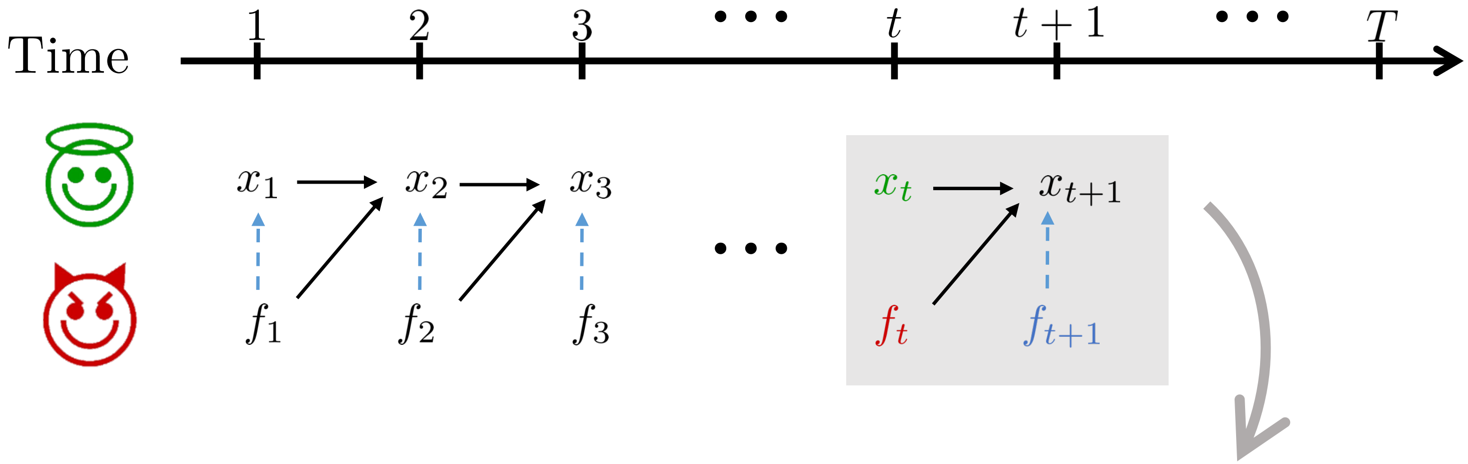

[1]
$$y_t = \Pi_{\mathcal{X}}\big(y_{t-1} - \eta_t \nabla f_t(x_t)\big)$$
$$x_{t+1} = \Pi_{\mathcal{X}}\big(y_t - \eta_{t+1} M_{t+1}\big)$$

[1] Rakhlin and Sridharan (2013)

Online Optimization with Prediction

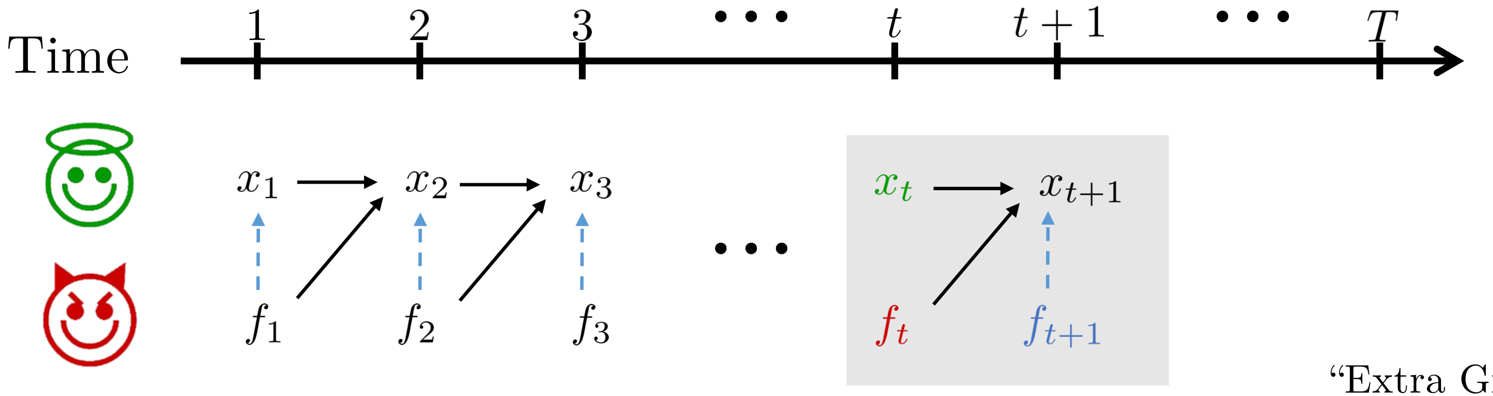

“Extra Gradient Descent”

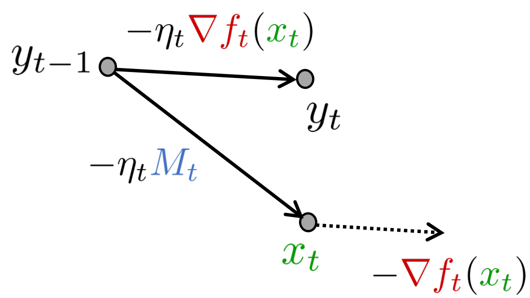

[1]
$$y_t = \Pi_{\mathcal{X}}\left(y_{t-1} - \eta_t \nabla f_t(x_t)\right)$$
$$x_{t+1} = \Pi_{\mathcal{X}}\left(y_t - \eta_{t+1} M_{t+1}\right)$$

[1] Rakhlin and Sridharan (2013)

Online Optimization with Prediction

[1]

$$y_t = \Pi_{\mathcal{X}}\big(y_{t-1} - \eta_t \nabla f_t(x_t)\big)$$
$$x_{t+1} = \Pi_{\mathcal{X}}\big(y_t - \eta_{t+1} M_{t+1}\big)$$

[1] Rakhlin and Sridharan (2013)

Online Optimization with Prediction

[1]
$$y_t = \Pi_{\mathcal{X}}\left(y_{t-1} - \eta_t \nabla f_t(x_t)\right)$$
$$x_{t+1} = \Pi_{\mathcal{X}}\left(y_t - \eta_{t+1} M_{t+1}\right)$$

- $\eta_t = O\left(1/\sqrt{\sum_{i=1}^{t-1} \|\nabla f_i(x_i) - M_i\|^2}\right)$

- $\text{Regret} = O\left(1 + \sqrt{\sum_{t=1}^{T} \|\nabla f_t(x_t) - M_t\|^2}\right)$

[1] Rakhlin and Sridharan (2013)

Online Optimization with Prediction

[1]

$$y_t = \Pi_{\mathcal{X}}\left(y_{t-1} - \eta_t \nabla f_t(x_t)\right)$$
$$x_{t+1} = \Pi_{\mathcal{X}}\left(y_t - \eta_{t+1} M_{t+1}\right)$$

- $\eta_t = O\left(1/\sqrt{\sum_{i=1}^{t-1} \|\nabla f_i(x_i) - M_i\|^2}\right)$

- $\text{Regret} = O\left(1 + \sqrt{\sum_{t=1}^{T} \|\nabla f_t(x_t) - M_t\|^2}\right)$

Constant regret
$\Rightarrow M_t = \nabla f_t(x_t)$

[1] Rakhlin and Sridharan (2013)

Online Optimization with Prediction

[1]

$$y_t = \Pi_{\mathcal{X}}\left(y_{t-1} - \eta_t \nabla f_t(x_t)\right)$$

$$x_{t+1} = \Pi_{\mathcal{X}}\left(y_t - \eta_{t+1} M_{t+1}\right)$$

- $\eta_t = O\left(1/\sqrt{\sum_{i=1}^{t-1} \|\nabla f_i(x_i) - M_i\|^2}\right)$
- Regret $= O\left(1 + \sqrt{\sum_{t=1}^{T} \|\nabla f_t(x_t) - M_t\|^2}\right)$

Constant regret $\Rightarrow M_t = \nabla f_t(x_t)$

Implicit Gradient Descent

$$x_{t+1} = \Pi_{\mathcal{X}}\left(y_t - \eta_{t+1} \nabla f_{t+1}(x_{t+1})\right)$$

[1] Rakhlin and Sridharan (2013)

Online Optimization with Prediction

[1]

$$y_t = \Pi_{\mathcal{X}}\left(y_{t-1} - \eta_t \nabla f_t(x_t)\right)$$

$$x_{t+1} = \Pi_{\mathcal{X}}\left(y_t - \eta_{t+1} M_{t+1}\right)$$

- $\eta_t = O\left(1/\sqrt{\sum_{i=1}^{t-1} \|\nabla f_i(x_i) - M_i\|^2 + 4\beta^2}\right)$

  $\nabla f_i(y_{i-1})$

- $\text{Regret} = O\left(1 + \sqrt{\sum_{t=1}^{T} \|\nabla f(x_t) - M_t\|^2}\right)$

  $\nabla f_t(y_{t-1})$

Constant regret

$\Rightarrow M_t = \nabla f(x_t)$

$\nabla f_t(y_{t-1})$

Explicit Gradient Descent

$$x_{t+1} = \Pi_{\mathcal{X}}\left(y_t - \eta_{t+1} \nabla f_{t+1}(x_{t+1})\right)$$

$\nabla f_{t+1}(y_t)$

[1] Rakhlin and Sridharan (2013)

[2] Zattoni Scroccaro, Kolarijani, and PME (2023)

Online Optimization: Convex Costs

[3]
$$y_t = \Pi_{\mathcal{X}}\big(y_{t-1} - \eta_t \nabla f_t(x_t)\big)$$
$$x_{t+1} = \Pi_{\mathcal{X}}\big(y_t - \eta_{t+1} \widetilde{\nabla f}_{t+1}(y_t)\big)$$

Gradient Prediction Error
$$D_t := \sum_{i=1}^{t} \|\nabla f_i(y_{i-1}) - \widetilde{\nabla f}_i(y_{i-1})\|^2$$

[1] Zinkevich (2003)
[2] Ho-Nguyen and Kılınç-Karzan (2019)
[3] Zattoni Scroccaro, Kolarijani, and PME (2023)

Online Optimization: Convex Costs

[3]
$$y_t = \Pi_{\mathcal{X}}\left(y_{t-1} - \eta_t \nabla f_t(x_t)\right)$$
$$x_{t+1} = \Pi_{\mathcal{X}}\left(y_t - \eta_{t+1} \widetilde{\nabla f}_{t+1}(y_t)\right)$$

Gradient Prediction Error
$$D_t := \sum_{i=1}^{t} \|\nabla f_i(y_{i-1}) - \widetilde{\nabla f}_i(y_{i-1})\|^2$$

	Worst-case [1]	**General prediction [3]**	**Perfect prediction [2]**
Prediction		$\widetilde{\nabla f}_{t+1}$	
η_t		$1/\sqrt{D_{t-1} + 4\beta^2}$	
Regret		$O(1 + \sqrt{D_T})$	

[1] Zinkevich (2003)
[2] Ho-Nguyen and Kılınç-Karzan (2019)
[3] Zattoni Scroccaro, Kolarijani, and PME (2023)

Online Optimization: Convex Costs

[3]
$$y_t = \Pi_{\mathcal{X}}\left(y_{t-1} - \eta_t \nabla f_t(x_t)\right)$$
$$x_{t+1} = \Pi_{\mathcal{X}}\left(y_t - \eta_{t+1} \widetilde{\nabla f}_{t+1}(y_t)\right)$$

Gradient Prediction Error

$$D_t := \sum_{i=1}^{t} \|\nabla f_i(y_{i-1}) - \widetilde{\nabla f}_i(y_{i-1})\|^2$$

	Worst-case [1]	**General prediction [3]**	**Perfect prediction [2]**
Prediction	$\widetilde{\nabla f}_{t+1} = 0$	$\widetilde{\nabla f}_{t+1}$	
η_t	$O(1/\sqrt{t})$	$1/\sqrt{D_{t-1} + 4\beta^2}$	
Regret	$O(\sqrt{T})$	$O(1 + \sqrt{D_T})$	

[1] Zinkevich (2003)
[2] Ho-Nguyen and Kılınç-Karzan (2019)
[3] Zattoni Scroccaro, Kolarijani, and PME (2023)

Online Optimization: Convex Costs

[3]
$$y_t = \Pi_{\mathcal{X}}\big(y_{t-1} - \eta_t \nabla f_t(x_t)\big)$$
$$x_{t+1} = \Pi_{\mathcal{X}}\big(y_t - \eta_{t+1} \widetilde{\nabla f}_{t+1}(y_t)\big)$$

Gradient Prediction Error
$$D_t := \sum_{i=1}^{t} \|\nabla f_i(y_{i-1}) - \widetilde{\nabla f}_i(y_{i-1})\|^2$$

	Worst-case [1]	**General prediction [3]**	**Perfect prediction [2]**
Prediction	$\widetilde{\nabla f}_{t+1} = 0$	$\widetilde{\nabla f}_{t+1}$	$\widetilde{\nabla f}_{t+1} = \nabla f_{t+1}$
η_t	$O(1/\sqrt{t})$	$1/\sqrt{D_{t-1} + 4\beta^2}$	$1/2\beta$
Regret	$O(\sqrt{T})$	$O(1 + \sqrt{D_T})$	$O(1)$

[1] Zinkevich (2003)
[2] Ho-Nguyen and Kılınç-Karzan (2019)
[3] Zattoni Scroccaro, Kolarijani, and PME (2023)

Online Optimization: Strongly Convex Costs

[3]
$$y_t = \Pi_{\mathcal{X}}\left(y_{t-1} - \eta_t \nabla f_t(x_t)\right)$$
$$x_{t+1} = \Pi_{\mathcal{X}}\left(y_t - \eta_{t+1} \widetilde{\nabla f}_{t+1}(y_t)\right)$$

Gradient Prediction Error $$D_t := \sum_{i=1}^{t} \|\nabla f_i(y_{i-1}) - \widetilde{\nabla f}_i(y_{i-1})\|^2$$

	Worst-case [1]	**General prediction [3]**	**Perfect prediction [2]**
Prediction	$\widetilde{\nabla f}_{t+1} = 0$	$\widetilde{\nabla f}_{t+1}$	$\widetilde{\nabla f}_{t+1} = \nabla f_{t+1}$
η_t	$O(1/t)$	$O(1/(D_{t-1} + 2\beta))$	$1/2\beta$
Regret	$O(\log(T))$	$O(1 + \log(1 + D_T))$	$O(1)$

Strongly convex costs

[1] Zinkevich (2003)
[2] Ho-Nguyen and Kılınç-Karzan (2019)
[3] Zattoni Scroccaro, Kolarijani, and PME (2023)

Numerical Example: Portfolio Selection

- Player is an investor. Nature is the stock market.
- The Player chooses $x_t \in \Delta_n$, a distributions of his/her wealth over n assets.
- Nature chooses the returns vector $r_t \in \mathbb{R}_+^n$
- At time t, the Player suffers the loss $f_t = -\log(\langle r_t, x_t \rangle)$
- $\widetilde{\nabla f}_t(y_{t-1}) = -\tilde{r}_t / \langle \tilde{r}_t, y_{t-1} \rangle$, where $\tilde{r}_t$ is the prediction of r_t

Numerical Example: Portfolio Selection

PREDICTION MODELS

- **perfect**: $\hat{r}_t = r_t$
- **previous**: $\hat{r}_t = r_{t-1}$
- **noisy**: $\hat{r}_t = r_t + w_t, \; w_t \sim N(0, I)$
- **random**: $\hat{r}_t \sim U(0, 2)$

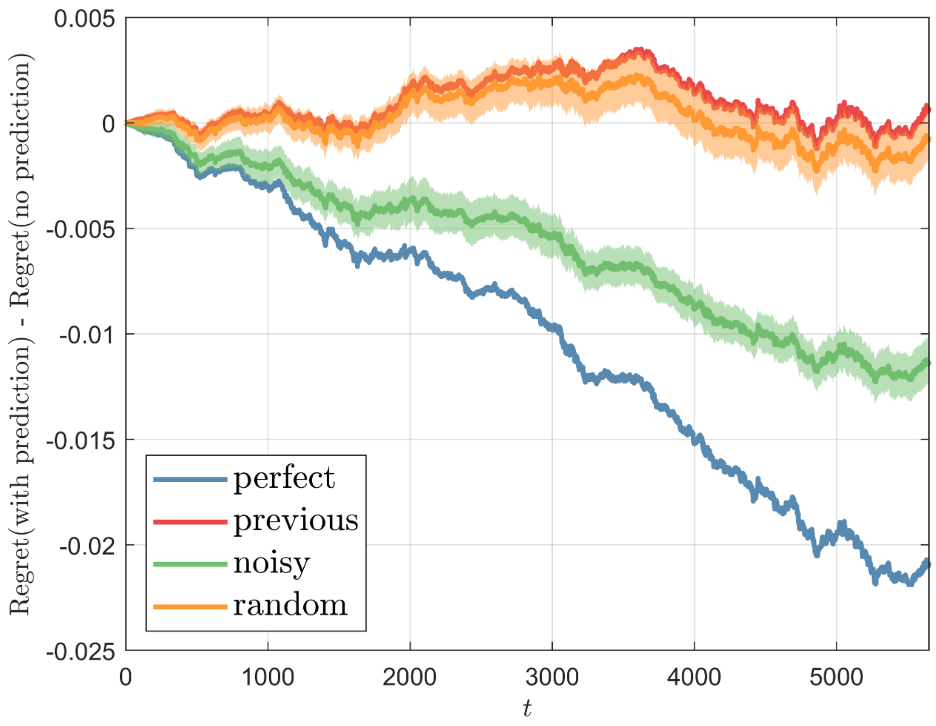

NYSE dataset.
Zattoni Scroccaro, Kolarijani, and PME (2022)

Other Extensions

- Non-smooth cost functions

$$f_t(x) = s_t(x) + \underbrace{r_t(x)}_{\text{Non-smooth}}$$

$\Rightarrow$ Extra Gradient Composite Mirror Descent

- Dynamic regrets

$$\sum_{t=1}^{T} f_t(x_t) - \underbrace{\min_{x \in \mathcal{X}}}_{\text{fixed decision}} \sum_{t=1}^{T} f_t(x) \quad \Longrightarrow \quad \sum_{t=1}^{T} f_t(x_t) - \sum_{t=1}^{T} f_t(\underbrace{u_t}_{\text{reference trajectory}})$$

[1] Zattoni Scroccaro, Kolarijani, and PME (2022)

Other Extensions

- Non-smooth cost functions

$$f_t(x) = s_t(x) + \underbrace{r_t(x)}_{\text{Non-smooth}}$$

$\Rightarrow$ Extra Gradient Composite Mirror Descent

- Dynamic regrets

$$\sum_{t=1}^{T} f_t(x_t) - \underbrace{\min_{x \in \mathcal{X}}}_{\text{fixed decision}} \sum_{t=1}^{T} f_t(x) \quad \Longrightarrow \quad \sum_{t=1}^{T} f_t(x_t) - \sum_{t=1}^{T} f_t\underbrace{(u_t)}_{\text{reference trajectory}}$$

[1] Zattoni Scroccaro, Kolarijani, and PME (2022)

References

- Zattoni Scroccaro, Kolarijani, and Mohajerin Esfahani, "Adaptive Composite Online Optimization: Predictions in Static and Dynamic Environments", *IEEE Transaction on Automatic Control (TAC)*, 2023.
- Mohajerin Esfahani, Sutter, and Lygeros "Performance Bounds for the Scenario Approach and an Extension to a Class of Non-convex Programs", *IEEE Transaction on Automatic Control (TAC)*, 2015
- Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent", *International Conference on Machine Learning (ICML)*, 2003.
- Hazan, Agarwal, and Kale, "Logarithmic regret algorithms for online convex optimization", Machine Learning, 2007.
- Rakhlin and Sridharan, "Optimization, Learning, and Games with Predictable Sequences", *Conference on Neural Information Processing Systems (NeurIPS)*, 2013.
- Ho-Nguyen and Kılınç-Karzan, "Exploiting problem structure in optimization under uncertainty via online convex optimization", *Mathematical Programming (MP)*, 2019.