NJUAI 南京大学 人工智能学院
SCHOOL OF ARTIFICIAL INTELLIGENCE, NANJING UNIVERSITY

LAMDA
Learning And Mining from DatA

On the Relation between Interval Regret and Dynamic Regret

Peng Zhao

LAMDA Group

School of AI, Nanjing University

2026-11-09 @ Università degli Studi di Milano

Ciao Milano!





Outline

- Online Learning
- Performance Metric
- Our Result
- Summary

Outline

- Online Learning
- Performance Metric
- Our Result
- Summary

Interaction in Machine Learning

- Many ML systems need to learn from ***interaction*** with environments…



personalized recommender



agentic system



embodied robotics

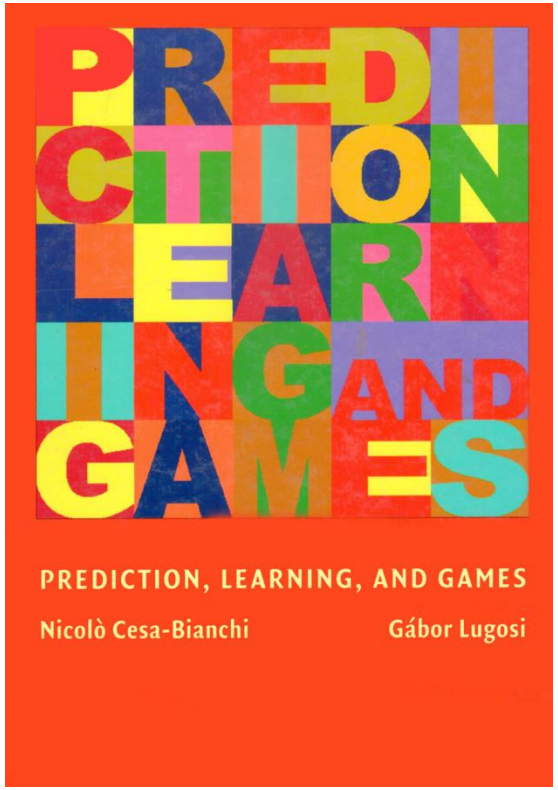
- Interactive ML: ***effective*** and ***efficient*** over **long horizon**
  - online horizon T can be very large (millions/billions of rounds), or even unbounded
  - impossible to store all history (memory cost, privacy, etc) and retrain from scratch

Online Learning

- Online Learning is ubiquitous in real applications when interacting with environments

At each round $t = 1, \ldots, T$:

1. The learner chooses a decision $\mathbf{x}_t \in \mathcal{X} \subseteq \mathbb{R}^d$.
2. The environment reveals a loss function f_t.
3. The learner incurs $f_t(\mathbf{x}_t)$ and uses feedback to update.



connecting areas: optimization, info theory, statistics; also, LLM-RL, agentic RL, attention mechanism, …

Regret

- The best in hindsight

  After T rounds, suppose we look back and imagine that we *did* know all $\{f_t\}_{t=1}^T$.

  Then the best decision in hindsight would be: $\mathbf{x}^\star \in \arg\min_{\mathbf{x}\in\mathcal{X}} \sum_{t=1}^T f_t(\mathbf{x})$.

- To benchmark with the offline model (optimal in hindsight)

Regret: $$\mathrm{REG}_T = \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \min_{\mathbf{x}\in\mathcal{X}} \sum_{t=1}^{T} f_t(\mathbf{x})$$

"If I had known all the outcomes ahead of time, I would have chosen differently — and I regret the extra loss I've accumulated."

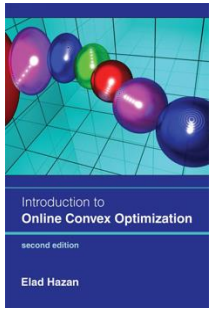


Caravaggio, The Cardsharps, c. 1596–1597, Kimbell Art Museum.

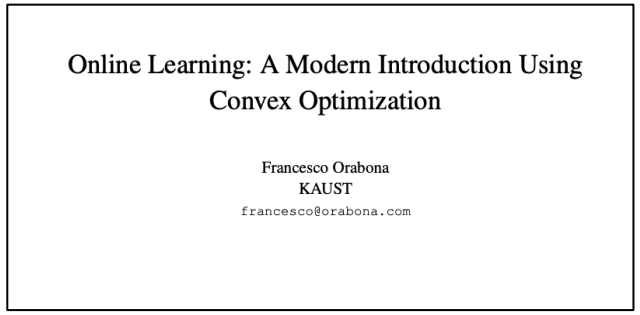
Non-stationary Environments

- There have been rich theories for regret minimization.

$$\text{REG}_T = \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \min_{\mathbf{x}\in\mathcal{X}} \sum_{t=1}^{T} f_t(\mathbf{x})$$







- **Non-stationary online learning**: *distribution shift* always happens when data are collected in real and open environments



manufacturing processes evolve over time, with changing products, equipment, and workflows…

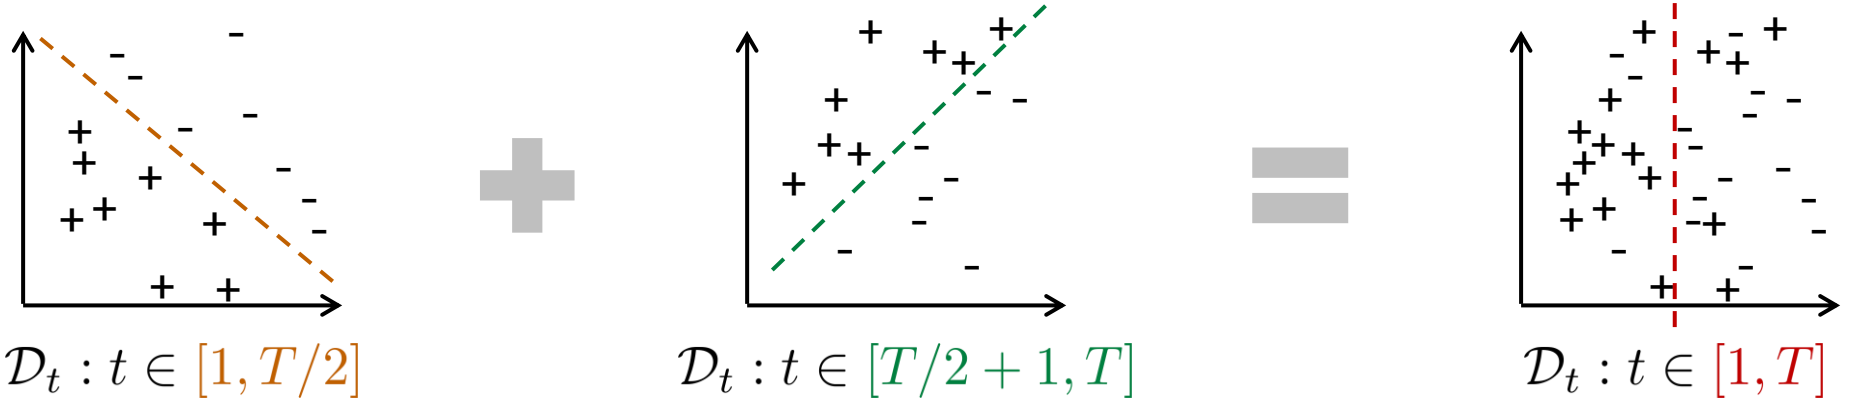
Outline

- Online Learning
- Performance Metric
- Our Result
- Summary

Non-stationary Online Learning

- **Challenge:** static regret can be ***misleading*** in non-stationary env. !

Consider online supervised learning: $f_t(\mathbf{x}) = \ell(\mathbf{x}; \mathbf{a}_t, y_t), \quad (\mathbf{a}_t, y_t) \sim \mathcal{D}_t.$



$\mathcal{D}_t : t \in [1, T/2]$ $\mathcal{D}_t : t \in [T/2+1, T]$ $\mathcal{D}_t : t \in [1, T]$

two stages: ***each stage the best online model incur zero error***

best offline model incurs constant error each round

$$\text{REG}_T \triangleq \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^{T} f_t(\mathbf{x})$$

optimal model might change over time in non-stationary environments

Non-stationary Online Learning

To guide algorithm design for non-stationary environments, there are two common metrics extensively studied in the literature:

- ***Interval*** Regret (also known as strongly adaptive regret):

  for any interval $I = [s, s+L-1] \subseteq [T]$, $\text{I-REG}_I \triangleq \sum_{t \in I} f_t(\mathbf{x}_t) - \min_{\mathbf{u} \in \mathcal{X}} \sum_{t \in I} f_t(\mathbf{u})$.

- ***Dynamic*** Regret (also known as tracking/switching/shifting regret for special cases):

  for any comparators $\mathbf{u}_1, \ldots, \mathbf{u}_T \in \mathcal{X}$, $\text{D-REG}_T(\{\mathbf{u}_t\}_{t=1}^T) \triangleq \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t)$.

Performance Metric

- **Interval Regret**: local time interval, locally fixed comparator

$$\text{I-REG}_I \triangleq \sum_{t\in I} f_t(\mathbf{x}_t) - \min_{\mathbf{u}\in\mathcal{X}} \sum_{t\in I} f_t(\mathbf{u})$$

ensure small regret on every interval

$$\text{I-REG}_I \le \widetilde{\mathcal{O}}(\sqrt{|I|}), \quad \text{(convex loss)},$$

$$\text{I-REG}_I \le \widetilde{\mathcal{O}}_d(\log|I|), \quad \text{(curved loss)},$$

simultaneously for any $I \subseteq [T]$

IR provides local guarantees in time.

- **Dynamic Regret**: full time horizon, globally changing comparators

$$\text{D-REG}_T(\{\mathbf{u}_t\}_{t=1}^T) \triangleq \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t)$$

ensure small regret against changing comparators

$$\text{D-REG}_T(\{\mathbf{u}_t\}_{t=1}^T) \le \mathcal{O}(P_T^{\frac{1}{2}} T^{\frac{1}{2}}), \quad \text{(convex)},$$

$$\text{D-REG}_T(\{\mathbf{u}_t\}_{t=1}^T) \le \widetilde{\mathcal{O}}_d(P_T^{\frac{2}{3}} T^{\frac{1}{3}}), \quad \text{(curved)},$$

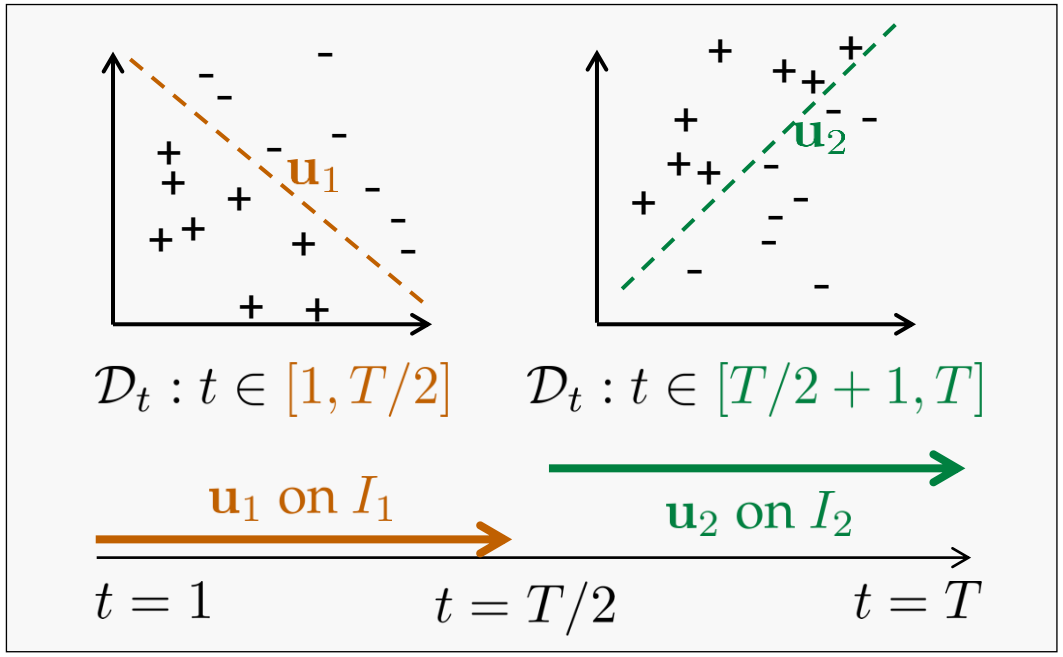
where $P_T = \sum_{t=2}^T \|\mathbf{u}_t - \mathbf{u}_{t-1}\|$ is path length.

DR requires global tracking capability.

Performance Metric

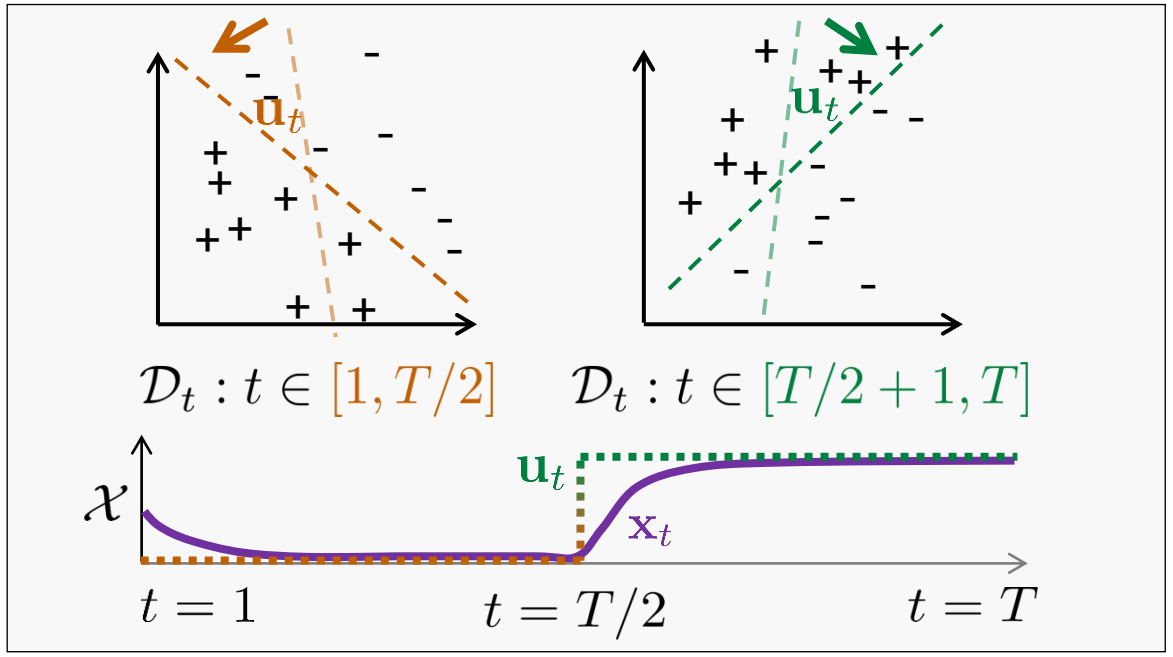
- **Interval Regret**: local time interval, locally fixed comparator

$$\text{I-REG}_I \triangleq \sum_{t \in I} f_t(\mathbf{x}_t) - \min_{\mathbf{u} \in \mathcal{X}} \sum_{t \in I} f_t(\mathbf{u})$$



- **Dynamic Regret**: full time horizon, globally changing comparators

$$\text{D-REG}_T(\{\mathbf{u}_t\}_{t=1}^T) \triangleq \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t)$$



Performance Metric

$$\text{I-REG}_I \triangleq \sum_{t\in I} f_t(\mathbf{x}_t) - \min_{\mathbf{u}\in\mathcal{X}} \sum_{t\in I} f_t(\mathbf{u}) \qquad \text{D-REG}_T(\{\mathbf{u}_t\}_{t=1}^T) \triangleq \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t)$$

loss type	Optimal Interval Regret (OIR)	Optimal Dynamic Regret (ODR)
convex	$\text{I-REG}_I = \widetilde{\mathcal{O}}\left(\sqrt{\lvert I\rvert}\right)$	$\text{D-REG}_T = \widetilde{\mathcal{O}}\left(\sqrt{T} + P_T^{\frac{1}{2}} T^{\frac{1}{2}}\right)$
curved	$\text{I-REG}_I = \widetilde{\mathcal{O}}_d\left(\log \lvert I\rvert\right)$	$\text{D-REG}_T = \widetilde{\mathcal{O}}_d\left(1 + P_T^{\frac{2}{3}} T^{\frac{1}{3}}\right)$

exponentially concave & strongly convex

where $P_T \triangleq \sum_{t=2}^T \|\mathbf{u}_t - \mathbf{u}_{t-1}\|$

- Interval and dynamic regret are largely studied ***separately***.
- In fact, the relation between interval and dynamic regret is largely ***unclear***.

Relation between Two Metrics?

An intuition on the metric-level implication:

Optimal Interval Regret (OIR) → Optimal Dynamic Regret (ODR)

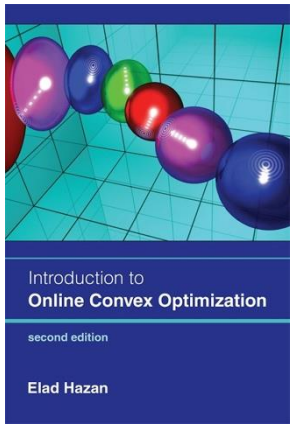
local properties → global properties

Zhang et al. (2018) and its textbook presentation **Hazan (2022)**:

Interval Regret

We now turn to another metric of performance that requires more advanced methods than we have seen thus far. This metric can be shown to be more general than dynamic regret, in the sense that the bounds we prove also imply low dynamic regret.

Optimal Interval Regret



Elad Hazan (2022). *"Introduction to Online Convex Optimization."* MIT Press.

Outline

- Online Learning
- Performance Metric
- Our Result
- Summary

Clarifying Literature

- **Zhang et al. (2018)** (cited by textbook **Hazan (2022)**): "OIR → ***worst-case*** ODR"

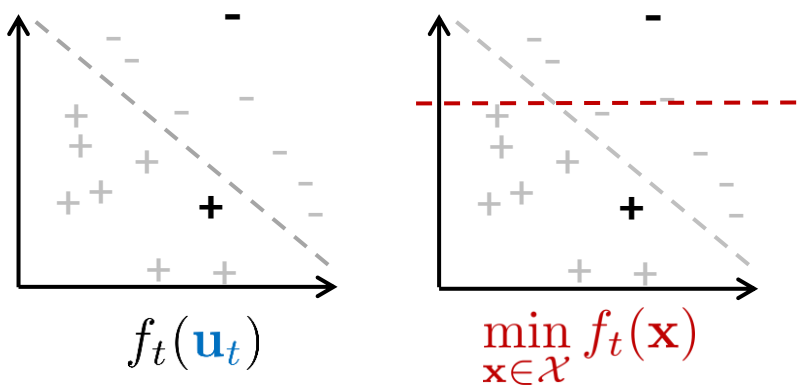
$$\text{D-REG}_T^{\text{worst-case}} \triangleq \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} \min_{\mathbf{x}\in\mathcal{X}} f_t(\mathbf{x})$$

Remark: ***worst-case*** dynamic regret is a special form of ***general*** dynamic regret;

but it encourages ***overfitting*** to the *individual loss function*

e.g., $\{f_t\}_{t=1}^T$ are drawn from underlying distributions, this measure ***overfits noise***

$$f_t(\mathbf{x}) = \ell(\mathbf{x}; \mathbf{a}_t, y_t), \quad (\mathbf{a}_t, y_t) \sim \mathcal{D}_t.$$



Clarifying Literature

Nevertheless, the intuition does hold on ***special domains***.

- **Luo & Schapire (2015)**: for linear loss on simplex, "OIR $\rightarrow$ ODR"

$$\text{OIR: I-REG}_I = \widetilde{\mathcal{O}}\left(\sqrt{|I|}\right) \rightarrow \text{ODR: D-REG}_T = \widetilde{\mathcal{O}}\left(\sqrt{T} + P_T^{\frac{1}{2}} T^{\frac{1}{2}}\right)$$

- **Baby & Wang (2021, 2022)**: for curved loss on cube (L_∞ ball), "OIR $\rightarrow$ ODR"

$$\text{OIR: I-REG}_I = \widetilde{\mathcal{O}}_d\left(\log |I|\right) \rightarrow \text{ODR: D-REG}_T = \widetilde{\mathcal{O}}_d\left(1 + P_T^{\frac{2}{3}} T^{\frac{1}{3}}\right)$$

Clarifying Literature

- **"general domain, but for a *special* metric implication"**
  - Zhang et al. (2018): "OIR → **worst-case** ODR".
- **"general metric implication, but for a *special* domain"**
  - Luo & Schapire (2015): "OIR → ODR" for linear loss on ***simplex***;
  - Baby & Wang (2021, 2022): "OIR → ODR" for curved loss on ***cube***.

In general, OIR ***not necessarily*** implies ODR, with unclear relation

Other Efforts: Optimizing Two Together



- For Expert setting, a ***generalized shifting regret*** is proposed:

3 A generalized shifting regret for the simplex

We now introduce a generalized notion of shifting regret which unifies and generalizes the notions of discounted regret (see [3, Section 2.11]), adaptive regret (see [8]), and shifting regret (see [2]). For a fixed horizon T, a sequence of discount factors $\beta_{t,T} \geqslant 0$ for $t = 1, \ldots, T$ assigns varying weights to the instantaneous losses suffered at each round. We compare the total loss of the forecaster with the loss of an arbitrary sequence of vectors $\boldsymbol{q}_1, \ldots, \boldsymbol{q}_T$ in the simplex Δ_d. Our goal is to bound the regret

$$\sum_{t=1}^{T} \beta_{t,T}\, \widehat{\boldsymbol{p}}_t^{\top} \boldsymbol{\ell}_t - \sum_{t=1}^{T} \beta_{t,T}\, \boldsymbol{q}_t^{\top} \boldsymbol{\ell}_t$$

Nicolò Cesa-Bianchi, Pierre Gaillard, Gabor Lugosi, and Gilles Stoltz. "Mirror Descent Meets Fixed Share (and Feels No Regret)". *NIPS 2012*.

- For OCO, people studies a metric called ***interval dynamic regret***:

$$\textsc{Interval-D-Reg}_T(|I|; \mathbf{u}_1, \ldots, \mathbf{u}_T) = \max_{[r, r+\tau-1] \subseteq T} \left\{ \sum_{t=r}^{r+\tau-1} f_t(\mathbf{x}_t) - \sum_{t=r}^{r+\tau-1} f_t(\mathbf{u}_t) \right\}.$$

Our Results

Guiding question:

Is there any ***metric-level implication*** from interval regret to dynamic regret?

Namely, if an online learning algorithm is equipped with optimal interval regret guarantees, does it automatically enjoy optimal dynamic regret?

- **Negative result**: we refute the "metric-level implication".
- **Positive result**: we construct an "algorithm-level connection".

I. Metric-Level Implication

Negative Result: we refute the "metric-level implication"

- We first try to extend argument of Luo & Schapire (2015) (OLO with ***simplex***),

  → partially successful: the implication works for ***polytope-like domains***

- Unfortunately, this does not hold in general, neither linear nor curved loss

  → Counterexamples on ***smooth convex bodies*** ($L_2/L_{2/3}$ ***ball***):

Linear: $\text{I-REG}_I = \widetilde{\mathcal{O}}\big(\sqrt{|I|}\big) \;\to\; \text{D-REG}_T = \Omega\big(P_T^{2/5} T^{3/5}\big) \succ \widetilde{\mathcal{O}}\big(P_T^{1/2} T^{1/2}\big)$

Curved: $\text{I-REG}_I = \widetilde{\mathcal{O}}_d\big(\log |I|\big) \;\to\; \text{D-REG}_T = \Omega_d\big(P_T^{3/5} T^{2/5}\big) \succ \widetilde{\mathcal{O}}_d\big(P_T^{2/3} T^{1/3}\big)$

Key finding: ***domain geometry*** essentially governs the metric-level implication.

II. How to Leverage Local Adaptivity?

Positive Result: we construct an "algorithm-level connection"

- ***Domain-conversion technique*** to bypass the constrained of ***domain geometry***

> For ODR $(\mathcal{X}, \{f_t\}_{t=1}^T)$, we run OIR algorithm on the "twisted problem"
>
> – an enlarged ball $\mathcal{Y} \supset \mathcal{X}$, and
>
> – a domain-converted surrogate loss function $h_t : \mathcal{Y} \to \mathbb{R}$

OIR on $\mathcal{Y}$ and $\{h_t\}_{t=1}^T$: I-REG$_I = \widetilde{\mathcal{O}}\big(\sqrt{|I|}\big)$ $\to$ ODR on $\mathcal{X}$ and $\{f_t\}_{t=1}^T$: D-REG$_T = \widetilde{\mathcal{O}}\big(P_T^{1/2} T^{1/2}\big)$

OIR on $\mathcal{Y}$ and $\{h_t\}_{t=1}^T$: I-REG$_I = \widetilde{\mathcal{O}}_d\big(\log |I|\big)$ $\to$ ODR on $\mathcal{X}$ and $\{f_t\}_{t=1}^T$: D-REG$_T = \widetilde{\mathcal{O}}_d\big(P_T^{2/3} T^{1/3}\big)$

- The argument can be extended to *higher-order difference* dynamic regret.

Technical Highlight

- I2D Decomposition
- Geometric Intuition
- Domain conversion and Improper Learning
- Extension to Higher-Order Path Length

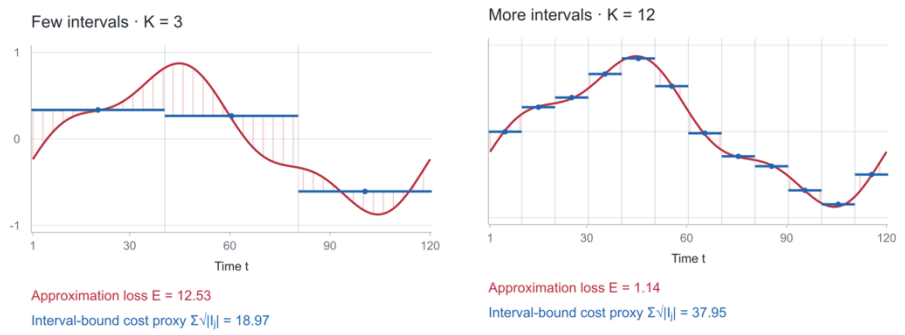
I2D Decomposition

It's natural to decompose dynamic regret via interval regret: splitting $[1, T]$ into disjoint blocks $\{I_j\}_{j=1}^{K}$, with $\mathbf{z}_j$ being comparator on I_j

$$\text{D-REG}_T = \sum_{j=1}^{K} \sum_{t \in I_j} (f_t(\mathbf{x}_t) - f_t(\mathbf{u}_t))$$

$$= \sum_{j=1}^{K} \sum_{t \in I_j} (f_t(\mathbf{x}_t) - f_t(\mathbf{z}_j) + f_t(\mathbf{z}_j) - f_t(\mathbf{u}_t))$$

$$= \sum_{j=1}^{K} \left(\text{I-REG}_{I_j}(\mathbf{z}_j) + \text{ERR}_{I_j}(\mathbf{z}_j)\right).$$



- Too many blocks → I-Reg blows
- Too few blocks → Approx. Error blows

Hard Case: A Geometric Example

Can two nearby minimizers share a good static comparator?

Feasible Domain

$$\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^2 : \|\mathbf{x}\|_2 \leq 1\}.$$

Comparators

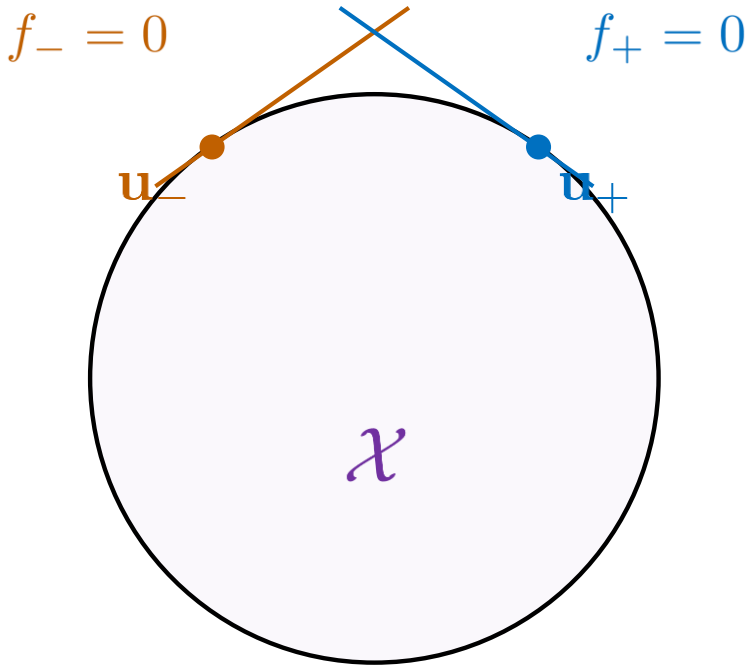
$$\mathbf{u}_\pm = (\pm \sin\theta, \cos\theta); \text{where } 0 < \theta < \pi/3$$

$$\|\mathbf{u}_+ - \mathbf{u}_-\|_2 = 2\sin\theta = \Theta(\theta) =: \epsilon$$

Online function

$$f_s(\mathbf{x}) = 1 - \langle \mathbf{u}_s, \mathbf{x} \rangle, \quad s \in \{-, +\}$$

$$\min_{\mathbf{x} \in \mathcal{X}} f_s(\mathbf{x}) = f_s(\mathbf{u}_s) = 0$$



A Fixed Comparator Pays a Gap

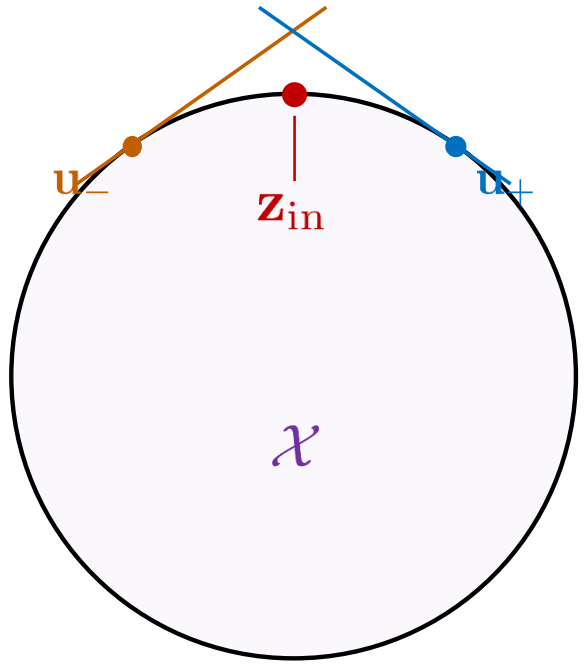
Suppose both losses occur equally often in one interval block.

The moving comparators pay zero loss $\sum_{t\in I} f_t(\mathbf{u}_t) = 0$

However, even the best fixed point must compromise

$$\frac{f_+(\mathbf{z}) + f_-(\mathbf{z})}{2} = 1 - (\cos\theta) z_2, \text{ taking minimum at } \mathbf{z}_{\text{in}} = (0, 1)$$

$$\Delta := \min_{\mathbf{z}\in\mathcal{X}} \frac{f_+(\mathbf{z}) + f_-(\mathbf{z})}{2} = 1 - \cos\theta \approx \frac{\theta^2}{2} > 0$$



Over the interval, approx. error gap accumulates: $\min_{\mathbf{z}\in\mathcal{X}} \sum_{t\in I} \big[f_t(\mathbf{z}) - f_t(\mathbf{u}_t)\big] = |I|\Delta$

From Interval to Dynamic Regret

The comparator $\{\mathbf{u}_t\}_{t=1}^T$ flips between $\mathbf{u}_-$ and $\mathbf{u}_+$ by time block of length B,

then $P_T = \Theta(\epsilon T/B)$

f_-	f_+	f_-	f_+	f_-	f_+

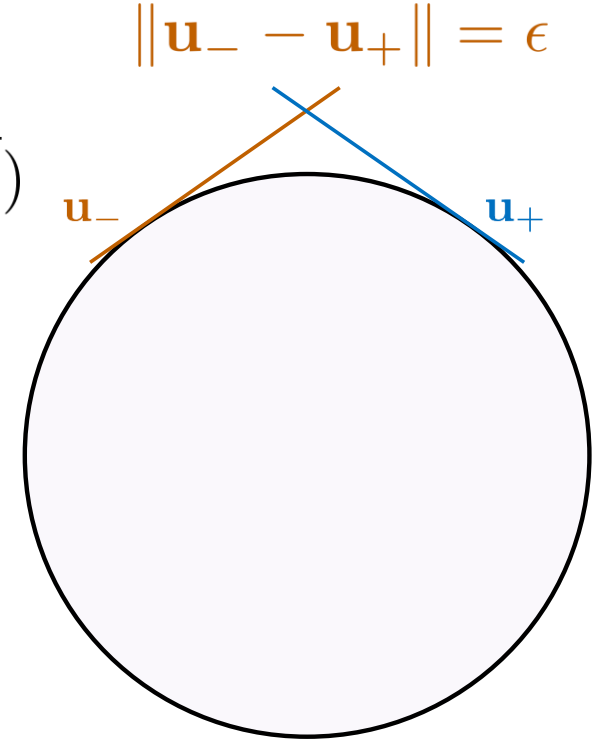
- If #short-interval (length $\leq B$) dominates, I-Reg blows

$$\text{D-REG}_T \sim \sum_{I:|I|\lesssim B} \text{I-REG}_I \sim (T/(2B)) \cdot \widetilde{\Theta}(\sqrt{B}) = \widetilde{\Theta}(\sqrt{P_T T/\epsilon})$$

- If #long-interval (length $> B$) dominates, approx. error blows

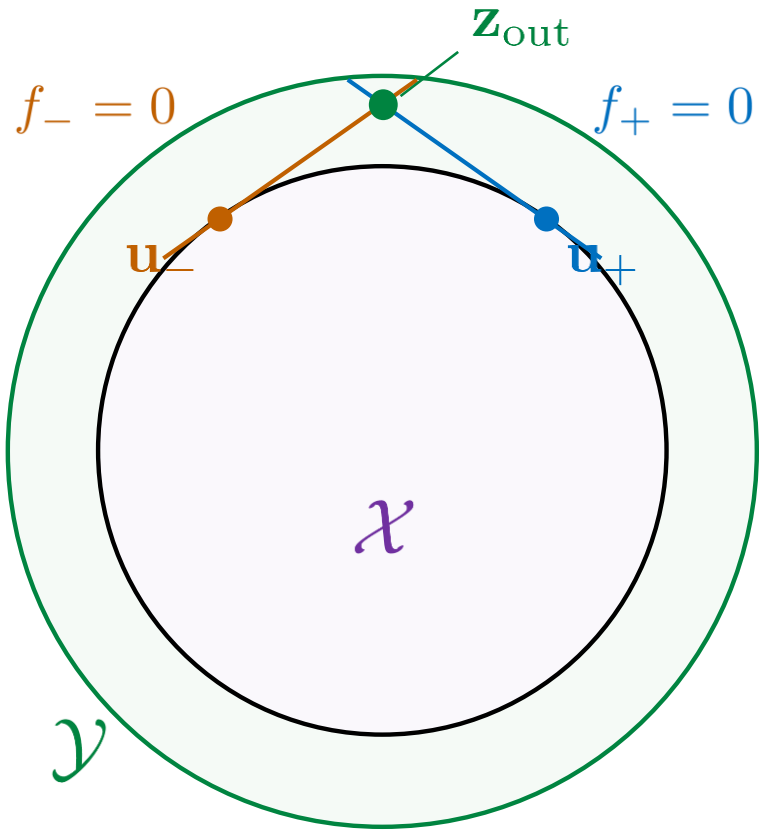
$$\text{D-REG}_T \sim \sum_{I:|I|\gtrsim B} \text{ERR}_I \sim (T/2) \cdot \epsilon^2$$

- Blancing $\epsilon = \Theta((P_T/T)^{1/5})$ gives $\text{D-REG}_T = \Omega\left(P_T^{2/5} T^{3/5}\right)$



Extra Space for Comparison

The two zero-loss lines intersect outside the original domain.



Same losses and
same moving comparator

The tangent intersection

$$\mathbf{z}_{\text{out}} = (0, \sec\theta) \notin \mathcal{X}$$

$$f_+(\mathbf{z}_{\text{out}}) = f_-(\mathbf{z}_{\text{out}}) = 0$$

One static point now matches both moving minima.

Domain conversion

Use surrogate losses h_t on the enlarged domain.

$$\mathbf{y}_t \in \mathcal{Y}, \quad \mathbf{x}_t = \Pi_{\mathcal{X}}(\mathbf{y}_t) \in \mathcal{X}$$

$$f_t(\mathbf{x}_t) - f_t(\mathbf{u}_t) \le h_t(\mathbf{y}_t) - h_t(\mathbf{u}_t)$$

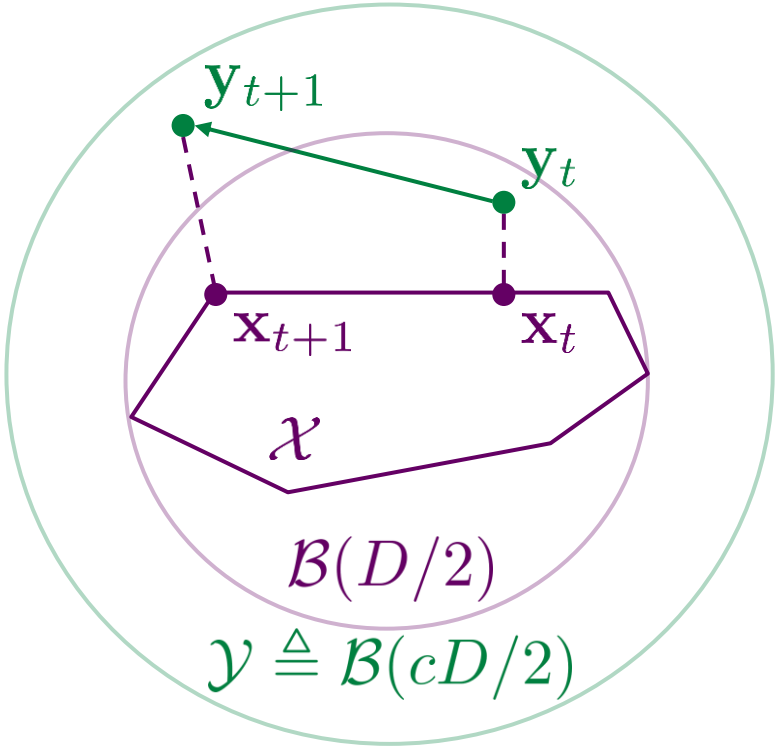
Upper Bound: Domain Conversion

How to circumvent *Domain Geometry*?

- *Domain Conversion*: run *OIR algorithm* on an enlarged ball over surrogate loss

  extremely(!) useful techniques **Cutkosky and Orabona (2018); Cutkosky (2020)**

$$\mathbf{x}_t = \Pi_{\mathcal{X}}[\mathbf{y}_t] \text{ (L_2-norm projection)}$$
$$\mathbf{g}_t = \nabla f_t(\mathbf{x}_t)$$
$$\mathbf{d}_t = \mathbf{g}_t + \frac{\max\{-\mathbf{g}_t^\top \mathbf{n}_t, 0\}}{\|\mathbf{n}_t\|_2^2}\mathbf{n}_t, \mathbf{n}_t = \mathbf{y}_t - \mathbf{x}_t$$
$$h_t(\mathbf{y}) = \mathbf{d}_t^\top \mathbf{y} \text{ (for convexity)}$$



Upper Bound: Domain Conversion

How to circumvent ***Domain Geometry***?

- ***Domain Conversion***: run ***OIR algorithm*** on an enlarged ball over surrogate loss

extremely(!) useful techniques **Cutkosky and Orabona (2018); Cutkosky (2020)**

Theorem. The domain-converted loss satisfies:

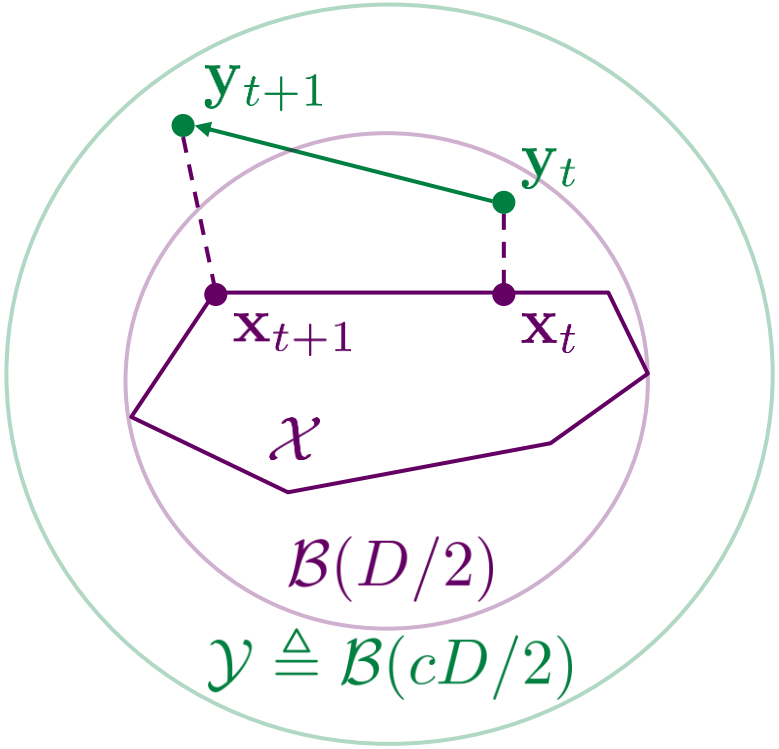
(i) h_t inherits regret: $f_t(\mathbf{x}_t) - f_t(\mathbf{u}_t) \leq h_t(\mathbf{y}_t) - h_t(\mathbf{u}_t)$;

(ii) h_t inherits Lipschitzness, convexity, exp-concavity, and strong convexity of f_t.

$\mathbf{y}_t \in \mathcal{Y}$

$\mathbf{u}_t \in \mathcal{X}$

→ *unfair competition (improper learning) over* h_t



Improper Learning Benefits

- We can exploit this improper learning to prove "***unfair***" I2D Decomposition

$$\text{D-REG}_T(\{h_t\}_{t=1}^T; \{\mathbf{u}_t\}_{t=1}^T) = \sum_{j=1}^K \text{I-REG}_{I_j}(\mathbf{z}_j), \qquad \text{ERR}_{I_j}(\mathbf{z}_j) \leq 0$$

Note: D-REG_T is evaluated with $h_t(\mathbf{y}_t) - h_t(\mathbf{u}_t)$ (improper learning)

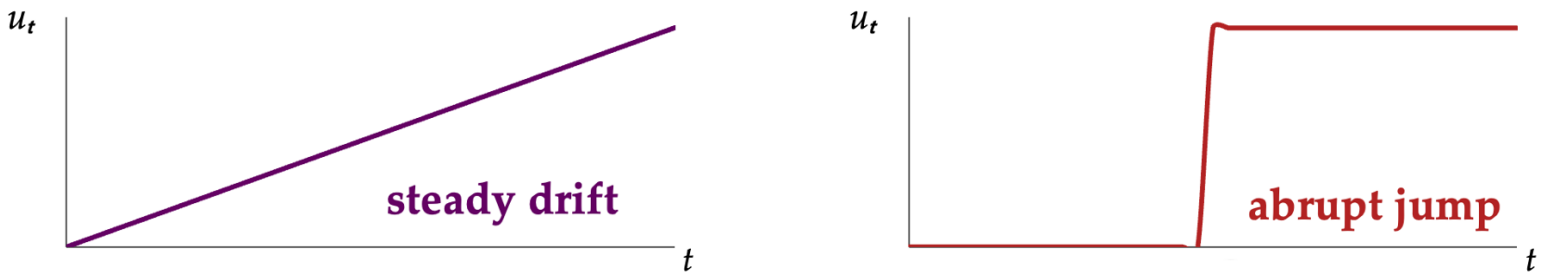
Using **Baby & Wang (2021)**'s partitions of $[1, T]$ into $K = \mathcal{O}(P_T)$ long intervals,

$$\begin{aligned}
\text{D-REG}_T &= \sum_{j=1}^K \left(\text{I-REG}_{I_j}(\mathbf{z}_j) + \text{ERR}_{I_j}(\mathbf{z}_j)\right) \\
&\leq \sum_{j=1}^K \text{I-REG}_{I_j}(\mathbf{z}_j) \\
&\leq \sum_{j=1}^K \widetilde{O}\left(\sqrt{|I_j|}\right) \leq \widetilde{\mathcal{O}}\left(\sqrt{KT}\right) = \widetilde{\mathcal{O}}\left(\sqrt{P_T T}\right)
\end{aligned}$$

Higher-order Path Length

- Path length is a common non-stationarity measure: but it measures total movement, but ***misses how comparators change over time***.

Same path length ($P_T = \sum_{t=2}^{T} |u_t - u_{t-1}| = 1$), **different non-stationarity.**



→ High-order smoothness: to measure changes in the increments.

Higher-order Differences

- First-order differences recover standard path length

$$\mathsf{D}u_t := u_t - u_{t-1}, \text{ and } P_T = \sum_{t=2}^{T} \|\mathsf{D}u_t\|$$

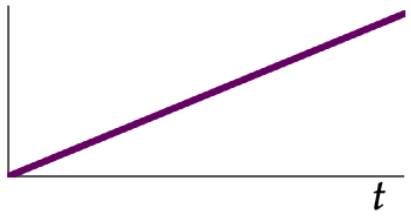
- Second differences compare consecutive increments

$$\mathsf{D}^2 u_t := \mathsf{D}u_t - \mathsf{D}u_{t-1} = (u_t - u_{t-1}) - (u_{t-1} - u_{t-2})$$

Repeat D to obtain higher-order differences.

- Higher-order smoothness: $P_{T,k} = T^{k-1} \sum_{t=k+1}^{T} \|\mathsf{D}^k \mathbf{u}_t\|_2, \qquad P_{T,1} = P_T$

Example: linear drift



$$\begin{aligned} u_t &: \quad 0,\ 1,\ 2,\ 3,\ \ldots \\ \mathsf{D}u_t &: \quad 1,\ 1,\ 1,\ \ldots \\ \mathsf{D}^2 u_t &: \quad 0,\ 0,\ \ldots \end{aligned}$$

Equal increments give zero second differences.

Higher-order Dynamic Regret

- Our method: Same wrapper with a polynomial interval learner (Baby and Wang, 2022) on an enlarged domain with domain-converted surrogate loss.

$$\text{Degree } \leq k-1: \quad \text{constants } (k=1),\ \text{lines } (k=2),\ \text{quadratics } (k=3).$$

- **convex loss**: $\text{D-REG}_T \leq \widetilde{\mathcal{O}}_k\left((1+P_{T,k}^{\frac{1}{2k}})\sqrt{T}\right)$
- **curved loss**: $\text{D-REG}_T \leq \widetilde{\mathcal{O}}_{d,k}\left(1+P_{T,k}^{\frac{2}{2k+1}}T^{\frac{1}{2k+1}}\right)$

$$\text{Linear drift with } P_T=1: \quad \widetilde{\mathcal{O}}(T^{1/3})\ (k=1) \quad \longrightarrow \quad \widetilde{\mathcal{O}}(1)\ (k=2)$$

Outline

- Online Learning
- Performance Metric
- Our Result
- Summary

Summary

- **Problem**: the relation between *interval regret (IR)* and *dynamic regret (DR)*, two crucial performance metrics in non-stationary online learning.
- **Our Result**: negative and positive sides
  - For an intuitive misunderstanding, **we refute "metric-level IR implies DR"** and reveal how *domain geometry* governs this metric-level implication.
  - To correctly leverage interval regret, we construct an efficient algorithm wrapper "**algorithm-level IR helps DR**" by *domain-conversion technique*.



[Wang-Z-Zhou, arXiv:2609.34423] On the Relation between Interval Regret and Dynamic Regret.

Thanks!