Fast Rates in Time-Varying Strongly Monotone Games

Yu-Hu Yan, Peng Zhao, Zhi-Hua Zhou

Contact

{yanyh, zhaop, zhouzh}@lamda.nju.edu.cn

NANJING UNIVERSITY

LAMDA Learning And Mining from DatA

ICML International Conference On Machine Learning 40 Years

Time-Varying Games

At each round $t = 1, 2, \ldots, T$:

- each player ($i \in [N]$) submits $x_{t,i} \in \mathcal{X}_i \subseteq \mathbb{R}^d$ respectively
- simultaneously, environments reveal a group of ***time-varying*** utility functions $u_{t,i} : \mathcal{X} \mapsto \mathbb{R}^+$ for each player, where $\mathcal{X} \triangleq \mathcal{X}_1 \times \ldots \times \mathcal{X}_N$
- the i-th player suffers loss $u_{t,i}(\mathbf{x}_t)$ and receives $v_{t,i}(\mathbf{x}_t) \triangleq \nabla_{x_{t,i}} u_{t,i}(x_{t,i}; x_{t,-i})$, where $\mathbf{x}_t \triangleq (x_{t,1}, \ldots, x_{t,N})$

Many real-world games are time-varying


Rush-hour traffic


Modern air combat


Stock market

Strongly Monotone Games

A general class containing many games of interest [1]

Monotone games: $\langle v(\mathbf{x}) - v(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle \geq 0$

Zero-sum games

Example: rock-paper-scissors game

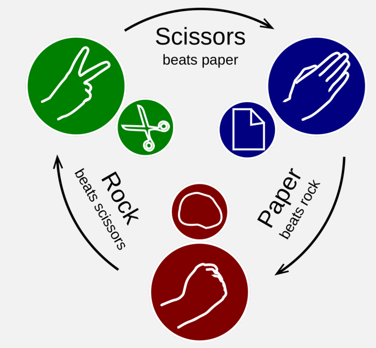

Time-varying zero-sum games: [2]

Strongly monotone games

$\langle v(\mathbf{x}) - v(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle \geq \mu \|\mathbf{x} - \mathbf{y}\|_2^2$

Example: Cournot competition

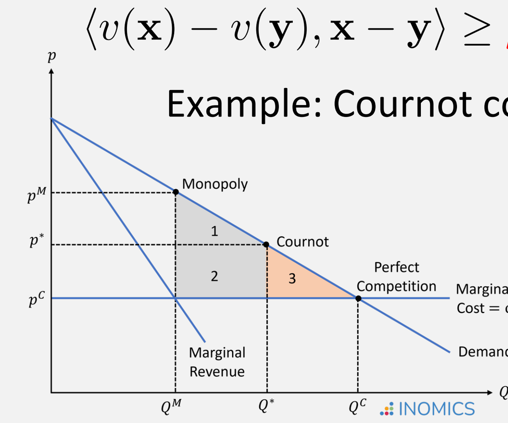

Time-varying strongly monotone games: **ours**

[1] Rosen, J. B.. Existence and uniqueness of equilibrium points for concave n-person games, Econometrica, 1965
[2] Zhang et al., No-Regret Learning in Time-Varying Zero-Sum Games, ICML'22

Performance and Non-Stationarity Measure

Measure the Performance:

The goal of time-varying games is chasing the time-varying Nash equilibriums $\{\mathbf{x}_t^\star\}_{t=1}^T$

Distance tracking error: $\textsc{Dist-Err} \triangleq \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_t^\star\|^2$

Measure the Environmental Non-Stationarity:

- *Path length*: $P_T \triangleq \sum_{t=2}^T \|\mathbf{x}_t^\star - \mathbf{x}_{t-1}^\star\|$
- *Gradient variation*: $V_T \triangleq \sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{X}} \|v_t(\mathbf{x}) - v_{t-1}(\mathbf{x})\|^2$
- *Gradient variance*: $W_T \triangleq \sum_{t=1}^T \sup_{\mathbf{x} \in \mathcal{X}} \|v_t(\mathbf{x}) - \bar{v}_T(\mathbf{x})\|$

Main Results:

Works	Time-Varying Game	Time-Invariant Game
Duvocelle et al. (2021)	$\mathcal{O}(\sqrt{T} + T^{2/3} P_T^{1/3})$	$\mathcal{O}(\sqrt{T})$
Non-Smooth [Theorem 3]	$\tilde{\mathcal{O}}(1 + \min\{T^{1/3} P_T^{2/3}, W_T\})$	$\tilde{\mathcal{O}}(1)$
Smooth [Theorem 5]	$\mathcal{O}(\min\{\sqrt{(1 + V_T + P_T)(1 + P_T)}, 1 + W_T\})$	$\mathcal{O}(1)$

Fate Rates in Non-Smooth Games - I

Revisiting Existing Results [3]**:**

Previous work uses periodical restarts to handle non-stationarity

$$\textsc{Dist-Err} \triangleq \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_t^\star\|^2 \leq \sum_{k=1}^K \sum_{t \in \Delta_k} \langle v(\mathbf{x}_t), \mathbf{x}_t - \mathbf{v}_k \rangle + \sum_{k=1}^K \sum_{t \in \Delta_k} \langle v(\mathbf{x}_t), \mathbf{v}_k - \mathbf{x}_t^\star \rangle,$$

Choosing the restart period as $\Delta = \min\{(T/P_T)^{2/3}, T\}$, previous work obtains a distance tracking error of $\mathcal{O}\left(\sqrt{T} + T^{2/3} P_T^{1/3}\right)$

Disadvantages:
- require known P_T
- restarts are ineffective when handling non-stationarity

Our Solution: tracking the non-stationarity directly

We handle non-stationarity directly via online ensemble

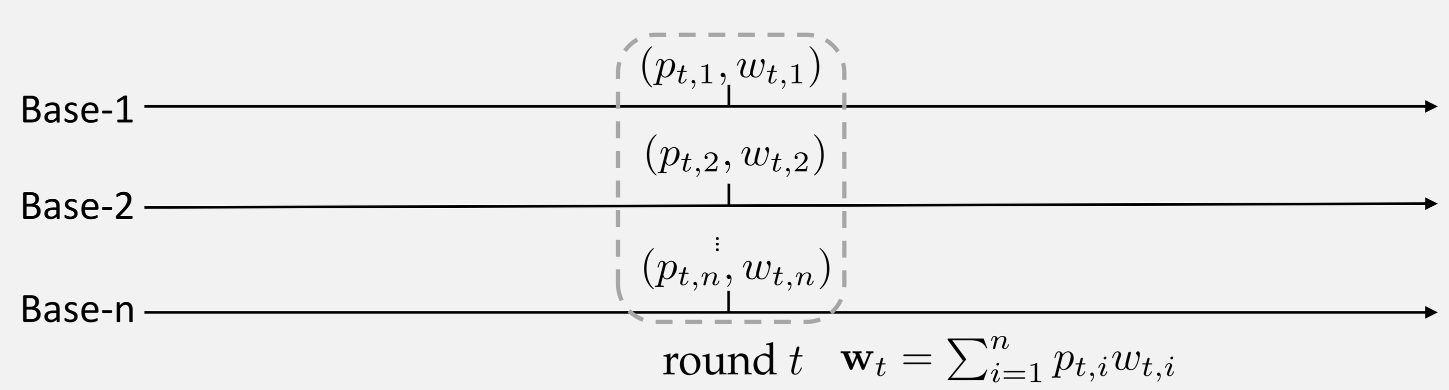

[3] Duvocelle et al.. Multi-Agent Online Learning in Time-Varying Games, MOR'21

Fate Rates in Non-Smooth Games - II

A Direct Solution:

$$\textsc{Dist-Err} \triangleq \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_t^\star\|^2 \leq \sum_{t=1}^T \langle v(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}_t^\star \rangle \leq \sum_{i=1}^N \sum_{t=1}^T \langle v_{t,i}(\mathbf{x}_t), x_{t,i} - x_{t,i}^\star \rangle,$$

(*a dynamic regret minimization problem*)

Algorithm: *use OGD as base algorithm, and Hedge as meta algorithm*

$\mathcal{O}\left(\sqrt{T(1+P_T)}\right) \leq \mathcal{O}\left(\sqrt{T} + T^{2/3} P_T^{1/3}\right)$ without knowing P_T

Improved result without exploring the problem structure

An Improved Solution: strong convexity from strong monotonicity

Proposition 1. *The distance tracking error can be upper-bounded by*

$$\mu \textsc{Dist-Err} \leq 2 \sum_{i=1}^N \sum_{t=1}^T \left(\ell_{t,i}(x_{t,i}) - \ell_{t,i}(x_{t,i}^\star)\right), \quad \textit{(surrogate loss)}$$

where $\ell_{t,i}(x) \triangleq \langle v_{t,i}(\mathbf{x}_t), x \rangle + \frac{\mu}{2}\|x - x_{t,i}\|^2$ *is μ-strongly convex.*

Methods for non-stationary online learning with strongly convex losses can be used

$$\tilde{\mathcal{O}}\left(1 + T^{1/3} P_T^{2/3}\right) \leq \mathcal{O}\left(\sqrt{T(1+P_T)}\right) \leq \mathcal{O}\left(\sqrt{T} + T^{2/3} P_T^{1/3}\right)$$

Improved result ***with*** exploring the problem structure:
- $\tilde{\mathcal{O}}(1)$ recovers the best known static bound
- $T^{1/3} P_T^{2/3}$ improves $T^{2/3} P_T^{1/3}$
- does not require P_T as input

The same algorithm also enjoys $\tilde{\mathcal{O}}(1 + W_T)$ gradient-variance bound

Faster Rates in Smooth Games

Smooth Games:

Smoothness: $\|v_t(\mathbf{x}) - v_t(\mathbf{y})\| \leq L \|\mathbf{x} - \mathbf{y}\|$

Common in literature: for example, *two-player zero-sum games* $f(\mathbf{x}, \mathbf{y}) = \mathbf{x}^\top A \mathbf{y}$

$$\|\nabla f(\mathbf{x}_1, \mathbf{y}_1) - \nabla f(\mathbf{x}_2, \mathbf{y}_2)\| = \|(A\mathbf{y}_1, -A\mathbf{x}_1) - (A\mathbf{y}_2, -A\mathbf{x}_2)\| \leq \|A\| \|(\mathbf{y}_1 - \mathbf{y}_2, \mathbf{x}_2 - \mathbf{x}_1)\|$$

Dynamic Regret bounded by Variation in Utilities (DRVU):

If each player runs a single algorithm:

We extend DRVU from *two-player finite games* [2] to *multi-layer convex games*

$$\sum_{t=1}^T \langle v_{t,i}(\mathbf{x}_t), x_{t,i} - x_{t,i}^\star \rangle \lesssim \frac{1 + P_{T,i}}{\eta_i} + \eta_i(1 + V_T) + \eta_i \sum_{j=1}^N S_j - \frac{1}{\eta_i} S_i$$

where $S_j \triangleq \sum_{t=2}^T \|x_{t,j} - x_{t-1,j}\|^2$

Regret summation brings cancellations and thus faster rates

If each player runs a two-layer algorithm: (due to unknown $P_{T,i}$)

Correction Term Injection

- What we want: $\sum_{t=1}^T \langle \boldsymbol{\ell}_t, \boldsymbol{p}_t - \boldsymbol{e}_{i^\star} \rangle \leq ???$
- What we do: $\sum_{t=1}^T \langle \boldsymbol{\ell}_t + \boldsymbol{b}_t, \boldsymbol{p}_t - \boldsymbol{e}_{i^\star} \rangle \leq R_T$ (*optimize on a corrected* loss)
- What we obtain: $\sum_{t=1}^T \langle \boldsymbol{\ell}_t, \boldsymbol{p}_t - \boldsymbol{e}_{i^\star} \rangle \leq R_T - \sum_{t=1}^T \sum_{i=1}^N p_{t,i} b_{t,i} + \sum_{t=1}^T b_{t,i^\star}$

Our Result: $\textsc{Dist-Err} \leq \mathcal{O}\left(\min\{\sqrt{(1 + V_T + P_T)(1 + P_T)}, 1 + W_T\}\right)$

Faster rates: fully independent of time horizon T

[2] Zhang et al., No-Regret Learning in Time-Varying Zero-Sum Games, ICML'22