ICML 40 Years International Conference On Machine Learning

LAMDA Learning And Mining from DatA

Fast Rates in Time-Varying Strongly Monotone Games


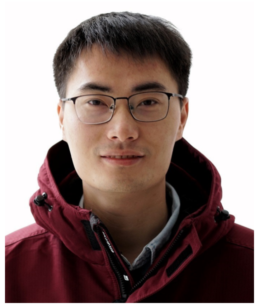


Yu-Hu Yan Peng Zhao Zhi-Hua Zhou

LAMDA Group
Nanjing University

Time-Varying Games

Learning And Mining from DatA

Nash equilibriums

a stable set of decisions where no player has incentives to deviate

Protocol

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 $\mathbf{x}_t \triangleq (x_{t,1}, \ldots, x_{t,N})$

The goal of time-varying games is chasing the time-varying Nash equilibriums

Monotone Games

- A general class containing many games of interest

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

Zero-sum games

Example: rock-paper-scissors game

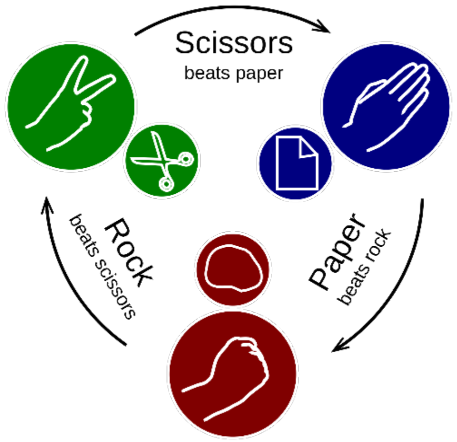

Time-varying zero-sum games: [Zhang et al., 2022]

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

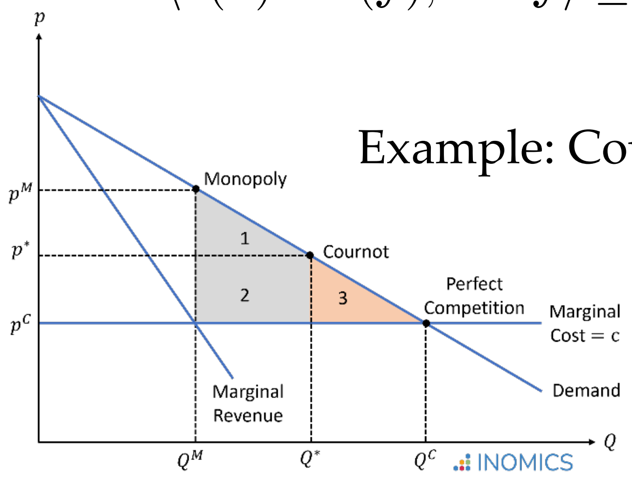

Time-varying strongly monotone games: **this work**

Measures and Results

Learning And Mining from DatA

- Performance Measure: *tracking error* $\sum_{t=1}^{T} \|\mathbf{x}_t - \mathbf{x}_t^\star\|^2$

measure the distance to Nash equilibriums

- Non-Stationarity Measure

  - *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*: $W_T \triangleq \sum_{t=1}^{T} \sup_{\mathbf{x}\in\mathcal{X}} \|v_t(\mathbf{x}) - \bar{v}_T(\mathbf{x})\|$

- Results

Setups	Works	Time-Varying Games	Time-Invariant Games
Non-Smooth Games	Duvocelle et al. (2023)	$\mathcal{O}(\sqrt{T} + T^{2/3} P_T^{1/3})$	$\mathcal{O}(\sqrt{T})$
	This Paper (Theorem 3)	$\widetilde{\mathcal{O}}(1 + \min\{T^{1/3} P_T^{2/3}, W_T\})$	$\widetilde{\mathcal{O}}(1)$
Smooth Games	This Paper (Theorem 5)	$\mathcal{O}(\min\{\sqrt{(1 + V_T + P_T)(1 + P_T)}, 1 + W_T\})$	$\mathcal{O}(1)$

Non-Smooth Games

☐ Existing Method [Duvocelle et al., 2021]

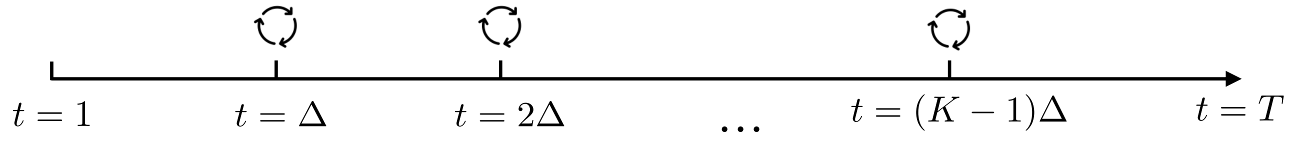

Algorithm: *single model with periodic restarts, can achieve* $\mathcal{O}\left(\sqrt{T} + T^{2/3} P_T^{1/3}\right)$

☐ Our Method: ***Online Ensemble***

$$\text{DIST-ERR} \triangleq \sum_{t=1}^{T} \|\mathbf{x}_t - \mathbf{x}_t^\star\|^2 \leq \sum_{t=1}^{T} \langle v_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}_t^\star \rangle \leq \sum_{i=1}^{N}\sum_{t=1}^{T} \langle v_{t,i}(\mathbf{x}_t), x_{t,i} - x_{t,i}^\star \rangle$$

(a dynamic regret minimization problem)

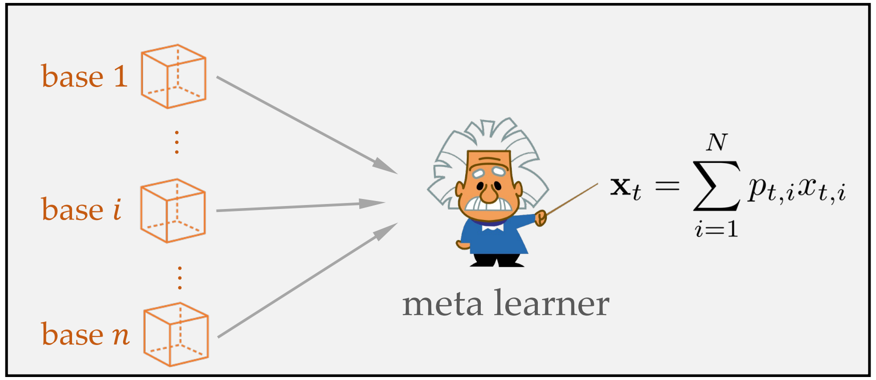

Algorithm: *OGD as base algorithm with different step sizes, Hedge as meta algorithm*

Result: *can achieve* $\mathcal{O}\left(\sqrt{T(1+P_T)}\right)$ *without knowing* $P_T \triangleq \sum_{t=2}^{T} \|\mathbf{x}_t^\star - \mathbf{x}_{t-1}^\star\|$ *in advance*

An initial improvement due to the refined non-stationarity handling.

Non-Smooth Games

◻ A further improvement by exploiting problem structure

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

$$\mu \sum_{t=1}^{T} \|\mathbf{x}_t - \mathbf{x}_t^\star\|^2 \leq 2 \sum_{i=1}^{N} \sum_{t=1}^{T} (\ell_{t,i}(x_{t,i}) - \ell_{t,i}(x_{t,i}^\star)),$$

(*strong convexity from strong monotonicity*)

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

Methods for non-stationary online learning with strongly convex losses can be used [Baby and Wang, 2022]

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

Advantages:
- $\widetilde{\mathcal{O}}(1)$ recovers the best-known static bound
- still does not require P_T as input
- $T^{1/3} P_T^{2/3}$ improves $T^{2/3} P_T^{1/3}$

An orthogonal result of $\widetilde{\mathcal{O}}(1 + W_T)$ can be found in the paper.

Smooth Games

☐ Exploiting smoothness for *faster* rates

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

Common in literature: e.g., satisfied by *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) [Zhang et al., 2022]

If each player runs a single-layer algorithm:

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

where $S_j \triangleq \sum_{t=2}^{T} \|x_{t,j} - x_{t-1,j}\|^2$ *Regret summation over all players brings cancellations and faster rates*

Smooth Games

Algorithm 2 TV-SMOG (smooth) for the i-th player

Input: step size pool $\mathcal{H}_i^\eta$ (D.18)
Initialize: M instances of OOGD $\mathcal{A}_1, \ldots, \mathcal{A}_M$ with step size $\eta_{i,j} \in \mathcal{H}_i^\eta$
for $t = 1, \ldots, T$ **do**
  Receive $x_{t,i,j}$ from $\mathcal{A}_j$
  Submit $x_{t,i} = \sum_{j=1}^M p_{t,i,j} x_{t,i,j}$
  Receive gradient feedback $v_{t,i}(\mathbf{x}_t)$
  Construct loss $\boldsymbol{\ell}_{t,i}$ and optimism $\boldsymbol{m}_{t,i}$ via (4.3)
  Update to $\boldsymbol{p}_{t+1,i}$ via (4.2) with learning rate $\varepsilon_{t,i}$ (D.7)
  for $j = 1, \ldots, M$ **do**
    $\mathcal{A}_j$ updates to $x_{t+1,i,j}$ with optimism $v_{t-1,i}(\mathbf{x}_{t-1})$
  end
end

Idea I: online ensemble (two-layer structure)

Idea II: inject correction term to bias towards more stable learners

Smooth Games

Learning And Mining from DatA

Algorithm 2 TV-SMOG (smooth) for the i-th player

Input: step size pool $\mathcal{H}_i^\eta$ (D.18)

Initialize: M instances of OOGD $\mathcal{A}_1, \ldots, \mathcal{A}_M$ with step size $\eta_{i,j} \in \mathcal{H}_i^\eta$

for $t = 1, \ldots, T$ **do**

Receive $x_{t,i,j}$ from $\mathcal{A}_j$

Submit $x_{t,i} = \sum_{j=1}^{M} p_{t,i,j} x_{t,i,j}$

Receive gradient feedback $v_{t,i}(\mathbf{x}_t)$

Construct loss $\boldsymbol{\ell}_{t,i}$ and optimism $\boldsymbol{m}_{t,i}$ via (4.3)

Update to $\boldsymbol{p}_{t+1,i}$ via (4.2) with learning rate $\varepsilon_{t,i}$ (D.7)

end

Injecting a correction term into feedback loss and optimism

$$\ell_{t,i,j} \triangleq \langle v_{t,i}(\mathbf{x}_t), x_{t,i,j} \rangle + \lambda \| x_{t,i,j} - x_{t-1,i,j} \|^2,$$

$$m_{t,i,j} \triangleq \langle v_{t-1,i}(\mathbf{x}_{t-1}), x_{t,i,j} \rangle + \lambda \| x_{t,i,j} - x_{t-1,i,j} \|^2.$$

Purpose: bias towards more stable base-learners to make the cancelation in the dynamic regret feasible

Idea I: online ensemble (two-layer structure)

Idea II: inject correction term to bias towards more stable learners

Result:

$$\mathcal{O}\left(\min\{\sqrt{(1 + V_T + P_T)(1 + P_T)}, 1 + W_T\}\right)$$

Faster rates: fully independent of T

Experiments

LAMDA
Learning And Mining from DatA

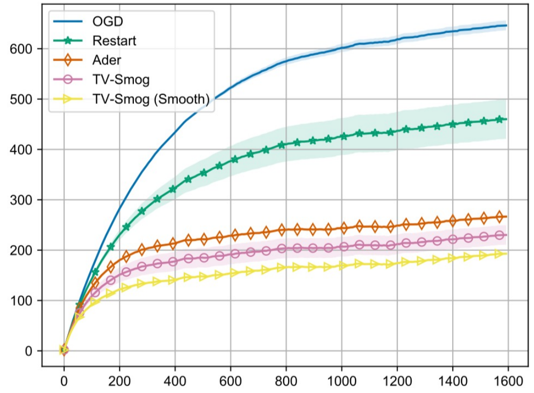

(a) `a1a`

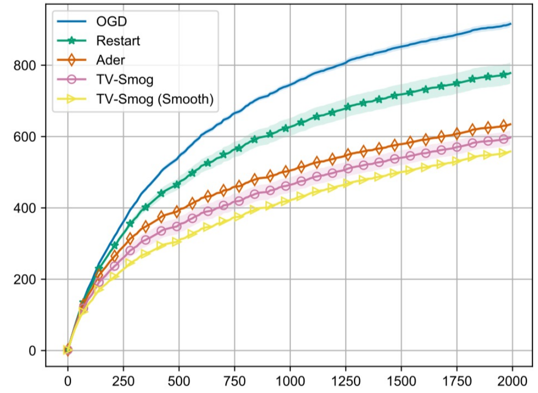

(b) `mushrooms`

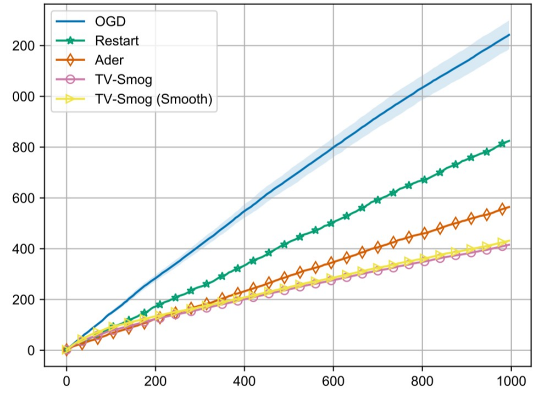

(c) `splice`

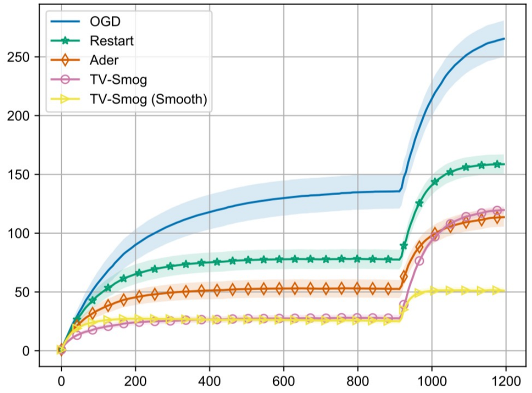

(d) `svmguide3`

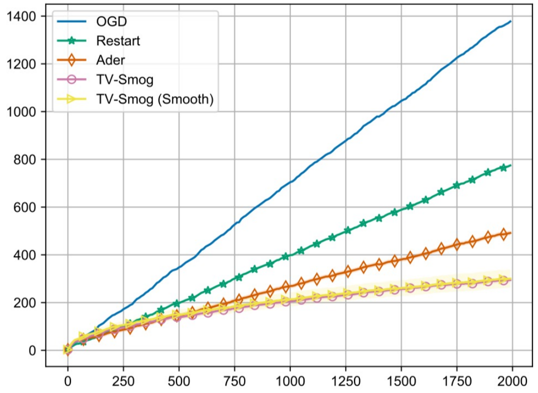

(e) `Cournot`

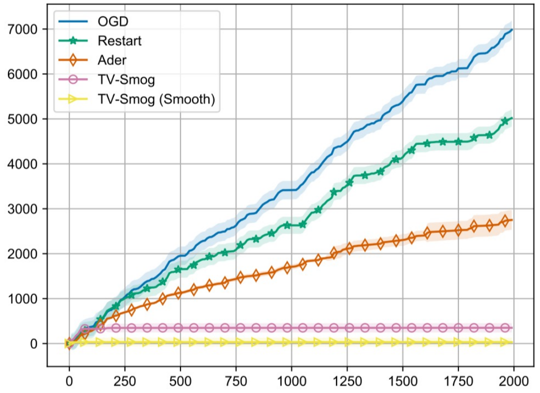

(f) `zero-sum`

- (a), (b), (c), (d): time-varying ℓ_2-regularized logistic regression
- (e): time-varying Cournot competition
- (f): time-varying zero-sum strongly convex-concave games

Summary

- **Problem:** time-varying strongly monotone games
- **Algorithms:** robust online algorithms for non-smooth and smooth games
- **Key ingredients:**
  - Online ensemble framework (suitable meta/base learners, correction, etc.)
  - Strong convexity extracted from strong monotonicity
- **Results:** best-known (fast-rate) tracking error guarantees for this problem

Thanks!