

Adaptive Regret of Convex and Smooth Functions

Lijun Zhang¹ Tie-Yan Liu² Zhi-Hua Zhou¹


¹National Key Laboratory for Novel Software Technology, Nanjing University

²Microsoft Research Asia


The 36th International Conference on Machine Learning (ICML 2019)

■ Online Convex Optimization [Zinkevich, 2003]

- 1: **for** $t = 1, 2, \dots, T$ **do**
- 2: Learner picks a decision $\mathbf{w}_t \in \mathcal{W}$
Adversary chooses a function $f_t(\cdot) : \mathcal{W} \mapsto \mathbb{R}$
- 3: Learner suffers loss $f_t(\mathbf{w}_t)$ and updates \mathbf{w}_t
- 4: **end for**



Learner



An example $(\mathbf{x}_t, y_t) \in \mathbb{R}^d \times \{\pm 1\}$
A loss $f_t(\mathbf{w}) = \max(1 - y_t \mathbf{w}^T \mathbf{x}_t, 0)$



Adversary

■ Cumulative Loss

$$\text{Cumulative Loss} = \sum_{t=1}^T f_t(\mathbf{w}_t)$$

Performance Measure

■ Regret

$$\text{Regret} = \underbrace{\sum_{t=1}^T f_t(\mathbf{w}_t)}_{\text{Cumulative Loss of Online Learner}} - \underbrace{\min_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T f_t(\mathbf{w})}_{\text{Minimal Loss of Offline Learner}}$$

Performance Measure

■ Regret

$$\text{Regret} = \underbrace{\sum_{t=1}^T f_t(\mathbf{w}_t)}_{\text{Cumulative Loss of Online Learner}} - \underbrace{\min_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T f_t(\mathbf{w})}_{\text{Minimal Loss of Offline Learner}}$$

■ Convex Functions [Zinkevich, 2003]

● Online Gradient Descent (OGD)

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

■ Convex and Smooth Functions [Srebro et al., 2010]

● OGD with prior knowledge

$$\text{Regret} = O\left(1 + \sqrt{F_*}\right)$$

$$\text{where } F_* = \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T f_t(\mathbf{w})$$

■ Exp-concave Functions [Hazan et al., 2007]

■ Strongly Convex Functions [Hazan et al., 2007]

■ Regret → Static Regret



$$\begin{aligned}\text{Regret} &= \sum_{t=1}^T f_t(\mathbf{w}_t) - \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T f_t(\mathbf{w}) \\ &= \sum_{t=1}^T f_t(\mathbf{w}_t) - \sum_{t=1}^T f_t(\mathbf{w}_*)\end{aligned}$$

where $\mathbf{w}_* \in \operatorname{argmin}_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T f_t(\mathbf{w})$

- \mathbf{w}_* is reasonably good during T rounds

■ Changing Environments

- Different decisions will be good in different periods
- E.g., recommendation, stock market

The Basic Idea

Minimize the regret over every interval $[r, s]$

$$\text{Regret}([r, s]) = \sum_{t=r}^s f_t(\mathbf{w}_t) - \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=r}^s f_t(\mathbf{w})$$

■ Weakly Adaptive Regret [Hazan and Seshadhri, 2007]

$$\text{WA-Regret}(T) = \max_{[r, s] \subseteq [T]} \text{Regret}([r, s])$$

- The maximal regret over all intervals

■ Strongly Adaptive Regret [Daniely et al., 2015]

$$\text{SA-Regret}(T, \tau) = \max_{[s, s+\tau-1] \subseteq [T]} \text{Regret}([s, s+\tau-1])$$

- The maximal regret over all intervals of length τ

- Convex Functions [Jun et al., 2017]

$$\text{Regret}([r, s]) = O\left(\sqrt{(s-r)\log s}\right)$$

$$\Rightarrow \text{SA-Regret}(T, \tau) = O\left(\sqrt{\tau \log T}\right)$$

- Exp-concave Functions [Hazan and Seshadhri, 2007]

- Strongly Convex Functions [Zhang et al., 2018]

- Convex Functions [Jun et al., 2017]

$$\text{Regret}([r, s]) = O\left(\sqrt{(s-r)\log s}\right)$$

$$\Rightarrow \text{SA-Regret}(T, \tau) = O\left(\sqrt{\tau \log T}\right)$$

- Exp-concave Functions [Hazan and Seshadhri, 2007]
- Strongly Convex Functions [Zhang et al., 2018]

Question

Can **smoothness** be exploited to boost the adaptive regret?

■ Convex and Smooth Functions

$$\text{Regret}([r, s]) = O\left(\sqrt{\left(\sum_{t=r}^s f_t(\mathbf{w})\right) \log s \cdot \log(s-r)}\right)$$

- Become tighter when $\sum_{t=r}^s f_t(\mathbf{w})$ is small

■ Convex Functions [Jun et al., 2017]

$$\text{Regret}([r, s]) = O\left(\sqrt{(s-r) \log s}\right)$$

■ Convex and Smooth Functions

$$\text{Regret}([r, s]) = O\left(\sqrt{\left(\sum_{t=r}^s f_t(\mathbf{w})\right)} \log s \cdot \log(s-r)\right)$$

- Become tighter when $\sum_{t=r}^s f_t(\mathbf{w})$ is small

■ Convex Functions [Jun et al., 2017]

$$\text{Regret}([r, s]) = O\left(\sqrt{(s-r) \log s}\right)$$

■ Convex and Smooth Functions

$$\text{Regret}([r, s]) = O\left(\sqrt{\left(\sum_{t=r}^s f_t(\mathbf{w})\right)} \log \sum_{t=1}^s f_t(\mathbf{w}) \cdot \log \sum_{t=r}^s f_t(\mathbf{w})\right)$$

- Fully problem-dependent

The Algorithm

■ An Expert-algorithm

- Scale-free online gradient descent [Orabona and Pál, 2018]
- Can exploit smoothness **automatically**

■ A Set of Intervals

- Compact geometric covering intervals [Daniely et al., 2015]

t	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	...		
\mathcal{C}_0	[]	[]	[]	[]	[]	[]	[]	[]	[]	[]	...
\mathcal{C}_1		[]		[]		[]		[]		[]	...
\mathcal{C}_2			[]				[]					...
\mathcal{C}_3								[]	...
\mathcal{C}_4																	[...

■ A Meta-algorithm

- AdaNormalHedge [Luo and Schapire, 2015]
- Attain a **small-loss** regret and support **sleeping** experts

Thanks!

Welcome to Our Poster @ Pacific Ballroom #161.



Daniely, A., Gonen, A., and Shalev-Shwartz, S. (2015).

Strongly adaptive online learning.

In Proceedings of the 32nd International Conference on Machine Learning, pages 1405–1411.



Hazan, E., Agarwal, A., and Kale, S. (2007).

Logarithmic regret algorithms for online convex optimization.

Machine Learning, 69(2-3):169–192.



Hazan, E. and Seshadhri, C. (2007).

Adaptive algorithms for online decision problems.

Electronic Colloquium on Computational Complexity, 88.



Jun, K.-S., Orabona, F., Wright, S., and Willett, R. (2017).

Improved strongly adaptive online learning using coin betting.

In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, pages 943–951.



Luo, H. and Schapire, R. E. (2015).

Achieving all with no parameters: Adanormalhedge.

In Proceedings of The 28th Conference on Learning Theory, pages 1286–1304.



Orabona, F. and Pál, D. (2018).
Scale-free online learning.
Theoretical Computer Science, 716:50–69.



Srebro, N., Sridharan, K., and Tewari, A. (2010).
Smoothness, low-noise and fast rates.
In Advances in Neural Information Processing Systems 23, pages 2199–2207.



Zhang, L., Yang, T., Jin, R., and Zhou, Z.-H. (2018).
Dynamic regret of strongly adaptive methods.
In Proceedings of the 35th International Conference on Machine Learning.



Zinkevich, M. (2003).
Online convex programming and generalized infinitesimal gradient ascent.
In Proceedings of the 20th International Conference on Machine Learning, pages 928–936.