Smoothed Online Learning

Lijun Zhang

Nanjing University, China

CSIAM-BDAI 2021

Outline

Learning And Mining from DatA

Outline

Background

- Perceptron [Rosenblatt, 1958]
- Prediction with Expert Advice [Littlestone and Warmuth, 1989]
- Online Convex Optimization [Zinkevich, 2003]

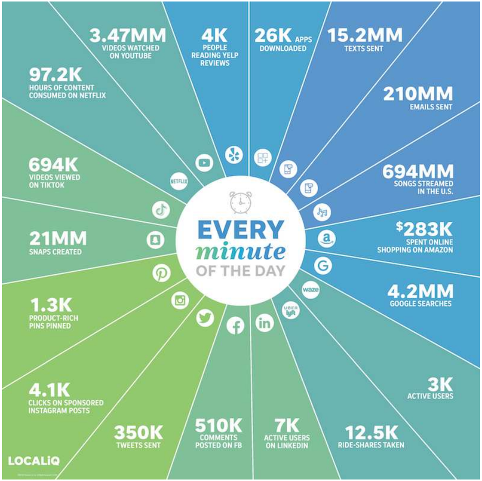

https://localiq.com/blog/what-happens-in-an-internet-minute-2021/

Online Learning [Shalev-Shwartz, 2011]

Online learning is the process of answering a sequence of questions given (maybe partial) knowledge of answers to previous questions and possibly additional information.

Online Learning [Shalev-Shwartz, 2011]

Online learning is the process of answering a sequence of questions given (maybe partial) knowledge of answers to previous questions and possibly additional information.

The Learning Procedure

1: **for** $t = 1, 2, \ldots, T$ **do**

4: **end for**

Online Learning [Shalev-Shwartz, 2011]

Online learning is the process of answering a sequence of questions given (maybe partial) knowledge of answers to previous questions and possibly additional information.

The Learning Procedure

for $t = 1, 2, \ldots, T$ **do**

  Learner picks a decision $\mathbf{w}_t \in \mathcal{W}$

  Adversary chooses a function $f_t(\cdot)$

end for

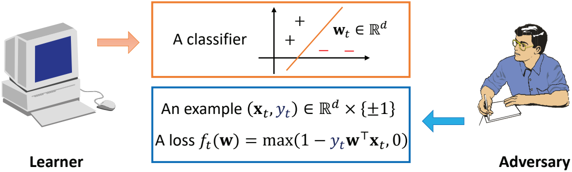

Online Learning [Shalev-Shwartz, 2011]

Online learning is the process of answering a sequence of questions given (maybe partial) knowledge of answers to previous questions and possibly additional information.

The Learning Procedure

for $t = 1, 2, \ldots, T$ **do**

  Learner picks a decision $\mathbf{w}_t \in \mathcal{W}$

  Adversary chooses a function $f_t(\cdot)$

  Learner suffers loss $f_t(\mathbf{w}_t)$ and updates $\mathbf{w}_t$

end for

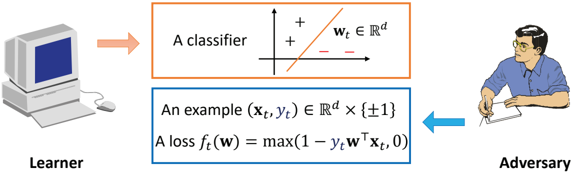

Online Learning [Shalev-Shwartz, 2011]

Online learning is the process of answering a sequence of questions given (maybe partial) knowledge of answers to previous questions and possibly additional information.

The Learning Procedure

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

Cumulative Loss

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

Outline

1. Online Learning
   - Regret
   - Dynamic Regret
2. Smoothed Online Learning
   - Competitive Ratio
   - Dynamic Regret with Switching Cost
3. Conclusion

LAMDA Learning And Mining from DatA

Regret

Cumulative Loss

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

Regret

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

Regret

Cumulative Loss

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

Regret

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

Hannan Consistent

$$\sum_{t=1}^{T} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=1}^{T} f_t(\mathbf{w}) = o(T), \text{ with probability } 1$$

Online Convex Optimization

- Convex Functions [Zinkevich, 2003]
  - Online Gradient Descent (OGD)

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

- Strongly Convex Functions [Hazan et al., 2007]
  - Online Gradient Descent (OGD)

$$\text{Regret}(T) = O\left(\log T\right)$$

- Exponentially Concave Functions [Hazan et al., 2007]
  - Online Newton Step (ONS)

$$\text{Regret}(T) = O\left(d \log T\right)$$

Outline

Learning And Mining from DatA

Learning in Changing Environments

Regret → *Static* Regret

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

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

- One of the decision is reasonably good during T rounds

Learning in Changing Environments

Regret → *Static Regret*

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

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

- One of the decision is reasonably good during T rounds

Changing Environments

Different decisions will be good in different periods

- Recommendation: the interests of a user could change
- Stock market: the best stock changes over time

Dynamic Regret

$$\text{D-Regret}(\mathbf{u}_1, \ldots, \mathbf{u}_T) = \sum_{t=1}^{T} f_t(\mathbf{w}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}_t)$$

where $\mathbf{u}_1, \ldots, \mathbf{u}_T \in \mathcal{W}$ is an arbitrary comparator sequence

Dynamic Regret

$$\text{D-Regret}(\mathbf{u}_1, \ldots, \mathbf{u}_T) = \sum_{t=1}^{T} f_t(\mathbf{w}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}_t)$$

where $\mathbf{u}_1, \ldots, \mathbf{u}_T \in \mathcal{W}$ is an arbitrary comparator sequence

- Online Gradient Descent (OGD) [Zinkevich, 2003]

$$\text{D-Regret}(\mathbf{u}_1, \ldots, \mathbf{u}_T) = O\left(\sqrt{T} \cdot (1 + P_T)\right)$$

where $P_T = \sum_{t=1}^{T} \|\mathbf{u}_{t+1} - \mathbf{u}_t\|_2$

Dynamic Regret

$$\text{D-Regret}(\mathbf{u}_1, \ldots, \mathbf{u}_T) = \sum_{t=1}^{T} f_t(\mathbf{w}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}_t)$$

where $\mathbf{u}_1, \ldots, \mathbf{u}_T \in \mathcal{W}$ is an arbitrary comparator sequence

- Online Gradient Descent (OGD) [Zinkevich, 2003]

$$\text{D-Regret}(\mathbf{u}_1, \ldots, \mathbf{u}_T) = O\left(\sqrt{T} \cdot (1 + P_T)\right)$$

where $P_T = \sum_{t=1}^{T} \|\mathbf{u}_{t+1} - \mathbf{u}_t\|_2$

- The First Lower Bound [Zhang et al., 2018]

$$\text{D-Regret}(\mathbf{u}_1, \ldots, \mathbf{u}_T) = \Omega\left(\sqrt{T} \cdot \sqrt{1 + P_T}\right)$$

- An Optimal Algorithm—Ader [Zhang et al., 2018]

$$\text{D-Regret}(\mathbf{u}_1, \ldots, \mathbf{u}_T) = O\left(\sqrt{T} \cdot \sqrt{1 + P_T}\right)$$

Outline

Smoothed Online Learning

The Learning Procedure

1: **for** $t = 1, 2, \ldots, T$ **do**

2: Learner picks a decision $\mathbf{w}_t \in \mathcal{W}$

Adversary chooses a function $f_t(\cdot)$

3: Learner suffers a hitting cost $f_t(\mathbf{w}_t)$,

and a switching cost $m(\mathbf{w}_t, \mathbf{w}_{t-1})$

4: **end for**

Smoothed Online Learning

The Learning Procedure

1: **for** $t = 1, 2, \ldots, T$ **do**
2: Learner picks a decision $\mathbf{w}_t \in \mathcal{W}$
Adversary chooses a function $f_t(\cdot)$
3: Learner suffers a hitting cost $f_t(\mathbf{w}_t)$,
and a switching cost $m(\mathbf{w}_t, \mathbf{w}_{t-1})$
4: **end for**

- Applications
  - Stock market: the transaction fee
  - Data center: the wear-and-tear cost
  - Store relocation: the decoration cost

LAMDA Learning And Mining from DatA

Smoothed Online Learning

The Learning Procedure

1: **for** $t = 1, 2, \ldots, T$ **do**
2: Learner picks a decision $\mathbf{w}_t \in \mathcal{W}$
Adversary chooses a function $f_t(\cdot)$
3: Learner suffers a hitting cost $f_t(\mathbf{w}_t)$,
and a switching cost $m(\mathbf{w}_t, \mathbf{w}_{t-1})$
4: **end for**

Cumulative Loss (Hitting Cost + Switching Cost)

$$\text{Cumulative Loss} = \sum_{t=1}^{T} f_t(\mathbf{w}_t) + m(\mathbf{w}_t, \mathbf{w}_{t-1})$$

Smoothed Online Learning

The Learning Procedure

1: **for** $t = 1, 2, \ldots, T$ **do**
2: Adversary chooses a function $f_t(\cdot)$,
then Learner picks a decision $\mathbf{w}_t \in \mathcal{W}$
3: Learner suffers a hitting cost $f_t(\mathbf{w}_t)$,
and a switching cost $m(\mathbf{w}_t, \mathbf{w}_{t-1})$
4: **end for**

The Lookahead Setting

The problem is nontrivial even when the learner can observe $f_t(\cdot)$ before deciding $\mathbf{w}_t$.

Outline

Competitive Ratio

$$\frac{\sum_{t=1}^{T}\left(f_t(\mathbf{w}_t) + m(\mathbf{w}_t, \mathbf{w}_{t-1})\right)}{\min_{\mathbf{u}_0, \mathbf{u}_1, \ldots, \mathbf{u}_T \in \mathcal{X}} \sum_{t=1}^{T}\left(f_t(\mathbf{u}_t) + m(\mathbf{u}_t, \mathbf{u}_{t-1})\right)}$$

Competitive Ratio

$$\frac{\sum_{t=1}^{T}\left(f_t(\mathbf{w}_t)+m(\mathbf{w}_t,\mathbf{w}_{t-1})\right)}{\min_{\mathbf{u}_0,\mathbf{u}_1,\ldots,\mathbf{u}_T\in\mathcal{X}}\sum_{t=1}^{T}\left(f_t(\mathbf{u}_t)+m(\mathbf{u}_t,\mathbf{u}_{t-1})\right)}$$

- Convex Body Chasing (CBC)
  - Select one point from convex bodies $\mathcal{W}_1,\ldots,\mathcal{W}_T\subseteq\mathbb{R}^d$

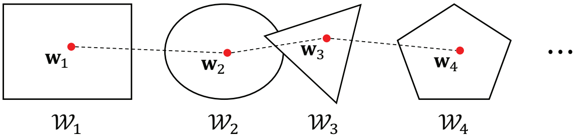

  - Minimize the total movement $\sum\|\mathbf{w}_t-\mathbf{w}_{t-1}\|$

Competitive Ratio

$$\frac{\sum_{t=1}^{T}\left(f_t(\mathbf{w}_t)+m(\mathbf{w}_t,\mathbf{w}_{t-1})\right)}{\min_{\mathbf{u}_0,\mathbf{u}_1,\ldots,\mathbf{u}_T\in\mathcal{X}}\sum_{t=1}^{T}\left(f_t(\mathbf{u}_t)+m(\mathbf{u}_t,\mathbf{u}_{t-1})\right)}$$

- Convex Body Chasing (CBC)
  - Select one point from convex bodies $\mathcal{W}_1,\ldots,\mathcal{W}_T\subseteq\mathbb{R}^d$

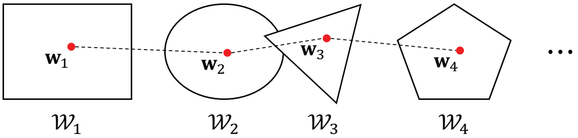

  - Minimize the total movement $\sum\|\mathbf{w}_t-\mathbf{w}_{t-1}\|$
  - Lower bound: $\Omega(\sqrt{d})$ [Friedman and Linial, 1993]
  - Upper bound: $O(\min(d,\sqrt{d\log T}))$ [Argue et al., 2020, Sellke, 2020]

Competitive Ratio

$$\frac{\sum_{t=1}^{T}\left(f_t(\mathbf{w}_t)+m(\mathbf{w}_t,\mathbf{w}_{t-1})\right)}{\min_{\mathbf{u}_0,\mathbf{u}_1,\ldots,\mathbf{u}_T\in\mathcal{X}}\sum_{t=1}^{T}\left(f_t(\mathbf{u}_t)+m(\mathbf{u}_t,\mathbf{u}_{t-1})\right)}$$

Research on Competitive Ratio

Identify sufficient conditions and develop algorithms for

dimension-free competitive ratio in lookahead setting

- Polyhedral functions [Chen et al., 2018, Lin et al., 2020]
- Quadratic growth functions [Goel et al., 2019, Lin et al., 2020]
- Strongly convex functions [Goel et al., 2019]

The function can not be too flat.

Polyhedral Functions

Definition

A function $f(\cdot) : \mathcal{W} \mapsto \mathbb{R}$ with minimizer $\mathbf{v}$ is α-polyhedral if

$$f(\mathbf{w}) - f(\mathbf{v}) \geq \alpha \|\mathbf{w} - \mathbf{v}\|, \ \forall \mathbf{w} \in \mathcal{W}.$$

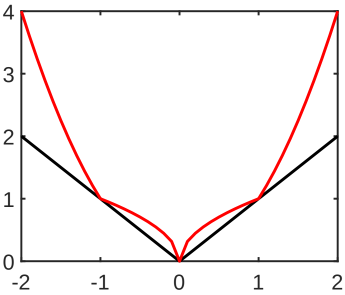

Could be non-convex

- Stochastic Network Optimization [Huang and Neely, 2011]
- Geographical Load Balancing [Lin et al., 2012]

State-of-the-Art

- α-polyhedral and convex functions [Chen et al., 2018]
  - Online balanced descent: balancing the two costs by iteratively projections
  - Competitive ratio: $3 + \frac{8}{\alpha}$

State-of-the-Art

- α-polyhedral and convex functions [Chen et al., 2018]
  - Online balanced descent: balancing the two costs by iteratively projections
  - Competitive ratio: $3 + \frac{8}{\alpha}$
- α-polyhedral functions [Lin et al., 2020]
  - The naive approach which *ignores* the switching cost

$$\mathbf{w}_t = \operatorname*{argmin}_{\mathbf{w} \in \mathcal{W}} f_t(\mathbf{w})$$

  - Competitive ratio: $1 + \frac{2}{\alpha}$

State-of-the-Art

- α-polyhedral and convex functions [Chen et al., 2018]
  - Online balanced descent: balancing the two costs by iteratively projections
  - Competitive ratio: $3 + \frac{8}{\alpha}$
- α-polyhedral functions [Lin et al., 2020]
  - The naive approach which *ignores* the switching cost

$$\mathbf{w}_t = \operatorname*{argmin}_{\mathbf{w} \in \mathcal{W}} f_t(\mathbf{w})$$

  - Competitive ratio: $1 + \frac{2}{\alpha}$
- α-polyhedral functions [Zhang et al., 2021]
  - Competitive ratio: $\max(1, \frac{2}{\alpha})$ for the naive approach

State-of-the-Art

- α-polyhedral and convex functions [Chen et al., 2018]
  - Online balanced descent: balancing the two costs by iteratively projections
  - Competitive ratio: $3 + \frac{8}{\alpha}$
- α-polyhedral functions [Lin et al., 2020]
  - The naive approach which *ignores* the switching cost

$$\mathbf{w}_t = \operatorname*{argmin}_{\mathbf{w} \in \mathcal{W}} f_t(\mathbf{w})$$

  - Competitive ratio: $1 + \frac{2}{\alpha}$
- α-polyhedral functions [Zhang et al., 2021]
  - Competitive ratio: $\max(1, \frac{2}{\alpha})$ for the naive approach

A counterintuitive fact

It seems safe to ignore the switching cost.

Quadratic Growth Functions

Definition

A function $f(\cdot) : \mathcal{W} \mapsto \mathbb{R}$ with minimizer $\mathbf{v}$ is λ-quadratic growth if

$$f(\mathbf{w}) - f(\mathbf{v}) \geq \frac{\lambda}{2}\|\mathbf{w} - \mathbf{v}\|^2, \ \forall \mathbf{w} \in \mathcal{W}.$$

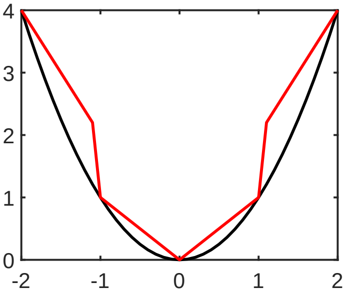

Could be non-convex

- A sufficient condition for linear convergence [Drusvyatskiy and Lewis, 2018, Necoara et al., 2019]
- Weaker than strong convexity [Hazan and Kale, 2011]

State-of-the-Art

- λ-quadratic growth functions [Lin et al., 2020]
  - The naive approach which *ignores* the switching cost

$$\mathbf{w}_t = \operatorname*{argmin}_{\mathbf{w} \in \mathcal{W}} f_t(\mathbf{w})$$

  - Competitive ratio: $\max(1 + \frac{6}{\lambda}, 4)$

Learning And Mining from DatA

State-of-the-Art

- λ-quadratic growth functions [Lin et al., 2020]
  - The naive approach which *ignores* the switching cost

$$\mathbf{w}_t = \operatorname*{argmin}_{\mathbf{w} \in \mathcal{W}} f_t(\mathbf{w})$$

  - Competitive ratio: $\max(1 + \frac{6}{\lambda}, 4)$

- λ-quadratic growth functions [Zhang et al., 2021]
  - Competitive ratio: $1 + \frac{4}{\lambda}$ for the naive approach
  - The order is optimal [Goel et al., 2019]

LAMDA Learning And Mining from DatA

State-of-the-Art

- λ-quadratic growth functions [Lin et al., 2020]
  - The naive approach which *ignores* the switching cost

$$\mathbf{w}_t = \operatorname*{argmin}_{\mathbf{w} \in \mathcal{W}} f_t(\mathbf{w})$$

  - Competitive ratio: $\max(1 + \frac{6}{\lambda}, 4)$

- λ-quadratic growth functions [Zhang et al., 2021]
  - Competitive ratio: $1 + \frac{4}{\lambda}$ for the naive approach
  - The order is optimal [Goel et al., 2019]

Discussions

- It is also unclear how to utilize the switching cost.
- But if convexity is also present, the switching cost plays an important role.

State-of-the-Art

- λ-quadratic growth and quasiconvex functions [Goel et al., 2019]
  - Greedy Online Balanced Descent (OBD)
  - Competitive ratio: $O(\frac{1}{\sqrt{\lambda}})$ as $\lambda \to 0$

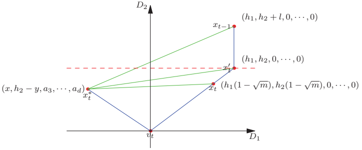

State-of-the-Art

- λ-quadratic growth and quasiconvex functions [Goel et al., 2019]
  - Greedy Online Balanced Descent (OBD)
  - Competitive ratio: $O(\frac{1}{\sqrt{\lambda}})$ as $\lambda \to 0$

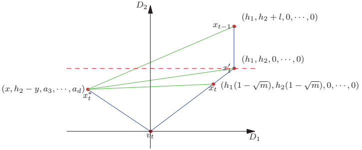

- λ-quadratic growth and convex functions [Zhang et al., 2021]
  - The greedy approach which minimizes the weighted sum

$$\mathbf{w}_t = \underset{\mathbf{w} \in \mathcal{W}}{\operatorname{argmin}} \left(f_t(\mathbf{w}) + \frac{\gamma}{2} \|\mathbf{w} - \mathbf{w}_{t-1}\|^2 \right)$$

  - Competitive ratio: $1 + \frac{2}{\sqrt{\lambda}}$
  - The order is optimal [Goel et al., 2019]

Outline

Dynamic Regret with Switching Cost

$$\sum_{t=1}^{T}\left(f_t(\mathbf{w}_t)+m(\mathbf{w}_t,\mathbf{w}_{t-1})\right)-\sum_{t=1}^{T}\left(f_t(\mathbf{u}_t)+m(\mathbf{u}_t,\mathbf{u}_{t-1})\right)$$

where $\mathbf{u}_1,\ldots,\mathbf{u}_T\in\mathcal{W}$ is an arbitrary comparator sequence

- The standard setting
- The lookahead setting

Dynamic Regret with Switching Cost

$$\sum_{t=1}^{T}\left(f_t(\mathbf{w}_t)+m(\mathbf{w}_t,\mathbf{w}_{t-1})\right)-\sum_{t=1}^{T}\left(f_t(\mathbf{u}_t)+m(\mathbf{u}_t,\mathbf{u}_{t-1})\right)$$

where $\mathbf{u}_1,\ldots,\mathbf{u}_T\in\mathcal{W}$ is an arbitrary comparator sequence

- The standard setting
- The lookahead setting

Dynamic Regret

$$\text{D-Regret}(\mathbf{u}_1,\ldots,\mathbf{u}_T)=\sum_{t=1}^{T}f_t(\mathbf{w}_t)-\sum_{t=1}^{T}f_t(\mathbf{u}_t)$$

where $\mathbf{u}_1,\ldots,\mathbf{u}_T\in\mathcal{W}$ is an arbitrary comparator sequence

The Standard Setting

Assumptions

1. All the functions f_t's are convex over their domain $\mathcal{W}$
2. The gradients of all functions are bounded by G
3. The diameter of the domain $\mathcal{W}$ is bounded by D

The Standard Setting

Assumptions

1. All the functions f_t's are convex over their domain $\mathcal{W}$
2. The gradients of all functions are bounded by G
3. The diameter of the domain $\mathcal{W}$ is bounded by D

- Smoothed Ader (SAder) [Zhang et al., 2021]

$$\sum_{t=1}^{T}\Big(f_t(\mathbf{w}_t)+\|\mathbf{w}_t-\mathbf{w}_{t-1}\|\Big)-\sum_{t=1}^{T}f_t(\mathbf{u}_t)=O\left(\sqrt{T}\cdot\sqrt{1+P_T}\right)$$

where $P_T=\sum_{t=1}^{T}\|\mathbf{u}_{t+1}-\mathbf{u}_t\|_2$

- Optimal according to the lower bound of dynamic regret [Zhang et al., 2018]
- The switching cost does not make the problem much harder, although we need to modify the algorithm

The Lookahead Setting

Assumptions

1. All the functions f_t's are convex over their domain $\mathcal{W}$
2. ~~The gradients of all functions are bounded by G~~
3. The diameter of the domain $\mathcal{W}$ is bounded by D

The Lookahead Setting

Assumptions

1. All the functions f_t's are convex over their domain $\mathcal{W}$
2. ~~The gradients of all functions are bounded by G~~
3. The diameter of the domain $\mathcal{W}$ is bounded by D

- Lookahead SAder [Zhang et al., 2021]

$$\sum_{t=1}^{T}\left(f_t(\mathbf{w}_t)+\|\mathbf{w}_t-\mathbf{w}_{t-1}\|\right)-\sum_{t=1}^{T}f_t(\mathbf{u}_t)=O\left(\sqrt{T}\cdot\sqrt{1+P_T}\right)$$

where $P_T=\sum_{t=1}^{T}\|\mathbf{u}_{t+1}-\mathbf{u}_t\|_2$

The Lookahead Setting

Assumptions

1. All the functions f_t's are convex over their domain $\mathcal{W}$
2. ~~The gradients of all functions are bounded by G~~
3. The diameter of the domain $\mathcal{W}$ is bounded by D

- Lookahead SAder [Zhang et al., 2021]

$$\sum_{t=1}^{T} \left(f_t(\mathbf{w}_t) + \|\mathbf{w}_t - \mathbf{w}_{t-1}\| \right) - \sum_{t=1}^{T} f_t(\mathbf{u}_t) = O\left(\sqrt{T} \cdot \sqrt{1 + P_T} \right)$$

where $P_T = \sum_{t=1}^{T} \|\mathbf{u}_{t+1} - \mathbf{u}_t\|_2$

- The first $\Omega(\sqrt{T} \cdot \sqrt{1 + P_T})$ lower bound for lookahead setting [Zhang et al., 2021]
  - Our lookahead SAder is optimal

The Lookahead Setting

Assumptions

1. All the functions f_t's are convex over their domain $\mathcal{W}$
2. ~~The gradients of all functions are bounded by G~~
3. The diameter of the domain $\mathcal{W}$ is bounded by D

- Lookahead SAder [Zhang et al., 2021]

$$\sum_{t=1}^{T}\left(f_t(\mathbf{w}_t)+\|\mathbf{w}_t-\mathbf{w}_{t-1}\|\right)-\sum_{t=1}^{T}f_t(\mathbf{u}_t)=O\left(\sqrt{T}\cdot\sqrt{1+P_T}\right)$$

where $P_T = \sum_{t=1}^{T}\|\mathbf{u}_{t+1}-\mathbf{u}_t\|_2$

- An $O(\sqrt{T}\cdot(1+P_T))$ Upper bound [Chen et al., 2018]
  - Suboptimal

Outline

Conclusion and Future Work

- Smoothed Online Learning
  - Minimize the sum of hitting cost and switching cost
  - Competitive ratio and dynamic regret with switching cost

Conclusion and Future Work

- Smoothed Online Learning
  - Minimize the sum of hitting cost and switching cost
  - Competitive ratio and dynamic regret with switching cost
- Future Work
  - Lower bounds for Competitive ratio
  - Improve dynamic regret with switching cost
  - Minimize the two measures simultaneously
  - The relation with continual learning

$$\sum_{t=1}^{T} \underbrace{f_t(\mathbf{w}_t)}_{\text{Perform well on each task}} + \underbrace{\|\mathbf{w}_t - \mathbf{w}_{t-1}\|}_{\text{Avoid catastrophic forgetting}}$$

Reference I

Thanks!

Argue, C., Gupta, A., Guruganesh, G., and Tang, Z. (2020).
Chasing convex bodies with linear competitive ratio.
In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1519–1524.

Chen, N., Goel, G., and Wierman, A. (2018).
Smoothed online convex optimization in high dimensions via online balanced descent.
In *Proceedings of the 31st Conference on Learning Theory*, pages 1574–1594.

Drusvyatskiy, D. and Lewis, A. S. (2018).
Error bounds, quadratic growth, and linear convergence of proximal methods.
Mathematics of Operations Research, 43(3):919–948.

Friedman, J. and Linial, N. (1993).
On convex body chasing.
Discrete & Computational Geometry, 9:293–321.

Goel, G., Lin, Y., Sun, H., and Wierman, A. (2019).
Beyond online balanced descent: An optimal algorithm for smoothed online optimization.
In *Advances in Neural Information Processing Systems 32*, pages 1875–1885.

Hazan, E., Agarwal, A., and Kale, S. (2007).
Logarithmic regret algorithms for online convex optimization.
Machine Learning, 69(2-3):169–192.

Reference II

Hazan, E. and Kale, S. (2011).
Beyond the regret minimization barrier: an optimal algorithm for stochastic strongly-convex optimization.
In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 421–436.

Huang, L. and Neely, M. J. (2011).
Delay reduction via lagrange multipliers in stochastic network optimization.
IEEE Transactions on Automatic Control, 56(4):842–857.

Lin, M., Liu, Z., Wierman, A., and Andrew, L. L. (2012).
Online algorithms for geographical load balancing.
In *Proceedings of the 2012 International Green Computing Conference*, pages 1–10.

Lin, Y., Goel, G., and Wierman, A. (2020).
Online optimization with predictions and non-convex losses.
Proceedings of the ACM on Measurement and Analysis of Computing Systems, 4(1):18:1–18:32.

Littlestone, N. and Warmuth, M. K. (1989).
The weighted majority algorithm.
In *Proceedings of the 30th Annual IEEE Symposium on Foundations of Computer Science*, pages 256–261.

Necoara, I., Nesterov, Y., and Glineur, F. (2019).
Linear convergence of first order methods for non-strongly convex optimization.
Mathematical Programming, 175(1):69–107.

Rosenblatt, F. (1958).
The perceptron: a probabilistic model for information storage and organization in the brain.
Psychological Review, 65:386–407.

Reference III

Sellke, M. (2020).
Chasing convex bodies optimally.
In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1509–1518.

Shalev-Shwartz, S. (2011).
Online learning and online convex optimization.
Foundations and Trends in Machine Learning, 4(2):107–194.

Zhang, L., Jiang, W., Lu, S., and Yang, T. (2021).
Revisiting smoothed online learning.
ArXiv e-prints, 2102.06933.

Zhang, L., Lu, S., and Zhou, Z.-H. (2018).
Adaptive online learning in dynamic environments.
In *Advances in Neural Information Processing Systems 31*, pages 1323–1333.

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

LAMDA Learning And Mining from DatA