Online Learning in Changing Environments

Lijun Zhang

Nanjing University, China

Early Career Spotlight Talk, the 29th International Joint Conference on Artificial Intelligence (IJCAI 2020)

Outline

Outline

Online Learning

- The Learning Process [Shalev-Shwartz, 2011]

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

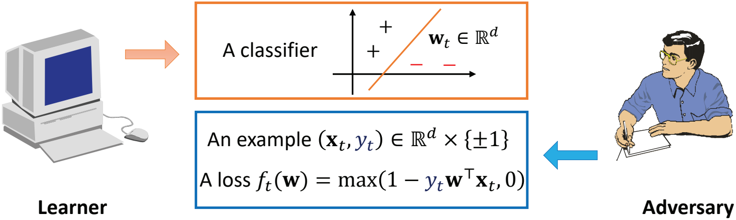

- Cumulative Loss

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

Performance Measure

- Regret

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

Performance Measure

- Regret

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

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

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

- 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)

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

Learning in Changing Environments

Regret → *Static* Regret

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

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

Outline

Definitions

- General Dynamic Regret [Zinkevich, 2003]

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

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

- Worst-case Dynamic Regret [Besbes et al., 2015]

$$\begin{aligned}\text{D-Regret}(\mathbf{w}_1^*, \ldots, \mathbf{w}_T^*) &= \sum_{t=1}^{T} f_t(\mathbf{w}_t) - \sum_{t=1}^{T} f_t(\mathbf{w}_t^*) \\ &= \sum_{t=1}^{T} f_t(\mathbf{w}_t) - \sum_{t=1}^{T} \min_{\mathbf{w}\in\mathcal{W}} f_t(\mathbf{w})\end{aligned}$$

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

Definitions

- General Dynamic Regret [Zinkevich, 2003]

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

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

- Worst-case Dynamic Regret [Besbes et al., 2015]

$$\begin{aligned}\text{D-Regret}(\mathbf{w}_1^*, \ldots, \mathbf{w}_T^*) &= \sum_{t=1}^{T} f_t(\mathbf{w}_t) - \sum_{t=1}^{T} f_t(\mathbf{w}_t^*) \\ &= \sum_{t=1}^{T} f_t(\mathbf{w}_t) - \sum_{t=1}^{T} \min_{\mathbf{w}\in\mathcal{W}} f_t(\mathbf{w})\end{aligned}$$

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

The Challenge

Sublinear Dynamic Regret is Impossible in General!

Worst-case Dynamic Regret

- Functional Variation [Besbes et al., 2015]
  - Suppose

$$F_T = \sum_{t=1}^{T} \sup_{\mathbf{w} \in \mathcal{W}} |f_{t+1}(\mathbf{w}) - f_t(\mathbf{w})| \leq V_T$$

  - Restarted Online Gradient Descent

$$\text{D-Regret}(\mathbf{w}_1^*, \ldots, \mathbf{w}_T^*) = \begin{cases} O\left(V_T^{1/3} T^{2/3}\right), \text{Convex} \\ O\left(\log T \sqrt{V_T T}\right), \text{Strongly Convex} \end{cases}$$

Worst-case Dynamic Regret

- Functional Variation [Besbes et al., 2015]
  - Suppose

$$F_T = \sum_{t=1}^{T} \sup_{\mathbf{w}\in\mathcal{W}} |f_{t+1}(\mathbf{w}) - f_t(\mathbf{w})| \leq V_T$$

  - Restarted Online Gradient Descent

$$\text{D-Regret}(\mathbf{w}_1^*, \ldots, \mathbf{w}_T^*) = \begin{cases} O\left(V_T^{1/3} T^{2/3}\right), & \text{Convex} \\ O\left(\log T \sqrt{V_T T}\right), & \text{Strongly Convex} \end{cases}$$

- Path-length [Mokhtari et al., 2016]
  - Strongly Convex and Smooth Functions

$$\text{D-Regret}(\mathbf{w}_1^*, \ldots, \mathbf{w}_T^*) = O\left(P_T^*\right)$$

where

$$P_T^* = \sum_{t=2}^{T} \|\mathbf{w}_t^* - \mathbf{w}_{t-1}^*\|_2$$

Learning And Mining from DatA

Our Contributions

- Path-length [Yang et al., 2016]

$$\text{D-Regret}(\mathbf{w}_1^*, \ldots, \mathbf{w}_T^*) = \begin{cases} O\left(P_T^*\right), \text{Convex and Smooth} \\ O\left(\sqrt{TB_T}\right), \text{Convex}, P_T^* \leq B_T \end{cases}$$

- Squared Path-length [Zhang et al., 2017]

$$S_T^* = \sum_{t=1}^{T} \|\mathbf{w}_{t+1}^* - \mathbf{w}_t^*\|_2^2$$

  - Strongly convex and smooth functions
  - Semi-strongly convex and smooth functions
  - Self-concordant functions

$$\text{D-Regret}(\mathbf{w}_1^*, \ldots, \mathbf{w}_T^*) = O\left(\min(P_T^*, S_T^*)\right)$$

Our Contributions

- Path-length [Yang et al., 2016]

$$\text{D-Regret}(\mathbf{w}_1^*, \ldots, \mathbf{w}_T^*) = \begin{cases} O\left(P_T^*\right), & \text{Convex and Smooth} \\ O\left(\sqrt{TB_T}\right), & \text{Convex}, P_T^* \leq B_T \end{cases}$$

- Squared Path-length [Zhang et al., 2017]

$$S_T^* = \sum_{t=1}^{T} \|\mathbf{w}_{t+1}^* - \mathbf{w}_t^*\|_2^2$$

  - Strongly convex and smooth functions
  - Semi-strongly convex and smooth functions
  - Self-concordant functions

$$\text{D-Regret}(\mathbf{w}_1^*, \ldots, \mathbf{w}_T^*) = O\left(\min(P_T^*, S_T^*)\right)$$

However, the worst-case dynamic regret is too pessimistic and may overfit in stationary environments.

General Dynamic Regret

- General 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

- Static Regret:

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

- Worst-case Dynamic Regret:

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

General Dynamic Regret

- General 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

- Static Regret:

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

- Worst-case Dynamic Regret:

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

The general dynamic regret can handle both the stationary and changing environments!

LAMDA Learning And Mining from DatA

Our Contributions

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

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

where

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

  - The bound automatically becomes tighter when P_T is small

Our Contributions

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

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

where

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

  - The bound automatically becomes tighter when P_T is small

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

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

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

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

LAMDA Learning And Mining from DatA

Our Contributions

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

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

where

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

  - The bound automatically becomes tighter when P_T is small

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

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

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

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

- Problem-dependent Algorithms [Zhao et al., 2020]

Ader [Zhang et al., 2018a]

The Basic Idea

- Discretize the possible values of $P_T \in [0, TD]$
- Create one expert (OGD) for each discrete P_T
- Prediction with expert advice

Ader [Zhang et al., 2018a]

The Basic Idea

- Discretize the possible values of $P_T \in [0, TD]$
- Create one expert (OGD) for each discrete P_T
- Prediction with expert advice

A Set of Experts

- Online Gradient Descent (OGD) with $\eta = 1$
- $\cdots$
- Online Gradient Descent (OGD) with $\eta = 1/\sqrt{T}$

$$\mathbf{w}_{t+1}^{\eta} = \Pi_{\mathcal{W}}\left[\mathbf{w}_t^{\eta} - \eta \nabla f_t(\mathbf{w}_t^{\eta})\right], \ \eta \in \mathcal{H}$$

Ader [Zhang et al., 2018a]

The Basic Idea

- Discretize the possible values of $P_T \in [0, TD]$
- Create one expert (OGD) for each discrete P_T
- Prediction with expert advice

A Set of Experts

- Online Gradient Descent (OGD) with $\eta = 1$
- $\cdots$
- Online Gradient Descent (OGD) with $\eta = 1/\sqrt{T}$

$$\mathbf{w}_{t+1}^{\eta} = \Pi_{\mathcal{W}}\left[\mathbf{w}_t^{\eta} - \eta \nabla f_t(\mathbf{w}_t^{\eta})\right], \ \eta \in \mathcal{H}$$

Aggregation

A Meta-algorithm

- The exponentially weighted average forecaster (Hedge)

$$\mathbf{w}_t = \sum_{\eta \in \mathcal{H}} \omega_t^{\eta} \mathbf{w}_t^{\eta}, \quad \omega_{t+1}^{\eta} = \frac{\omega_t^{\eta} e^{-\alpha f_t(\mathbf{w}_t^{\eta})}}{\sum_{\mu \in \mathcal{H}} \omega_t^{\mu} e^{-\alpha f_t(\mathbf{w}_t^{\mu})}}$$

Outline

Definitions

- Weakly Adaptive Regret [Hazan and Seshadhri, 2007]

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

  - The maximum regret over any interval

Definitions

- Weakly Adaptive Regret [Hazan and Seshadhri, 2007]

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

  - The maximum regret over any interval

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

$$\text{SA-Regret}(T,\tau) = \max_{[r,r+\tau-1]\subseteq[T]} \left(\sum_{t=r}^{r+\tau-1} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=r}^{r+\tau-1} f_t(\mathbf{w}) \right)$$

  - The maximum regret over any interval of length τ

Definitions

- Weakly Adaptive Regret [Hazan and Seshadhri, 2007]

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

  - The maximum regret over any interval

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

$$\text{SA-Regret}(T,\tau) = \max_{[r,r+\tau-1]\subseteq[T]} \left(\sum_{t=r}^{r+\tau-1} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=r}^{r+\tau-1} f_t(\mathbf{w}) \right)$$

  - The maximum regret over any interval of length τ

$$f_1(\cdot),\ f_2(\cdot), \ldots, f_\tau(\cdot), f_{\tau+1}(\cdot), \ldots, f_s(\cdot),\ f_{s+1}(\cdot), \ldots, f_{s+\tau-1}(\cdot), f_{s+\tau}(\cdot), \ldots$$

LAMDA Learning And Mining from DatA

Definitions

- Weakly Adaptive Regret [Hazan and Seshadhri, 2007]

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

  - The maximum regret over any interval

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

$$\text{SA-Regret}(T,\tau) = \max_{[r,r+\tau-1]\subseteq[T]} \left(\sum_{t=r}^{r+\tau-1} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=r}^{r+\tau-1} f_t(\mathbf{w}) \right)$$

  - The maximum regret over any interval of length τ

$$\underbrace{\boxed{f_1(\cdot),\, f_2(\cdot), \ldots, f_\tau(\cdot)}}_{\sum_{t=1}^{\tau} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=1}^{\tau} f_t(\mathbf{w})}, f_{\tau+1}(\cdot), \ldots, f_s(\cdot),\, f_{s+1}(\cdot), \ldots, f_{s+\tau-1}(\cdot), f_{s+\tau}(\cdot), \ldots$$

LAMDA Learning And Mining from DatA

Definitions

- Weakly Adaptive Regret [Hazan and Seshadhri, 2007]

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

  - The maximum regret over any interval

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

$$\text{SA-Regret}(T,\tau) = \max_{[r,r+\tau-1]\subseteq[T]} \left(\sum_{t=r}^{r+\tau-1} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=r}^{r+\tau-1} f_t(\mathbf{w}) \right)$$

  - The maximum regret over any interval of length τ

$$f_1(\cdot), \boxed{\overbrace{f_2(\cdot), \ldots, f_\tau(\cdot), f_{\tau+1}(\cdot)}^{\sum_{t=2}^{\tau+1} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=2}^{\tau+1} f_t(\mathbf{w})}}, \ldots, f_s(\cdot), f_{s+1}(\cdot), \ldots, f_{s+\tau-1}(\cdot), f_{s+\tau}(\cdot), \ldots$$

$$\underbrace{f_1(\cdot), f_2(\cdot), \ldots, f_\tau(\cdot)}_{\sum_{t=1}^{\tau} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=1}^{\tau} f_t(\mathbf{w})}$$

LAMDA Learning And Mining from DatA

Definitions

- Weakly Adaptive Regret [Hazan and Seshadhri, 2007]

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

  - The maximum regret over any interval

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

$$\text{SA-Regret}(T,\tau) = \max_{[r,r+\tau-1]\subseteq[T]} \left(\sum_{t=r}^{r+\tau-1} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=r}^{r+\tau-1} f_t(\mathbf{w}) \right)$$

  - The maximum regret over any interval of length τ

$$\underbrace{f_1(\cdot), \overbrace{f_2(\cdot), \ldots, f_\tau(\cdot)}^{\sum_{t=2}^{\tau+1} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=2}^{\tau+1} f_t(\mathbf{w})}}_{\sum_{t=1}^{\tau} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=1}^{\tau} f_t(\mathbf{w})}, f_{\tau+1}(\cdot), \ldots, \underbrace{\boxed{f_s(\cdot), f_{s+1}(\cdot), \ldots, f_{s+\tau-1}(\cdot)}}_{\sum_{t=s}^{s+\tau-1} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=s}^{s+\tau-1} f_t(\mathbf{w})}, f_{s+\tau}(\cdot), \ldots$$

Definitions

- Weakly Adaptive Regret [Hazan and Seshadhri, 2007]

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

  - The maximum regret over any interval

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

$$\text{SA-Regret}(T,\tau) = \max_{[r,r+\tau-1]\subseteq[T]} \left(\sum_{t=r}^{r+\tau-1} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=r}^{r+\tau-1} f_t(\mathbf{w}) \right)$$

  - The maximum regret over any interval of length τ

$$\rlap{$\underbrace{}_{\sum_{t=1}^{\tau} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=1}^{\tau} f_t(\mathbf{w})}$}f_1(\cdot), \overbrace{f_2(\cdot), \ldots, f_\tau(\cdot), f_{\tau+1}(\cdot)}^{\sum_{t=2}^{\tau+1} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=2}^{\tau+1} f_t(\mathbf{w})}, \ldots, \rlap{$\underbrace{\phantom{f_s(\cdot), f_{s+1}(\cdot), \ldots, f_{s+\tau-1}(\cdot)}}_{\sum_{t=s}^{s+\tau-1} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=s}^{s+\tau-1} f_t(\mathbf{w})}$}f_s(\cdot), \overbrace{\boxed{f_{s+1}(\cdot), \ldots, f_{s+\tau-1}(\cdot), f_{s+\tau}(\cdot)}}^{\sum_{t=s+1}^{s+\tau} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=s+1}^{s+\tau} f_t(\mathbf{w})}, \ldots$$

Learning And Mining from DatA

Existing Results

- Follow the leading history (FLH) [Hazan and Seshadhri, 2007]

$$\text{WA-Regret}(T) = \begin{cases} O\left(d \log^2 T\right), & \text{Exponentially Concave} \\ O\left(\sqrt{T \log^3 T}\right), & \text{Convex} \end{cases}$$

  - The $O(\sqrt{T \log^3 T})$ bound is meaningless for short intervals

Learning And Mining from DatA

Existing Results

- Follow the leading history (FLH) [Hazan and Seshadhri, 2007]

$$\text{WA-Regret}(T) = \begin{cases} O\left(d \log^2 T\right), & \text{Exponentially Concave} \\ O\left(\sqrt{T \log^3 T}\right), & \text{Convex} \end{cases}$$

  - The $O(\sqrt{T \log^3 T})$ bound is meaningless for short intervals

- Convex Functions

$$\text{SA-Regret}(T, \tau) = \begin{cases} O\left(\sqrt{\tau} \log T\right), & \text{[Daniely et al., 2015]} \\ O\left(\sqrt{\tau \log T}\right), & \text{[Jun et al., 2017a]} \end{cases}$$

Existing Results

- Follow the leading history (FLH) [Hazan and Seshadhri, 2007]

$$\text{WA-Regret}(T) = \begin{cases} O\left(d \log^2 T\right), & \text{Exponentially Concave} \\ O\left(\sqrt{T \log^3 T}\right), & \text{Convex} \end{cases}$$

  - The $O(\sqrt{T \log^3 T})$ bound is meaningless for short intervals

- Convex Functions

$$\text{SA-Regret}(T, \tau) = \begin{cases} O\left(\sqrt{\tau} \log T\right), & \text{[Daniely et al., 2015]} \\ O\left(\sqrt{\tau \log T}\right), & \text{[Jun et al., 2017a]} \end{cases}$$

- Convex and Smooth Functions [Jun et al., 2017b]

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

Our Contributions

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

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

- Efficient Algorithms for Adaptive Regret [Wang et al., 2018]
  - Reduce the # of gradient evaluations from $O(\log t)$ to 1

Our Contributions

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

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

- Efficient Algorithms for Adaptive Regret [Wang et al., 2018]
  - Reduce the # of gradient evaluations from $O(\log t)$ to 1
- Convex and Smooth Functions [Zhang et al., 2019a]

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

Our Contributions

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

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

- Efficient Algorithms for Adaptive Regret [Wang et al., 2018]
    - Reduce the # of gradient evaluations from $O(\log t)$ to 1

- Convex and Smooth Functions [Zhang et al., 2019a]

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

- A Universal Algorithm–UMA [Zhang et al., 2019b]

$$\text{Regret}\left([r,s]\right) = \begin{cases} O\left(\sqrt{(s-r)\log s}\right), & \text{Convex} \\ O\left(\log(s-r)\log T\right), & \text{Strongly Convex} \\ O\left(d\log(s-r)\log T\right), & \text{Exponentially Concave} \end{cases}$$

UMA [Zhang et al., 2019b]

- Two Types of Experts
  - $E^{\eta}_{\mathcal{J}}$: ONS for exp-concave surrogate loss:

$$\ell_t^{\eta}(\mathbf{w}) = -\eta\langle\nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w}\rangle + \eta^2\langle\nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w}\rangle^2$$

  - $\widehat{E}^{\eta}_{\mathcal{J}}$: OGD for strongly convex surrogate loss:

$$\hat{\ell}_t^{\eta}(\mathbf{w}) = -\eta\langle\nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w}\rangle + \eta^2 G^2\|\mathbf{w}_t - \mathbf{w}\|^2$$

UMA [Zhang et al., 2019b]

- Two Types of Experts
  - $E_{\mathcal{J}}^{\eta}$: ONS for exp-concave surrogate loss:

$$\ell_t^{\eta}(\mathbf{w}) = -\eta\langle\nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w}\rangle + \eta^2\langle\nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w}\rangle^2$$

  - $\widehat{E}_{\mathcal{J}}^{\eta}$: OGD for strongly convex surrogate loss:

$$\hat{\ell}_t^{\eta}(\mathbf{w}) = -\eta\langle\nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w}\rangle + \eta^2 G^2\|\mathbf{w}_t - \mathbf{w}\|^2$$

- A Set of Intervals: GC intervals [Daniely et al., 2015]

t	1	2	3	4	5	6	7	8	9	10	11	$\cdots$
$\mathcal{I}_0$	[ ]	[ ]	[ ]	[ ]	[ ]	[ ]	[ ]	[ ]	[ ]	[ ]	[ ]	$\cdots$
$\mathcal{I}_1$		[	]	[	]	[	]	[	]	[	]	$\cdots$
$\mathcal{I}_2$				[			]	[			]	$\cdots$

LAMDA Learning And Mining from DatA

UMA [Zhang et al., 2019b]

- Two Types of Experts
  - $E_{\mathcal{J}}^{\eta}$: ONS for exp-concave surrogate loss:

$$\ell_t^{\eta}(\mathbf{w}) = -\eta\langle\nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w}\rangle + \eta^2\langle\nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w}\rangle^2$$

  - $\widehat{E}_{\mathcal{J}}^{\eta}$: OGD for strongly convex surrogate loss:

$$\widehat{\ell}_t^{\eta}(\mathbf{w}) = -\eta\langle\nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w}\rangle + \eta^2 G^2\|\mathbf{w}_t - \mathbf{w}\|^2$$

- A Set of Intervals: GC intervals [Daniely et al., 2015]

t	1	2	3	4	5	6	7	8	9	10	11	$\cdots$
$\mathcal{I}_0$	[ ]	[ ]	[ ]	[ ]	[ ]	[ ]	[ ]	[ ]	[ ]	[ ]	[ ]	$\cdots$
$\mathcal{I}_1$		[	] [	] [	] [	] [	] [	] [	] [	] [	]	$\cdots$
$\mathcal{I}_2$				[			]	[		$\mathcal{J}$	]	$\cdots$

$$\begin{cases} E_{\mathcal{J}}^{\eta}, \eta \in \mathcal{S}(|\mathcal{J}|) \\ \widehat{E}_{\mathcal{J}}^{\eta}, \eta \in \mathcal{S}(|\mathcal{J}|) \end{cases}$$

$$\mathcal{S}(\tau) = \left\{ \frac{2^{-i}}{5DG} \;\middle|\; i = 0, 1, \ldots, \left\lceil \frac{1}{2}\log_2 \tau \right\rceil \right\}$$

UMA [Zhang et al., 2019b]

- Two Types of Experts
  - $E_{\mathcal{J}}^{\eta}$: ONS for exp-concave surrogate loss:

$$\ell_t^{\eta}(\mathbf{w}) = -\eta\langle \nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w}\rangle + \eta^2 \langle \nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w}\rangle^2$$

  - $\widehat{E}_{\mathcal{J}}^{\eta}$: OGD for strongly convex surrogate loss:

$$\hat{\ell}_t^{\eta}(\mathbf{w}) = -\eta\langle \nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w}\rangle + \eta^2 G^2 \|\mathbf{w}_t - \mathbf{w}\|^2$$

- A Set of Intervals: GC intervals [Daniely et al., 2015]

$$\begin{array}{lcccccccccccc}
t & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & \cdots \\
\mathcal{I}_0 & [\;] & [\;] & [\;] & [\;] & [\;] & [\;] & [\;] & [\;] & [\;] & [\;] & [\;] & \cdots \\
\mathcal{I}_1 & & [&] & [&] & [&] & [&] & [&] & \cdots \\
\mathcal{I}_2 & & & & [& & &] & \multicolumn{4}{c}{\boxed{[\quad \mathcal{J} \quad]}} & \cdots
\end{array}
\qquad
\boxed{\begin{cases} E_{\mathcal{J}}^{\eta}, \eta \in \mathcal{S}(|\mathcal{J}|) \\ \widehat{E}_{\mathcal{J}}^{\eta}, \eta \in \mathcal{S}(|\mathcal{J}|) \end{cases}}$$

$$\mathcal{S}(\tau) = \left\{ \frac{2^{-i}}{5DG} \;\middle|\; i = 0, 1, \ldots, \left\lceil \frac{1}{2}\log_2 \tau \right\rceil \right\}$$

- A Meta-algorithm
  - TWEA [van Erven and Koolen, 2016] with sleeping experts

Outline

The Relationship between Two Metrics [Zhang et al., 2018b]

An Upper Bound of the Worst-case Dynamic Regret

$$\text{D-Regret}(\mathbf{w}_1^*, \ldots, \mathbf{w}_T^*) \leq \min_{1 \leq \tau \leq T} \left(\frac{\text{SA-Regret}(T, \tau) T}{\tau} + 2\tau F_T \right)$$

where

$$F_T = \sum_{t=1}^{T} \sup_{\mathbf{w} \in \mathcal{W}} |f_{t+1}(\mathbf{w}) - f_t(\mathbf{w})|$$

The Relationship between Two Metrics [Zhang et al., 2018b]

An Upper Bound of the Worst-case Dynamic Regret

$$\text{D-Regret}(\mathbf{w}_1^*, \ldots, \mathbf{w}_T^*) \leq \min_{1 \leq \tau \leq T} \left(\frac{\text{SA-Regret}(T, \tau) T}{\tau} + 2\tau F_T \right)$$

where

$$F_T = \sum_{t=1}^{T} \sup_{\mathbf{w} \in \mathcal{W}} |f_{t+1}(\mathbf{w}) - f_t(\mathbf{w})|$$

- Dynamic Regret of Strongly Adaptive Methods

$$\text{D-Regret}(\mathbf{w}_1^*, \ldots, \mathbf{w}_T^*)$$

$$= \begin{cases} O\left(\max\left\{\sqrt{T \log T}, T^{2/3} F_T^{1/3} \log^{1/3} T\right\}\right), & \text{Convex} \\ O\left(\max\left\{\log T, \sqrt{T F_T \log T}\right\}\right), & \text{Strongly Convex} \\ O\left(d \cdot \max\left\{\log T, \sqrt{T F_T \log T}\right\}\right), & \text{Exponentially Concave} \end{cases}$$

- Do not need any prior knowledge of the functional variation

Minimizing Two Metrics Simultaneously [Zhang et al., 2020]

- State-of-the-art Results
  - Dynamic Regret [Zhang et al., 2018a]

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

  - Adaptive Regret [Jun et al., 2017a]

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

LAMDA Learning And Mining from DatA

Minimizing Two Metrics Simultaneously [Zhang et al., 2020]

- State-of-the-art Results
  - Dynamic Regret [Zhang et al., 2018a]

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

  - Adaptive Regret [Jun et al., 2017a]

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

- Adaptive Online learning with Dynamic regret (AOD)

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

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

Minimizing Two Metrics Simultaneously [Zhang et al., 2020]

- State-of-the-art Results
  - Dynamic Regret [Zhang et al., 2018a]

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

  - Adaptive Regret [Jun et al., 2017a]

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

- Adaptive Online learning with Dynamic regret (AOD)

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

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

- Adaptive Online learning based on Ader (AOA)

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

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

Conclusion and Future Work

Conclusion

- Dynamic Regret [Yang et al., 2016, Zhang et al., 2017, Zhang et al., 2018a, Zhao et al., 2020, Zhang et al., 2020]

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

- Adaptive Regret [Zhang et al., 2018b, Wang et al., 2018, Zhang et al., 2019a, Zhang et al., 2019b]

$$\text{SA-Regret}(T, \tau) = \max_{[r, r+\tau-1] \subseteq [T]} \left(\sum_{t=r}^{r+\tau-1} f_t(\mathbf{w}_t) - \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=r}^{r+\tau-1} f_t(\mathbf{w}) \right)$$

Conclusion and Future Work

Conclusion

- Dynamic Regret [Yang et al., 2016, Zhang et al., 2017, Zhang et al., 2018a, Zhao et al., 2020, Zhang et al., 2020]

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

- Adaptive Regret [Zhang et al., 2018b, Wang et al., 2018, Zhang et al., 2019a, Zhang et al., 2019b]

$$\text{SA-Regret}(T, \tau) = \max_{[r, r+\tau-1] \subseteq [T]} \left(\sum_{t=r}^{r+\tau-1} f_t(\mathbf{w}_t) - \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=r}^{r+\tau-1} f_t(\mathbf{w}) \right)$$

Future Work

- General dynamic regret of strongly convex functions
- General dynamic regret of exponentially concave functions
- Adaptive Regret v.s. Dynamic Regret

Reference I

Thanks!

Besbes, O., Gur, Y., and Zeevi, A. (2015).
Non-stationary stochastic optimization.
Operations Research, 63(5):1227–1244.

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. (2017a).
Improved strongly adaptive online learning using coin betting.
In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pages 943–951.

Jun, K.-S., Orabona, F., Wright, S., and Willett, R. (2017b).
Online learning for changing environments using coin betting.
Electronic Journal of Statistics, 11(2):5282–5310.

Reference II

Mokhtari, A., Shahrampour, S., Jadbabaie, A., and Ribeiro, A. (2016).
Online optimization in dynamic environments: Improved regret rates for strongly convex problems.
In *Proceedings of the 55th IEEE Conference on Decision and Control*, pages 7195–7201.

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

van Erven, T. and Koolen, W. M. (2016).
MetaGrad: Multiple learning rates in online learning.
In *Advances in Neural Information Processing Systems 29*, pages 3666–3674.

Wang, G., Zhao, D., and Zhang, L. (2018).
Minimizing adaptive regret with one gradient per iteration.
In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 2762–2768.

Yang, T., Zhang, L., Jin, R., and Yi, J. (2016).
Tracking slowly moving clairvoyant: Optimal dynamic regret of online learning with true and noisy gradient.
In *Proceedings of the 33rd International Conference on Machine Learning*, pages 449–457.

Zhang, L., Liu, T.-Y., and Zhou, Z.-H. (2019a).
Adaptive regret of convex and smooth functions.
In *Proceedings of the 36th International Conference on Machine Learning*, pages 7414–7423.

Zhang, L., Lu, S., and Yang, T. (2020).
Minimizing dynamic regret and adaptive regret simultaneously.
In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*.

Reference III

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

Zhang, L., Wang, G., Tu, W.-W., and Zhou, Z.-H. (2019b).
Dual adaptivity: A universal algorithm for minimizing the adaptive regret of convex functions.
ArXiv e-prints, arXiv:1906.10851.

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

Zhang, L., Yang, T., Yi, J., Jin, R., and Zhou, Z.-H. (2017).
Improved dynamic regret for non-degenerate functions.
In *Advances in Neural Information Processing Systems 30*, pages 732–741.

Zhao, P., Zhang, Y.-J., Zhang, L., and Zhou, Z.-H. (2020).
Dynamic regret of convex and smooth functions.
In *Advances in Neural Information Processing Systems 33*.

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