Full-information Online Learning

Lijun Zhang

Nanjing University, China

June 2, 2017

Outline

1. Introduction
   - Definitions
   - Regret
2. Prediction with Expert Advice
3. Online Convex Optimization
   - Convex Functions
   - Strongly Convex Functions
   - Exponentially Concave Functions

Outline

What Happens in an Internet Minute?

http://www.dailyinfographic.com/what-happens-in-an-internet-minute-infographic

Outline

Online Algorithm vs Online Learning

Online Algorithm [Karp, 1992]

An online algorithm is one that receives a sequence of requests and performs an immediate action in response to each request.

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.

Learning And Mining from DatA

Online Algorithm vs Online Learning

Online Algorithm [Karp, 1992]

An online algorithm is one that receives a sequence of requests and performs an immediate action in response to each request.

- Computer Vision, Machine Learning, Data Mining
- Theoretical Computer Science, Computer Networks

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 Algorithm vs Online Learning

Online Algorithm [Karp, 1992]

An online algorithm is one that receives a sequence of requests and performs an immediate action in response to each request.

- Computer Vision, Machine Learning, Data Mining
- Theoretical Computer Science, Computer Networks

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.

- Machine Learning, Game Theory, Information Theory

Online Algorithm vs Online Learning

Online Algorithm [Karp, 1992]

An online algorithm is one that receives a sequence of requests and performs an immediate action in response to each request.

- The Ski Rental Problem: in the t-th day

$$\text{Rent } (1\$) \text{ or } \text{ Buy } (10\$)?$$

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 Algorithm vs Online Learning

Online Algorithm [Karp, 1992]

An online algorithm is one that receives a sequence of requests and performs an immediate action in response to each request.

- The Ski Rental Problem: in the t-th day

$$\text{Rent } (1\$) \text{ or } \text{ Buy } (10\$)?$$

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 Classification: in the t-th round

$$\text{Recevie } \mathbf{x}_t \in \mathbb{R}^d, \text{ predict } \hat{y}_t, \text{ observe } y_t$$

Full-information vs Bandit

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.

Full-information vs Bandit

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.

- Full-information Online Learning

Full-information vs Bandit

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.

- Bandit Online Learning

Formal Definitions

Online Learning

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

4: **end for**

Formal Definitions

Online Learning

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

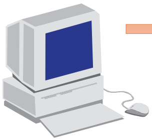

Learner

A classifier

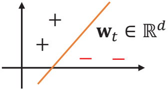

An example $(\mathbf{x}_t, y_t) \in \mathbb{R}^d \times \{\pm 1\}$

A loss $f_t(\mathbf{w}) = \max(1 - y_t \mathbf{w}^\top \mathbf{x}_t, 0)$

Adversary

Formal Definitions

Online Learning

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

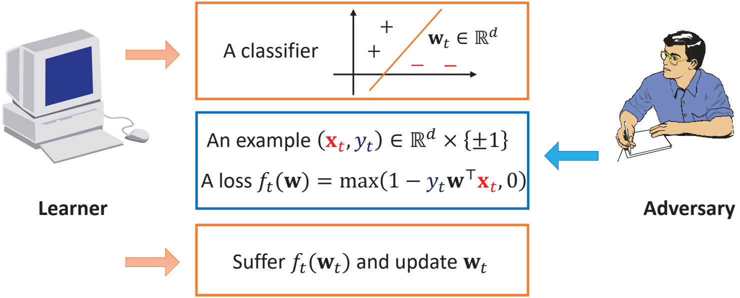

Formal Definitions

Online Learning

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

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

Formal Definitions

Online Learning

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

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

- Full-information Online Learning
  - Learner observes the function $f_t(\cdot)$
  - Online first-order optimization
- Bandit Online Learning
  - Learner only observes the value of $f_t(\mathbf{w}_t)$
  - Online zero-order optimization

Outline

Regret

Cumulative Loss

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

Regret

Cumulative Loss

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

Regret

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

Regret

Cumulative Loss

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

Regret

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

Regret

Cumulative Loss

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

Regret

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

Hannan Consistent

$$\limsup_{T\to\infty} \frac{1}{T}\left(\sum_{t=1}^{T} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=1}^{T} f_t(\mathbf{w})\right) = 0, \text{ with probability } 1$$

Regret

Cumulative Loss

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

Regret

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

Hannan Consistent

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

Outline

LAMDA Learning And Mining from DatA

Definitions

- The Prediction Protocol

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

end for

Definitions

- The Prediction Protocol

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

    The environment chooses the next outcome y_t and the expert advice $\{f_{i,t} : i \in [N]\}$

end for

Definitions

- The Prediction Protocol

for $t = 1, \ldots, T$ **do**
    The environment chooses the next outcome y_t and the expert advice $\{f_{i,t} : i \in [N]\}$
    The expert advice is revealed to the forecaster
end for

Definitions

- The Prediction Protocol

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

  The environment chooses the next outcome y_t and the expert advice $\{f_{i,t} : i \in [N]\}$

  The expert advice is revealed to the forecaster

  The forecaster chooses the prediction $\hat{p}_t$

end for

Definitions

- The Prediction Protocol

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

    The environment chooses the next outcome y_t and the expert advice $\{f_{i,t} : i \in [N]\}$

    The expert advice is revealed to the forecaster

    The forecaster chooses the prediction $\hat{p}_t$

    The environment reveals the next outcome y_t

end for

Definitions

- The Prediction Protocol

for $t = 1, \ldots, T$ **do**
    The environment chooses the next outcome y_t and the expert advice $\{f_{i,t} : i \in [N]\}$
    The expert advice is revealed to the forecaster
    The forecaster chooses the prediction $\hat{p}_t$
    The environment reveals the next outcome y_t
    The forecaster incurs loss $\ell(\hat{p}_t, y_t)$ and each expert i incurs loss $\ell(f_{i,t}, y_t)$
end for

Definitions

- The Prediction Protocol

for $t = 1, \ldots, T$ **do**
    The environment chooses the next outcome y_t and the expert advice $\{f_{i,t} : i \in [N]\}$
    The expert advice is revealed to the forecaster
    The forecaster chooses the prediction $\hat{p}_t$
    The environment reveals the next outcome y_t
    The forecaster incurs loss $\ell(\hat{p}_t, y_t)$ and each expert i incurs loss $\ell(f_{i,t}, y_t)$
end for

- Regret with respect to expect i

$$R_{i,T} = \sum_{t=1}^{T} \left(\ell(\widehat{p}_t, y_t) - \ell(f_{i,t}, y_t)\right) = \widehat{L}_T - L_{i,T}$$

Online Algorithms

- Weighted Average Prediction

$$\widehat{p}_t = \frac{\sum_{i=1}^N w_{i,t-1} f_{i,t}}{\sum_{j=1}^N w_{j,t-1}}$$

where $w_{i,t-1}$ is the weight assigned to expert i

Online Algorithms

- Weighted Average Prediction

$$\widehat{p}_t = \frac{\sum_{i=1}^N w_{i,t-1} f_{i,t}}{\sum_{j=1}^N w_{j,t-1}}$$

where $w_{i,t-1}$ is the weight assigned to expert i

- Polynomially Weighted Average Forecaster

$$w_{i,t-1} = \frac{2(R_{i,t-1})_+^{p-1}}{\|(\mathbf{R}_{t-1})_+\|_p^{p-2}}$$

and

$$\widehat{p}_t = \frac{\sum_{i=1}^N \left(\sum_{s=1}^{t-1} \left(\ell(\widehat{p}_s, y_s) - \ell(f_{i,s}, y_s)\right)\right)_+^{p-1} f_{i,t}}{\sum_{j=1}^N \left(\sum_{s=1}^{t-1} \left(\ell(\widehat{p}_s, y_s) - \ell(f_{j,s}, y_s)\right)\right)_+^{p-1}}$$

Online Algorithms

- Weighted Average Prediction

$$\widehat{p}_t = \frac{\sum_{i=1}^N w_{i,t-1} f_{i,t}}{\sum_{j=1}^N w_{j,t-1}}$$

where $w_{i,t-1}$ is the weight assigned to expert i

- Exponentially Weighted Average Forecaster

$$w_{i,t-1} = \frac{e^{\eta R_{i,t-1}}}{\sum_{j=1}^N e^{\eta R_{j,t-1}}}$$

and

$$\widehat{p}_t = \frac{\sum_{i=1}^N \exp\left(\eta(\widehat{L}_{t-1} - L_{i,t-1})\right) f_{i,t}}{\sum_{j=1}^N \exp\left(\eta(\widehat{L}_{t-1} - L_{j,t-1})\right)} = \frac{\sum_{i=1}^N e^{-\eta L_{i,t-1}} f_{i,t}}{\sum_{j=1}^N e^{-\eta L_{j,t-1}}}$$

Regret Bounds I

Corollary 2.1 of [Cesa-Bianchi and Lugosi, 2006]

Assume that the loss function ℓ is convex in its first argument and that it takes values in $[0, 1]$. Then, for any sequence $y_1, y_2, \ldots$ of outcomes and for any $T \geq 1$, the regret of the polynomially weighted average forecaster satisfies

$$\widehat{L}_T - \min_{i=1,\ldots,N} L_{i,T} \leq \sqrt{T(p-1)N^{2/p}}$$

LAMDA Learning And Mining from DatA

Regret Bounds I

Corollary 2.1 of [Cesa-Bianchi and Lugosi, 2006]

Assume that the loss function ℓ is convex in its first argument and that it takes values in $[0, 1]$. Then, for any sequence $y_1, y_2, \ldots$ of outcomes and for any $T \geq 1$, the regret of the polynomially weighted average forecaster satisfies

$$\widehat{L}_T - \min_{i=1,\ldots,N} L_{i,T} \leq \sqrt{T(p-1)N^{2/p}}$$

When $p = 2$, we have

$$\widehat{L}_T - \min_{i=1,\ldots,N} L_{i,T} \leq \sqrt{TN}$$

When $p = 2 \ln N$, we have

$$\widehat{L}_T - \min_{i=1,\ldots,N} L_{i,T} \leq \sqrt{Te(2\ln N - 1)}$$

Regret Bounds II

Corollary 2.2 of [Cesa-Bianchi and Lugosi, 2006]

Assume that the loss function ℓ is convex in its first argument and that it takes values in $[0, 1]$. Then, for any sequence $y_1, y_2, \ldots$ of outcomes and for any $T \geq 1$, the regret of the exponentially weighted average forecaster satisfies

$$\widehat{L}_T - \min_{i=1,\ldots,N} L_{i,T} \leq \frac{\ln N}{\eta} + \frac{\eta T}{2}$$

Regret Bounds II

Corollary 2.2 of [Cesa-Bianchi and Lugosi, 2006]

Assume that the loss function ℓ is convex in its first argument and that it takes values in $[0,1]$. Then, for any sequence $y_1, y_2, \ldots$ of outcomes and for any $T \geq 1$, the regret of the exponentially weighted average forecaster satisfies

$$\widehat{L}_T - \min_{i=1,\ldots,N} L_{i,T} \leq \frac{\ln N}{\eta} + \frac{\eta T}{2}$$

When $\eta = \sqrt{2 \ln N / T}$, we have

$$\widehat{L}_T - \min_{i=1,\ldots,N} L_{i,T} \leq \sqrt{2T \ln N}$$

Extensions

- Tighter Regret Bounds
  - Small Losses

$$\widehat{L}_T - \min_{i=1,\dots,N} L_{i,T} \leq \sqrt{2L_T^* \ln N} + \ln N$$

where $L_T^* = \min_{i=1,\dots,N} L_{i,T}$

  - Exp-concave Losses

$$\widehat{L}_T - \min_{i=1,\dots,N} L_{i,T} \leq \frac{\ln N}{\eta}$$

Extensions

- Tighter Regret Bounds
  - Small Losses

$$\widehat{L}_T - \min_{i=1,\ldots,N} L_{i,T} \leq \sqrt{2L_T^* \ln N} + \ln N$$

where $L_T^* = \min_{i=1,\ldots,N} L_{i,T}$

  - Exp-concave Losses

$$\widehat{L}_T - \min_{i=1,\ldots,N} L_{i,T} \leq \frac{\ln N}{\eta}$$

- Tracking Regret

$$R(i_1, \ldots, i_T) = \sum_{t=1}^{T} \left(\ell(\widehat{p}_t, y_t) - \ell(f_{i_t,t}, y_t) \right)$$

LAMDA Learning And Mining from DatA

Outline

LAMDA Learning And Mining from DatA

Definitions

- Online Convex Optimization

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**

where $\mathcal{W}$ and f_t's are convex

Definitions

- Online Convex Optimization

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**

where $\mathcal{W}$ and f_t's are convex

- Online Gradient Descent

$$\mathbf{w}_{t+1} = \Pi_{\mathcal{W}}\left(\mathbf{w}_t - \eta_t \nabla f_t(\mathbf{w}_t)\right)$$

where $\Pi_{\mathcal{W}}(\cdot)$ is the projection operator

Outline

Analysis I

Define $\mathbf{w}'_{t+1} = \mathbf{w}_t - \eta_t \nabla f_t(\mathbf{w}_t)$. For any $\mathbf{w} \in \mathcal{W}$, we have

$$
\begin{aligned}
f_t(\mathbf{w}_t) - f_t(\mathbf{w}) \leq & \langle \nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w} \rangle = \frac{1}{\eta_t} \langle \mathbf{w}_t - \mathbf{w}'_{t+1}, \mathbf{w}_t - \mathbf{w} \rangle \\
= & \frac{1}{2\eta_t} \left(\|\mathbf{w}_t - \mathbf{w}\|_2^2 - \|\mathbf{w}'_{t+1} - \mathbf{w}\|_2^2 + \|\mathbf{w}_t - \mathbf{w}'_{t+1}\|_2^2 \right) \\
= & \frac{1}{2\eta_t} \left(\|\mathbf{w}_t - \mathbf{w}\|_2^2 - \|\mathbf{w}'_{t+1} - \mathbf{w}\|_2^2 \right) + \frac{\eta_t}{2} \|\nabla f_t(\mathbf{w}_t)\|_2^2 \\
\leq & \frac{1}{2\eta_t} \left(\|\mathbf{w}_t - \mathbf{w}\|_2^2 - \|\mathbf{w}_{t+1} - \mathbf{w}\|_2^2 \right) + \frac{\eta_t}{2} \|\nabla f_t(\mathbf{w}_t)\|_2^2
\end{aligned}
$$

By adding the inequalities of all iterations, we have

$$
\begin{aligned}
\sum_{t=1}^{T} f_t(\mathbf{w}_t) - f_t(\mathbf{w}) \leq & \frac{1}{2\eta_1} \|\mathbf{w}_1 - \mathbf{w}\|_2^2 - \frac{1}{2\eta_T} \|\mathbf{w}_{T+1} - \mathbf{w}\|_2^2 \\
& + \frac{1}{2} \sum_{t=2}^{T} \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \|\mathbf{w}_t - \mathbf{w}\|_2^2 + \frac{1}{2} \sum_{t=1}^{T} \eta_t \|\nabla f_t(\mathbf{w}_t)\|_2^2
\end{aligned}
$$

Analysis II

Assuming

$$\eta_t \le \eta_{t-1}, \ \|\mathbf{w}_t - \mathbf{w}\|_2^2 \le D^2, \text{ and } \|\nabla f_t(\mathbf{w}_t)\|_2^2 \le G^2,$$

we have

$$\begin{aligned}\sum_{t=1}^{T} f_t(\mathbf{w}_t) - f_t(\mathbf{w}) \le & \frac{D^2}{2\eta_1} + \frac{D^2}{2}\sum_{t=2}^{T}\left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}}\right) + \frac{G^2}{2}\sum_{t=1}^{T}\eta_t \\ = & \frac{D^2}{2\eta_T} + \frac{G^2}{2}\sum_{t=1}^{T}\eta_t\end{aligned}$$

By setting

$$\eta_t = \frac{1}{\sqrt{t}},$$

we have

$$\sum_{t=1}^{T} f_t(\mathbf{w}_t) - f_t(\mathbf{w}) \le \sqrt{T}\left(\frac{D^2}{2} + G^2\right)$$

Outline

1. Introduction
   - Definitions
   - Regret
2. Prediction with Expert Advice
3. Online Convex Optimization
   - Convex Functions
   - Strongly Convex Functions
   - Exponentially Concave Functions

Analysis I

Define $\mathbf{w}'_{t+1} = \mathbf{w}_t - \eta_t \nabla f_t(\mathbf{w}_t)$. For any $\mathbf{w} \in \mathcal{W}$, we have

$$
\begin{aligned}
& f_t(\mathbf{w}_t) - f_t(\mathbf{w}) \\
\leq & \langle \nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w} \rangle - \frac{\lambda}{2} \|\mathbf{w}_t - \mathbf{w}\|_2^2 \\
\leq & \frac{1}{2\eta_t} \left(\|\mathbf{w}_t - \mathbf{w}\|_2^2 - \|\mathbf{w}_{t+1} - \mathbf{w}\|_2^2 \right) + \frac{\eta_t}{2} \|\nabla f_t(\mathbf{w}_t)\|_2^2 - \frac{\lambda}{2} \|\mathbf{w}_t - \mathbf{w}\|_2^2
\end{aligned}
$$

By adding the inequalities of all iterations, we have

$$
\begin{aligned}
\sum_{t=1}^{T} f_t(\mathbf{w}_t) - f_t(\mathbf{w}) \leq & \frac{1}{2\eta_1} \|\mathbf{w}_1 - \mathbf{w}\|_2^2 - \frac{\lambda}{2} \|\mathbf{w}_1 - \mathbf{w}\|_2^2 - \frac{1}{2\eta_T} \|\mathbf{w}_{T+1} - \mathbf{w}\|_2^2 \\
& + \frac{1}{2} \sum_{t=2}^{T} \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} - \lambda \right) \|\mathbf{w}_t - \mathbf{w}\|_2^2 + \frac{1}{2} \sum_{t=1}^{T} \eta_t \|\nabla f_t(\mathbf{w}_t)\|_2^2
\end{aligned}
$$

Analysis II

Assuming

$$\|\nabla f_t(\mathbf{w}_t)\|_2^2 \leq G^2,$$

we have

$$\sum_{t=1}^{T} f_t(\mathbf{w}_t) - f_t(\mathbf{w}) \leq \frac{1}{2\eta_1}\|\mathbf{w}_1 - \mathbf{w}\|_2^2 - \frac{\lambda}{2}\|\mathbf{w}_1 - \mathbf{w}\|_2^2 + \frac{1}{2}\sum_{t=2}^{T}\left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} - \lambda\right)\|\mathbf{w}_t - \mathbf{w}\|_2^2 + \frac{G^2}{2}\sum_{t=1}^{T}\eta_t$$

By setting

$$\eta_t = \frac{1}{\lambda t},$$

we have

$$\sum_{t=1}^{T} f_t(\mathbf{w}_t) - f_t(\mathbf{w}) \leq \frac{G^2}{2\lambda}\sum_{t=1}^{T}\frac{1}{t} \leq \frac{G^2}{2\lambda}(\log T + 1)$$

Outline

Definitions

Exponential Concavity

A function $f(\cdot) : \mathcal{W} \mapsto \mathbb{R}$ is α-exp-concave if $\exp(-\alpha f(\cdot))$ is concave over $\mathcal{W}$.

Definitions

Exponential Concavity

A function $f(\cdot) : \mathcal{W} \mapsto \mathbb{R}$ is α-exp-concave if $\exp(-\alpha f(\cdot))$ is concave over $\mathcal{W}$.

- Logistic Loss

$$f(\mathbf{w}) = \log\left(1 + \exp(-y\mathbf{x}^\top \mathbf{w})\right)$$

- Square Loss

$$f(\mathbf{w}) = (\mathbf{x}^\top \mathbf{w} - y)^2$$

- Negative Logarithm Loss

$$f(\mathbf{w}) = -\log(\mathbf{x}^\top \mathbf{w})$$

Online Newton Step

The Algorithm

1: $A_1 = \frac{1}{\beta^2 D^2} I$

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

3:

$$A_{t+1} = A_t + \nabla f_t(\mathbf{w}_t)[\nabla f_t(\mathbf{w}_t)]^\top$$

4:

$$\begin{aligned} \mathbf{w}_{t+1} =& \Pi_{\mathcal{W}}^{A_t}\left(\mathbf{w}_t - \frac{1}{\beta} A_t^{-1} \nabla f_t(\mathbf{w}_t)\right) \\ =& \operatorname*{argmin}_{\mathbf{w} \in \mathcal{W}} \left(\mathbf{w} - \mathbf{w}'_{t+1}\right) A_t \left(\mathbf{w} - \mathbf{w}'_{t+1}\right)^\top \end{aligned}$$

where $\mathbf{w}'_{t+1} = \mathbf{w}_t - \frac{1}{\beta} A_t^{-1} \nabla f_t(\mathbf{w}_t)$

5: **end for**

Regret Bound

Theorem 2 of [Hazan et al., 2007]

Assume that for all t , the loss function $f_t : \mathcal{W} \subseteq \mathbb{R}^d \mapsto \mathbb{R}$ is α-exp-concave and has the property that

$$\|\nabla f_t(\mathbf{w})\|_2 \leq G, \ \forall \mathbf{w} \in \mathcal{W}, t \in [T]$$

Then, online Newton Step with

$$\beta = \frac{1}{2} \min \left\{ \frac{1}{4GD}, \alpha \right\}$$

has the following regret bound

$$\sum_{t=1}^{T} f_t(\mathbf{w}_t) - \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^{T} f_t(\mathbf{w}) \leq 5 \left(\frac{1}{\alpha} + GD \right) d \log T$$

Online-to-batch Conversion

- Statistical Assumption
  - $f_1, f_2, \ldots$ are sampled independently from $\mathcal{P}$
  - Define

$$F(\mathbf{w}) = \mathrm{E}_{f \sim \mathcal{P}} [f(\mathbf{w})]$$

Online-to-batch Conversion

- Statistical Assumption
  - $f_1, f_2, \ldots$ are sampled independently from $\mathcal{P}$
  - Define

$$F(\mathbf{w}) = \mathrm{E}_{f \sim \mathcal{P}} \left[f(\mathbf{w}) \right]$$

- Regret Bound

$$\sum_{t=1}^{T} f_t(\mathbf{w}_t) - \sum_{t=1}^{T} f_t(\mathbf{w}) \leq C$$

Taking expectation over both sides, we have

$$\mathrm{E}\left[\sum_{t=1}^{T} F(\mathbf{w}_t) \right] - \sum_{t=1}^{T} F(\mathbf{w}) \leq C$$

Online-to-batch Conversion

- Statistical Assumption
  - $f_1, f_2, \ldots$ are sampled independently from $\mathcal{P}$
  - Define

$$F(\mathbf{w}) = \mathrm{E}_{f \sim \mathcal{P}}\left[f(\mathbf{w})\right]$$

- Regret Bound

$$\sum_{t=1}^{T} f_t(\mathbf{w}_t) - \sum_{t=1}^{T} f_t(\mathbf{w}) \leq C$$

Taking expectation over both sides, we have

$$\mathrm{E}\left[\sum_{t=1}^{T} F(\mathbf{w}_t)\right] - \sum_{t=1}^{T} F(\mathbf{w}) \leq C$$

- Risk Bound

$$\mathrm{E}\left[F(\bar{\mathbf{w}})\right] - F(\mathbf{w}) \leq \frac{C}{T}, \text{ where } \bar{\mathbf{w}} = \frac{1}{T}\sum_{t=1}^{T} \mathbf{w}_t$$

Examples of Conversion

- Convex Functions

$$\mathrm{E}\left[F(\bar{\mathbf{w}})\right] - F(\mathbf{w}) \leq \frac{1}{\sqrt{T}}\left(\frac{D^2}{2} + G^2\right)$$

High-probability Bound [Cesa-bianchi et al., 2002]

LAMDA Learning And Mining from DatA

Examples of Conversion

- Convex Functions

$$\mathrm{E}\left[F(\bar{\mathbf{w}})\right] - F(\mathbf{w}) \leq \frac{1}{\sqrt{T}}\left(\frac{D^2}{2} + G^2\right)$$

High-probability Bound [Cesa-bianchi et al., 2002]

- Strongly Convex Functions

$$\mathrm{E}\left[F(\bar{\mathbf{w}})\right] - F(\mathbf{w}) \leq \frac{G^2}{2\lambda T}(\log T + 1)$$

High-probability Bound [Kakade and Tewari, 2009]

Examples of Conversion

- Convex Functions

$$\mathrm{E}\left[F(\bar{\mathbf{w}})\right] - F(\mathbf{w}) \leq \frac{1}{\sqrt{T}} \left(\frac{D^2}{2} + G^2 \right)$$

High-probability Bound [Cesa-bianchi et al., 2002]

- Strongly Convex Functions

$$\mathrm{E}\left[F(\bar{\mathbf{w}})\right] - F(\mathbf{w}) \leq \frac{G^2}{2\lambda T}(\log T + 1)$$

High-probability Bound [Kakade and Tewari, 2009]

- Exponentially Concave Functions

$$\mathrm{E}\left[F(\bar{\mathbf{w}})\right] - F(\mathbf{w}) \leq 5 \left(\frac{1}{\alpha} + GD \right) \frac{d \log T}{T}$$

High-probability Bound [Mahdavi et al., 2015]

Extensions

- Faster Rates [Srebro et al., 2010]

$$\sum_{t=1}^{T} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=1}^{T} f_t(\mathbf{w}) \leq O(\sqrt{Tf_*})$$

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

Extensions

- Faster Rates [Srebro et al., 2010]

$$\sum_{t=1}^{T} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=1}^{T} f_t(\mathbf{w}) \leq O(\sqrt{Tf_*})$$

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

- Dynamic Regret [Zinkevich, 2003, Yang et al., 2016]

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

Extensions

- Faster Rates [Srebro et al., 2010]

$$\sum_{t=1}^{T} f_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \sum_{t=1}^{T} f_t(\mathbf{w}) \leq O(\sqrt{Tf_*})$$

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

- Dynamic Regret [Zinkevich, 2003, Yang et al., 2016]

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

- Adaptive Regret

[Hazan and Seshadhri, 2007, Daniely et al., 2015]

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

Reference I

Cesa-bianchi, N., Conconi, A., and Gentile, C. (2002).
On the generalization ability of on-line learning algorithms.
In *Advances in Neural Information Processing Systems 14*, pages 359–366.

Cesa-Bianchi, N. and Lugosi, G. (2006).
Prediction, Learning, and Games.
Cambridge University Press.

Daniely, A., Gonen, A., and Shalev-Shwartz, S. (2015).
Strongly adaptive online learning.
In *Proceedings of The 32nd International Conference on Machine Learning*.

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.

Kakade, S. M. and Tewari, A. (2009).
On the generalization ability of online strongly convex programming algorithms.
In *Advances in Neural Information Processing Systems 21*, pages 801–808.

Reference II

Karp, R. M. (1992).
On-line algorithms versus off-line algorithms: How much is it worth to know the future?
In *Proceedings of the IFIP 12th World Computer Congress on Algorithms, Software, Architecture – Information Processing*, pages 416–429.

Mahdavi, M., Zhang, L., and Jin, R. (2015).
Lower and upper bounds on the generalization of stochastic exponentially concave optimization.
In *Proceedings of the 28th Conference on Learning Theory*.

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

Srebro, N., Sridharan, K., and Tewari, A. (2010).
Optimistic rates for learning with a smooth loss.
ArXiv e-prints, arXiv:1009.3896.

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*.

Reference III

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