



Lecture 6. Prediction with Expert Advice

Advanced Optimization (Fall 2023)

Peng Zhao zhaop@lamda.nju.edu.cn Nanjing University

Outline

- Problem Setup
- Algorithms

Connection to Online Convex Optimization

Motivation

• Consider that we are making predictions based on external experts.



Advanced Optimization (Fall 2023)

Prediction with Expert Advice

• Other examples include



Advanced Optimization (Fall 2023)

PEA Problem Setup



PEA Problem Setup



Advanced Optimization (Fall 2023)

PEA: Formulization

• The online learner (player) aims to make the prediction based by combining *N* experts' advice.

```
At each round t = 1, 2, \cdots
```

(1) the player first picks a weight p_t from a simplex Δ_N ;

- (2) and simultaneously environments pick a loss vector $\ell_t \in \mathbb{R}^N$;
- (3) the player suffers loss $f_t(\mathbf{p}_t) \triangleq \langle \mathbf{p}_t, \boldsymbol{\ell}_t \rangle$, observes $\boldsymbol{\ell}_t$ and updates the model.

The feasible domain is the (N-1)-dim simplex $\Delta_N = \{ \boldsymbol{p} \in \mathbb{R}^N \mid p_i \ge 0, \sum_{i=1}^N p_i = 1 \}.$

We typically assume that $0 \leq \ell_{t,i} \leq 1$ holds for all $t \in [T]$ and $i \in [N]$.

PEA: Formulization

• The online learner (player) aims to make the prediction based by combining *N* experts' advice.

```
At each round t = 1, 2, \cdots
```

- (1) the player first picks a weight p_t from a simplex Δ_N ;
- (2) and simultaneously environments pick a loss vector $\ell_t \in \mathbb{R}^N$;
- (3) the player suffers loss $f_t(\mathbf{p}_t) \triangleq \langle \mathbf{p}_t, \boldsymbol{\ell}_t \rangle$, observes $\boldsymbol{\ell}_t$ and updates the model.
- The goal is to minimize the regret with respect to the *best expert*:

$$\operatorname{Regret}_{T} \triangleq \sum_{t=1}^{T} \langle \boldsymbol{p}_{t}, \boldsymbol{\ell}_{t} \rangle - \min_{\boldsymbol{p} \in \Delta_{N}} \sum_{t=1}^{T} \langle \boldsymbol{p}, \boldsymbol{\ell}_{t} \rangle = \sum_{t=1}^{T} \langle \boldsymbol{p}_{t}, \boldsymbol{\ell}_{t} \rangle - \min_{i \in [N]} \sum_{t=1}^{T} \ell_{t,i}$$

Advanced Optimization (Fall 2023)

A Natural Solution

• Follow the Leader (FTL)

Select the expert that *performs best so far*, specifically,

$$\boldsymbol{p}_{t}^{\text{FTL}} = \underset{\boldsymbol{p} \in \Delta_{N}}{\operatorname{arg\,min}} \langle \boldsymbol{p}, L_{t-1} \rangle = \underset{i \in [N]}{\operatorname{arg\,min}} L_{t-1,i}$$

where $L_{t-1} \in \mathbb{R}^{N}$ is the cumulative loss vector with $L_{t-1,i} \triangleq \sum_{s=1}^{t-1} \ell_{s,i}$.



Advanced Optimization (Fall 2023)

A Natural Solution

• Follow the Leader (FTL)

Select the expert that *performs best so far*, specifically,

$$p_t^{\text{FTL}} = \underset{p \in \Delta_N}{\operatorname{arg min}} \langle p, L_{t-1} \rangle = \underset{i \in [N]}{\operatorname{arg min}} L_{t-1,i}$$

where $L_{t-1} \in \mathbb{R}^N$ is the cumulative loss vector with $L_{t-1,i} \triangleq \sum_{s=1}^{t-1} \ell_{s,i}$.

 \implies Pitfall: decision is actually a one-hot vector, which can be very *unstable*.

Replacing the 'max' operation in FTL by '*softmax*'.

Hedge: Algorithm

• Hedge: replacing the 'max' operation in FTL by '*softmax*'.

At each round $t = 1, 2, \cdots$ (1) compute $p_t \in \Delta_N$ such that $p_{t,i} \propto \exp(-\eta L_{t-1,i})$ for $i \in [N]$ (2) the player submits p_t , suffers loss $\langle p_t, \ell_t \rangle$, and observes loss $\ell_t \in \mathbb{R}^N$ (3) update $L_t = L_{t-1} + \ell_t$

 $egin{aligned} \mathbf{FTL} & \mathbf{update} \ oldsymbol{p}_t^{ ext{FTL}} = rg\max_{oldsymbol{p} \in \Delta_N} ig\langle oldsymbol{p}, -oldsymbol{L}_{t-1} ig
angle \end{aligned}$

Hedge update $p_{t,i} \propto \exp(-\eta L_{t-1,i}), \forall i \in [N]$

Lazy and Greedy Updates

• Hedge algorithm

$$p_{t+1,i} \propto \exp\left(-\eta L_{t,i}\right), \forall i \in [N] \qquad L_{t,i} = \sum_{s=1}^{t} \ell_{s,i}, \ \forall i \in [N]$$

• Another equivalent update (when the learning rate η is *fixed*)

 $p_{t+1,i} \propto p_{t,i} \exp{(-\eta \ell_{t,i})}, \forall i \in [N]$ greedy update

where we set the uniform initialization as $p_{0,i} = 1/N$, $\forall i \in [N]$.

> But the two updates can be *significantly different when learning rate is changing*.

Hedge: Regret Bound

Theorem 1. Suppose that $\forall t \in [T]$ and $i \in [N], 0 \leq \ell_{t,i} \leq 1$, then Hedge with learning rate η guarantees

$$\operatorname{Regret}_T \leq \frac{\ln N}{\eta} + \eta T = \mathcal{O}(\sqrt{T \log N}),$$

where the last equality is by setting η optimally as $\sqrt{(\ln N)/T}$.

Proof. We present a 'potential-based' proof here, where the potential is defined as

$$\Phi_t \triangleq \frac{1}{\eta} \ln \left(\sum_{i=1}^N \exp\left(-\eta L_{t,i}\right) \right).$$

Proof of Hedge Regret Bound

Proof. $\Phi_t - \Phi_{t-1} = \frac{1}{\eta} \ln \left(\frac{\sum_{i=1}^N \exp\left(-\eta L_{t,i}\right)}{\sum_{i=1}^N \exp\left(-\eta L_{t-1,i}\right)} \right) \qquad \Phi_t \triangleq \frac{1}{\eta} \ln \left(\sum_{i=1}^N \exp\left(-\eta L_{t,i}\right) \right)$ $= \frac{1}{\eta} \ln \left(\sum_{i=1}^{N} \left(\frac{\exp\left(-\eta L_{t-1,i}\right)}{\sum_{i=1}^{N} \exp\left(-\eta L_{t-1,i}\right)} \exp\left(-\eta \ell_{t,i}\right) \right) \right)$ $= \frac{1}{n} \ln \left(\sum_{i=1}^{N} p_{t,i} \exp\left(-\eta \ell_{t,i}\right) \right) \qquad (\text{update step of } \boldsymbol{p}_t)$ $\leq \frac{1}{\eta} \ln \left(\sum_{i=1}^{N} p_{t,i} \left(1 - \eta \ell_{t,i} + \eta^2 \ell_{t,i}^2 \right) \right) \quad (\forall x \geq 0, e^{-x} \leq 1 - x + x^2)$ $= \frac{1}{\eta} \ln \left(1 - \eta \left\langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \right\rangle + \eta^2 \sum_{i=1}^N p_{t,i} \ell_{t,i}^2 \right) \right)$

Advanced Optimization (Fall 2023)

Proof of Hedge Regret Bound

Proof.
$$\Phi_t - \Phi_{t-1} = \frac{1}{\eta} \ln \left(\frac{\sum_{i=1}^N \exp(-\eta L_{t,i})}{\sum_{i=1}^N \exp(-\eta L_{t-1,i})} \right)$$

$$\leq -\langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle + \eta \sum_{i=1}^N p_{t,i} \boldsymbol{\ell}_{t,i}^2 \qquad (\ln(1+x) \leq x)$$

Summing over *t*, we have

$$\begin{split} \sum_{t=1}^{T} \langle \boldsymbol{p}_{t}, \boldsymbol{\ell}_{t} \rangle &\leq \Phi_{0} - \Phi_{T} + \eta \sum_{t=1}^{T} \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^{2} \qquad \Phi_{t} \triangleq \frac{1}{\eta} \ln \left(\sum_{i=1}^{N} \exp\left(-\eta L_{t,i}\right) \right) \\ &\leq \frac{\ln N}{\eta} - \frac{1}{\eta} \ln \left(\exp\left(-\eta L_{T,i^{\star}}\right) \right) + \eta \sum_{t=1}^{T} \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^{2} \\ &\leq \frac{\ln N}{\eta} + L_{T,i^{\star}} + \eta \sum_{t=1}^{T} \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^{2} \end{split}$$

Advanced Optimization (Fall 2023)

Proof of Hedge Regret Bound

Proof.

$$\sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle \leq \frac{\ln N}{\eta} + L_{T,i^{\star}} + \eta \sum_{t=1}^{T} \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^2$$

Rearranging the term gives

$$\begin{split} \sum_{t=1}^{T} \langle \boldsymbol{p}_{t}, \boldsymbol{\ell}_{t} \rangle - L_{T,i^{\star}} &\leq \frac{\ln N}{\eta} + \eta \sum_{t=1}^{T} \sum_{i=1}^{N} p_{t,i} \boldsymbol{\ell}_{t,i}^{2} \\ &\leq \frac{\ln N}{\eta} + \eta T \qquad (\boldsymbol{\ell}_{t,i} \leq 1) \end{split}$$

Thus, setting $\eta = \sqrt{\ln N/T}$ yields
 $\operatorname{Regret}_{T} &\leq \frac{\ln N}{\eta} + \eta T = 2\sqrt{T \ln N}. \end{split}$

Advanced Optimization (Fall 2023)

• As above, we have proved the regret bound for Hedge:

 $\operatorname{Regret}_T \le 2\sqrt{T \ln N}$

• A natural question: can we further improve the bound?

Theorem 2 (Lower Bound of PEA). *For any algorithm A, we have that*

$$\sup_{T,N} \max_{\ell_1,\ldots,\ell_T} \frac{\operatorname{Regret}_T}{\sqrt{T \ln N}} \geq \frac{1}{\sqrt{2}}.$$

Hedge achieves minimax optimal regret (up to a constant of $2\sqrt{2}$ *) for PEA.*

Theorem 2 (Lower Bound of PEA). *For any algorithm A, we have that*

$$\sup_{T,N} \max_{\boldsymbol{\ell}_1,\ldots,\boldsymbol{\ell}_T} \frac{\operatorname{Regret}_T}{\sqrt{T \ln N}} \geq \frac{1}{\sqrt{2}}.$$

Proof. We construct the 'hard' instance by randomization. Let \mathcal{D} be the uniform distribution over $\{0, 1\}$. We have

$$\max_{\boldsymbol{\ell}_{1},...,\boldsymbol{\ell}_{T}} \operatorname{Regret}_{T} \geq \mathbb{E}_{\boldsymbol{\ell}_{1},...,\boldsymbol{\ell}_{T}} \operatorname{Ind}_{\mathcal{D}^{N}} [\operatorname{Regret}_{T}]$$
(conditional expectation decomposition)
$$= \sum_{t=1}^{T} \mathbb{E}_{\boldsymbol{\ell}_{1},...,\boldsymbol{\ell}_{t-1}} \mathbb{E}_{\boldsymbol{\ell}_{t}} \left[\langle p_{t}, \boldsymbol{\ell}_{t} \rangle \mid \boldsymbol{\ell}_{t-1}, \ldots, \boldsymbol{\ell}_{1} \right] - \mathbb{E}_{\boldsymbol{\ell}_{1},...,\boldsymbol{\ell}_{T}} \left[\min_{i \in [N]} \sum_{t=1}^{T} \boldsymbol{\ell}_{t,i} \right]$$
$$= \sum_{t=1}^{T} \mathbb{E}_{\boldsymbol{\ell}_{1},...,\boldsymbol{\ell}_{t-1}} \langle p_{t}, \mathbb{E}_{\boldsymbol{\ell}_{t}} \left[\boldsymbol{\ell}_{t} \mid \boldsymbol{\ell}_{t-1}, \ldots, \boldsymbol{\ell}_{1} \right] \rangle - \mathbb{E}_{\boldsymbol{\ell}_{1},...,\boldsymbol{\ell}_{T}} \left[\min_{i \in [N]} \sum_{t=1}^{T} \boldsymbol{\ell}_{t,i} \right]$$

Advanced Optimization (Fall 2023)

Theorem 2 (Lower Bound of PEA). *For any algorithm A, we have that*

$$\sup_{T,N} \max_{\boldsymbol{\ell}_1,\ldots,\boldsymbol{\ell}_T} \frac{\operatorname{Regret}_T}{\sqrt{T \ln N}} \geq \frac{1}{\sqrt{2}}.$$

Proof.
$$\max_{\boldsymbol{\ell}_1,\ldots,\boldsymbol{\ell}_T} \operatorname{Regret}_T \geq \sum_{t=1}^T \mathbb{E}_{\boldsymbol{\ell}_1,\ldots,\boldsymbol{\ell}_{t-1}} \langle \boldsymbol{p}_t, \mathbb{E}_{\boldsymbol{\ell}_t} \left[\boldsymbol{\ell}_t \mid \boldsymbol{\ell}_{t-1},\ldots,\boldsymbol{\ell}_1 \right] \rangle - \mathbb{E}_{\boldsymbol{\ell}_1,\ldots,\boldsymbol{\ell}_T} \left[\min_{i \in [N]} \sum_{t=1}^T \ell_{t,i} \right]$$

$$= T/2 - \mathbb{E}_{\boldsymbol{\ell}_1, \dots, \boldsymbol{\ell}_T} \left[\min_{i \in [N]} \sum_{t=1}^T \ell_{t,i} \right] = \mathbb{E}_{\boldsymbol{\ell}_1, \dots, \boldsymbol{\ell}_T} \left[\max_{i \in [N]} \sum_{t=1}^T \left(\frac{1}{2} - \ell_{t,i} \right) \right]$$
$$= \frac{1}{2} \mathbb{E}_{\sigma_1, \dots, \sigma_T} \left[\max_{i \in [N]} \sum_{t=1}^T \sigma_{t,i} \right], \qquad (\ell_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{D} \text{ with } \mathcal{D} \text{ be the uniform distribution over } \{0, 1\})$$

(σ_t for $i \in [N], t \in [T]$ are i.i.d. Rademacher random variables)

Advanced Optimization (Fall 2023)

Theorem 2 (Lower Bound of PEA). *For any algorithm A, we have that*

 $\sup_{T,N} \max_{\ell_1,\ldots,\ell_T} \frac{\operatorname{Regret}_T}{\sqrt{T \ln N}} \geq \frac{1}{\sqrt{2}}.$

Proof.
$$\max_{\boldsymbol{\ell}_1,\ldots,\boldsymbol{\ell}_T} \operatorname{Regret}_T \geq \frac{1}{2} \mathbb{E}_{\sigma_1,\ldots,\sigma_T} \left[\max_{i \in [N]} \sum_{t=1}^T \sigma_{t,i} \right]$$

($\sigma_{t,i}$ for $i \in [N], t \in [T]$ are i.i.d. Rademacher random variables)

Using the result from probability theory (*Prediction, Learning, and Games,* Chapter 3.7) of Rademacher variables,

$$\square \qquad \lim_{T \to \infty} \lim_{N \to \infty} \frac{\mathbb{E}_{\sigma_1, \dots, \sigma_T} \left[\max_{i \in [N]} \sum_{t=1}^T \sigma_{t, i} \right]}{\sqrt{T \ln N}} = \sqrt{2}. \qquad \square$$

Advanced Optimization (Fall 2023)

Upper Bound and Lower Bound

Theorem 1. Suppose that $\forall t \in [T]$ and $i \in [N], 0 \leq \ell_{t,i} \leq 1$, then Hedge with *learning rate* η *guarantees*

$$\operatorname{Regret}_T \leq \frac{\ln N}{\eta} + \eta T = \mathcal{O}(\sqrt{T \log N}),$$

where the last equality is by setting η optimally as $\sqrt{(\ln N)/T}$.

Theorem 2 (Lower Bound of PEA). For any algorithm \mathcal{A} , we have that $\sup_{T,N} \max_{\ell_1,\ldots,\ell_T} \frac{\operatorname{Regret}_T}{\sqrt{T \ln N}} \geq \frac{1}{\sqrt{2}}.$

Advanced Optimization (Fall 2023)

Prediction with Expert Advice: history bits



AGGREGATING STRATEGIES

371

Volodimir G. Vovk^{*} Research Council for Cybernetics 40 ulitsa Vavilova, Moscow 117333, USSR

The following situation is considered. At each moment of discrete time a decision maker, who does not know the current state of Nature but knows all its past states, must make a decision. The decision together with the current state of Nature determines the loss of the decision maker. The performance of the decision maker is measured by his total loss. We suppose there is a pool of the decision maker's potential strategies one of which is believed to perform well, and construct an "aggregating" strategy for which the total loss is not much bigger than the total loss under strategies in the pool, whatever states of Nature. Our construction generalizes both the Weighted Majority Algorithm of N.Littlestone and M.K. Warmuth and the Bayesian rule.

NOTATION

 $\mathbb{N},\ \mathbb{Q}$ and \mathbb{R} stand for the sets of positive integers. rational numbers and real numbers respectively, \mathbb{B} symbolizes the set (0.1). We put

$$\mathbb{B}^{n} = \bigcup_{i \leq n} \mathbb{B}^{i}, \mathbb{B}^{\leq n} = \bigcup_{i \leq n} \mathbb{B}^{i},$$

The empty sequence is denoted by D. The notation for logarithms is in (natural), 1b (binary) and \log_{λ} (base λ). The integer part of a real number t is denoted by $\lfloor t \rfloor$. For $A \subseteq \mathbb{R}^2$, con A is the convex hull of A.

1. UNIFORM MATCHES

We are working within (the finite horizon variant of) A.P.Dawid's "prequential" (predictive sequential) framework (See (Dawid, 1980); in detail it is described in (Dawid, 1980). Nature and a decision maker function in discrete time $(c_1,\ldots,n-1)$. Nature sequentially finds itself in states s_0 .

 s_1,\ldots,s_{n-1} comprising the string $s=s_0s_1\ldots s_{n-1}.$ For simplicity we suppose $s\in\mathbb{B}^n$. At each moment i the decision maker does not know the current state s_i of Nature but knows

^MAddress for correspondence: 9-3-451 ulitsa Ramenki, Moscow 117607, USSR.

Volodimir G. Vovk Royal Holloway, University of London

Nick Littlestone and Manfred K. Warmuth. "The Weighted Majority Algorithm." FOCS 1989: 256-261. Volodimir G. Vovk. "Aggregating Strategies." COLT 1990: 371-383.

Prediction with Expert Advice: history bits

Yoav Freund

Robert Schapire

Goldel Prize 2003

This paper introduced AdaBoost, an adaptive algorithm to improve the accuracy of hypotheses in machine learning. The algorithm demonstrated novel possibilities in analyzing data and is a permanent contribution to science even beyond computer science. JOURNAL OF COMPUTER AND SYSTEM SCIENCES 55, 119–139 (1997) ARTICLE NO. SS971504

A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting*

Yoav Freund and Robert E. Schapire⁴

AT&T Labs, 180 Park Avenue, Florham Park, New Jersey 07932

Received December 19, 1996

In the first part of the paper we consider the problem of dynamically apportioning resources among a set of options in a worst-case on-line framework. The model we study can be interpreted as a broad, abstract extension of the well-studied on-line prediction model to a general decision-theoretic setting. We show that the multiplicative weightupdate Littlestone-Warmuth rule can be adapted to this model, yielding bounds that are slightly weaker in some cases, but applicable to a considerably more general class of learning problems. We show how the resulting learning algorithm can be applied to a variety of problems, including gambling, multiple-outcome prediction, repeated games, and prediction of points in \mathbb{R}^n . In the second part of the paper we apply the multiplicative weight-update technique to derive a new boosting algorithm. This boosting algorithm does not require any prior knowledge about the performance of the weak learning algorithm. We also study generalizations of the new boosting algorithm to the problem of learning functions whose range, rather than being binary, is an arbitrary finite set or a bounded segment of the real line. © 1997 Academic Press

converting a "weak" PAC learning algorithm that performs just slightly better than random guessing into one with arbitrarily high accuracy.

We formalize our *on-line allocation model* as follows. The allocation agent A has N options or *strategies* to choose from; we number these using the integers 1, ..., N. At each time step t = 1, 2, ..., T, the allocator A decides on a distribution \mathbf{p}^t over the strategies; that is $p_i^t \ge 0$ is the amount allocated to strategy *i*, and $\sum_{i=1}^{N} p_i^t = 1$. Each strategy *i* then suffers some *loss* ℓ_i^t which is determined by the (possibly adversarial) "environment." The loss suffered by A is then $\sum_{i=1}^{n} p_i^t \ell_i^t = \mathbf{p}^t \cdot \ell^t$, i.e., the average loss of the strategies with respect to A's chosen allocation rule. We call this loss function the *mixture loss*.

In this paper, we always assume that the loss suffered by any strategy is bounded so that, without loss of generality, $\ell_i^t \in [0, 1]$. Besides this condition, we make no assumptions

Reference: Y. Freund and R. Schapire. A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting. JCSS 1997.

PEA vs. OCO

At each round $t = 1, 2, \cdots$

Prediction with Expert Advice

- (1) the player first picks a weight p_t from a simplex Δ_N ;
- (2) and simultaneously environments pick an loss vector $\ell_t \in \mathbb{R}^N$;
- (3) the player suffers loss $f_t(\mathbf{p}_t) \triangleq \langle \mathbf{p}_t, \boldsymbol{\ell}_t \rangle$, observes $\boldsymbol{\ell}_t$ and updates the model.

require domain to be a simplex
$$\mathcal{X} = \Delta_N$$
 integrable linear loss $f_t(\mathbf{x}) \triangleq \langle \mathbf{x}, \boldsymbol{\ell}_t \rangle$ is a *special case* of OCO!

At each round $t = 1, 2, \cdots$

Online Convex Optimization

- (1) the player first picks a model $\mathbf{x}_t \in \mathcal{X}$;
- (2) and simultaneously environments pick an online function $f_t : \mathcal{X} \to \mathbb{R}$;
- (3) the player suffers loss $f_t(\mathbf{x}_t)$, observes f_t and updates the model.

Deploying OGD to PEA

• PEA is a special case of OCO:

Why not directly deploy OGD (proposed in last lecture) to address PEA?

Theorem 4 (Regret bound for OGD). Under Assumption 1, 2 and 3, online gradient descent (OGD) with step sizes $\eta_t = \frac{D}{G\sqrt{t}}$ for $t \in [T]$ guarantees:

$$\operatorname{Regret}_{T} = \sum_{t=1}^{T} f_{t}(\mathbf{x}_{t}) - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^{T} f_{t}(\mathbf{x}) \leq \frac{3}{2} GD\sqrt{T}.$$

Regret guarantee:
$$D = \max_{\mathbf{x}, \mathbf{y} \in \Delta_N} \|\mathbf{x} - \mathbf{y}\|_2 = \sqrt{2}$$
 $G = \max_{\boldsymbol{\ell}_t \in \mathbb{R}^N} \|\boldsymbol{\ell}_t\|_2 = \sqrt{N}$
 \implies Regret $T = \sum_{t=1}^T \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle - \min_{\boldsymbol{p} \in \Delta_N} \sum_{t=1}^T \langle \boldsymbol{p}, \boldsymbol{\ell}_t \rangle \le \mathcal{O}(\sqrt{TN})$

Advanced Optimization (Fall 2023)

Deploying OGD to PEA

• OGD for PEA Problem:

$$D = \max_{\mathbf{x}, \mathbf{y} \in \Delta_N} \|\mathbf{x} - \mathbf{y}\|_2 = \sqrt{2} \qquad G = \max_{\boldsymbol{\ell}_t \in \mathbb{R}^N} \|\boldsymbol{\ell}_t\|_2 = \sqrt{N}$$
$$\implies \text{Regret}_T = \sum_{t=1}^T \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle - \min_{\boldsymbol{p} \in \Delta_N} \sum_{t=1}^T \langle \boldsymbol{p}, \boldsymbol{\ell}_t \rangle \le \mathcal{O}(\sqrt{TN})$$

- A natural question: is the O(√TN) regret bound tight enough?
 recall that the lower bound of PEA is Ω(√T log N)
 - OGD is not optimal with respect to *N* (number of experts)

Deploying OGD to PEA

• PEA is a special case of OCO:

Why not directly deploy OGD (proposed in last lecture) to address PEA?

Theorem 4 (Regret bound for OGD). Under Assumption 1, 2 and 3, online gradient descent (OGD) with step sizes $\eta_t = \frac{D}{G\sqrt{t}}$ for $t \in [T]$ guarantees:

$$\operatorname{Regret}_{T} = \sum_{t=1}^{T} f_{t}(\mathbf{x}_{t}) - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^{T} f_{t}(\mathbf{x}) \leq \frac{3}{2} GD\sqrt{T}.$$

Regret guarantee:
$$D = \max_{\mathbf{x}, \mathbf{y} \in \Delta_N} \|\mathbf{x} - \mathbf{y}\|_2 = \sqrt{2}$$

 \longrightarrow Regret $T = \sum_{t=1}^T \langle \mathbf{p}_t, \boldsymbol{\ell}_t \rangle - \min_{\mathbf{p} \in \Delta_N} \sum_{t=1}^T \langle \mathbf{p}, \boldsymbol{\ell}_t \rangle \le \mathcal{O}(\sqrt{TN})$

Advanced Optimization (Fall 2023)

Why OGD Fails for PEA

• PEA has a special structure whereas general OCO doesn't have.

Convex Problem

Domain: convex set \mathcal{X}

Online function: convex function f_t

Lower Bound: $\Omega(GD\sqrt{T})$

Why OGD Fails for PEA

• Remember that for the general OCO, we linearized the function to analyze the first gradient descent lemma:

• So, linearized loss is not the essence, but the *simplex domain* of the PEA problem is worthy specifically considering.

Why OGD Fails for PEA?

• Recall that for general OCO, we update the model as follows:

General Online Convex Optimization				
OGD:	Proximal type update:			
$\mathbf{x}_{t+1} = \Pi_{\mathcal{X}} \left[\mathbf{x}_t - \eta_t \nabla f(\mathbf{x}_t) \right]$	$\mathbf{x}_{t+1} = \underset{\mathbf{x}\in\mathcal{X}}{\arg\min} \left\{ \langle \mathbf{x}, \eta_t \nabla f_t(\mathbf{x}_t) \rangle + \frac{1}{2} \left\ \mathbf{x} - \mathbf{x}_t \right\ _2^2 \right\}$			

• In PEA, is it proper to use 2-norm (ball) to measure distance?

Advanced Optimization (Fall 2023)

Proximal View

• Recall that for general OCO, we update the model as follows:

General Online Convex Optimization				
OGD:	Proximal type update:			
$\mathbf{x}_{t+1} = \Pi_{\mathcal{X}} \left[\mathbf{x}_t - \eta_t \nabla f(\mathbf{x}_t) \right]$	$\mathbf{x}_{t+1} = \underset{\mathbf{x}\in\mathcal{X}}{\arg\min} \left\{ \langle \mathbf{x}, \eta_t \nabla f_t(\mathbf{x}_t) \rangle + \frac{1}{2} \left\ \mathbf{x} - \mathbf{x}_t \right\ _2^2 \right\}$			

• In PEA, is it proper to use 2-norm (ball) to measure distance?

Proximal View

We need to find an alternative distance measure for the *special structure* in PEA.

• Intuitively, for Euclidean space, 2-norm is the most natural measure:

$$\|\mathbf{x} - \mathbf{y}\|_2^2$$

- For PEA problem
 - the decision can be viewed as a distribution within the simplex
 - for two distributions *P* and *Q*, **KL divergence** is a natural measure:

$$\mathrm{KL}(P \| Q) = \sum_{x \in \mathcal{X}} P(x) \log \left(\frac{P(x)}{Q(x)} \right)$$

Reinvent Hedge Algorithm

Theorem 3. Consider $f_t(\mathbf{p}) = \langle \mathbf{p}, \boldsymbol{\ell}_t \rangle$. An online learning algorithm that updates the model following

$$\boldsymbol{p}_{t+1} = \underset{\boldsymbol{p} \in \Delta_N}{\operatorname{arg min}} \left\{ \eta \langle \boldsymbol{p}, \nabla f_t(\boldsymbol{p}_t) \rangle + \frac{\mathrm{KL}(\boldsymbol{p} \| \boldsymbol{p}_t)}{\mathrm{KL}(\boldsymbol{p} \| \boldsymbol{p}_t)} \right\}$$

is equal to Hedge update, i.e.,

$$p_{t+1,i} \propto p_{t,i} \exp(-\eta \ell_{t,i})$$
 for all $i \in [N]$.

$$\begin{array}{ll} \textit{Proof.} \quad \textit{p}_{t+1} = \mathop{\arg\min}_{p \in \Delta_N} \eta \langle \textit{p}, \nabla f_t(\textit{p}_t) \rangle + \textit{KL}(\textit{p} || \textit{p}_t) \\ = \mathop{\arg\min}_{p \in \Delta_N} \eta \langle \textit{p}, \nabla f_t(\textit{p}_t) \rangle - \sum_{i=1}^N p_i \ln\left(\frac{p_{t,i}}{p_i}\right) & \text{(definition of KL divergence)} \\ F(\textit{p}) & \Box \end{array}$$

Advanced Optimization (Fall 2023)

Reinvent Hedge Algorithm

• Proximal update rule for OGD:

$$\mathbf{x}_{t+1} = \underset{\mathbf{x}\in\mathcal{X}}{\operatorname{arg min}} \left\{ \eta_t \langle \mathbf{x}, \nabla f_t(\mathbf{x}_t) \rangle + \frac{1}{2} \left\| \mathbf{x} - \mathbf{x}_t \right\|_2^2 \right\}$$

• Proximal update rule for Hedge:

$$\mathbf{x}_{t+1} = \underset{\mathbf{x}\in\mathcal{X}}{\operatorname{arg\,min}} \left\{ \eta_t \langle \mathbf{x}, \nabla f_t(\mathbf{x}_t) \rangle + \frac{\operatorname{KL}(\mathbf{x} \| \mathbf{x}_t)}{\operatorname{KL}(\mathbf{x} \| \mathbf{x}_t)} \right\}$$

• More possibility: changing the distance measure to a more general form using *Bregman divergence*

$$\mathbf{x}_{t+1} = \underset{\mathbf{x}\in\mathcal{X}}{\operatorname{arg min}} \left\{ \eta_t \langle \mathbf{x}, \nabla f_t(\mathbf{x}_t) \rangle + \mathcal{D}_{\psi}(\mathbf{x}, \mathbf{x}_t) \right\}$$

Bregman Divergence

Definition 1 (Bregman Divergence). Let ψ be a strongly convex and differentiable function over a convex set \mathcal{X} , then for any $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, the bregman divergence \mathcal{D}_{ψ} associated to ψ is defined as

$$\mathcal{D}_{\psi}(\mathbf{x}, \mathbf{y}) = \psi(\mathbf{x}) - \psi(\mathbf{y}) - \langle \nabla \psi(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle.$$

- Bregman divergence measures the difference of a function and its linear approximation.
- Using second-order Taylor expansion, we know

$$\mathcal{D}_{\psi}(\mathbf{x}, \mathbf{y}) = \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_{\nabla^2 \psi(\boldsymbol{\xi})}^2$$

for some $\boldsymbol{\xi} \in [\mathbf{x}, \mathbf{y}]$.

Bregman Divergence

Definition 1 (Bregman Divergence). Let ψ be a strongly convex and differentiable function over a convex set \mathcal{X} , then for any $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, the bregman divergence \mathcal{D}_{ψ} associated to ψ is defined as

 $\mathcal{D}_{\psi}(\mathbf{x}, \mathbf{y}) = \psi(\mathbf{x}) - \psi(\mathbf{y}) - \langle \nabla \psi(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle.$

Table 1: Choice of $\psi(\cdot)$ and the corresponding Bregman divergence.

	$\psi(\mathbf{x})$	$\mathcal{D}_\psi(\mathbf{x},\mathbf{y})$
Squared L_2 -distance	$\ \mathbf{x}\ _2^2$	$\ \mathbf{x} - \mathbf{y}\ _2^2$
Mahalanobis distance	$\left\ \mathbf{x} ight\ _{A}^{2}$	$\left\ \mathbf{x} - \mathbf{y} ight\ _{A}^{2}$
Negative entropy	$\sum_i x_i \log x_i$	$KL(\mathbf{x}\ \mathbf{y})$

Advanced Optimization (Fall 2023)

Online Mirror Descent

Online Mirror Descent

At each round $t = 1, 2, \cdots$

$$\mathbf{x}_{t+1} = \underset{\mathbf{x}\in\mathcal{X}}{\arg\min} \left\{ \eta_t \langle \mathbf{x}, \nabla f_t(\mathbf{x}_t) \rangle + \mathcal{D}_{\psi}(\mathbf{x}, \mathbf{x}_t) \right\}$$

where $\mathcal{D}_{\psi}(\mathbf{x}, \mathbf{y}) = \psi(\mathbf{x}) - \psi(\mathbf{y}) - \langle \nabla \psi(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle$ is the Bregman divergence.

- $\psi(\cdot)$ is a required to be strongly convex and differentiable over a convex set \mathcal{X} .
- Strong convexity of ψ will ensure the uniqueness of the minimization problem, and actually we further need some analytical assumptions (see later mirror map definition) to ensure the solutions' feasibility in domain \mathcal{X} .

Online Mirror Descent

• So we can unify OGD and Hedge from the same view of OMD.

$$\mathbf{x}_{t+1} = \underset{\mathbf{x}\in\mathcal{X}}{\operatorname{arg min}} \left\{ \eta_t \langle \mathbf{x}, \nabla f_t(\mathbf{x}_t) \rangle + \mathcal{D}_{\psi}(\mathbf{x}, \mathbf{x}_t) \right\}$$

Algo.	OMD/proximal form	$\psi(\cdot)$	η_t	Regret_T
OGD	$\mathbf{x}_{t+1} = \underset{\mathbf{x} \in \mathcal{X}}{\operatorname{arg min}} \left\{ \eta_t \langle \mathbf{x}, \nabla f_t(\mathbf{x}_t) \rangle + \frac{1}{2} \left\ \mathbf{x} - \mathbf{x}_t \right\ _2^2 \right\}$	$\ \mathbf{x}\ _2^2$	$\frac{1}{\sqrt{t}}$	$\mathcal{O}(\sqrt{T})$
Hedge	$\mathbf{x}_{t+1} = \underset{\mathbf{x} \in \Delta_N}{\operatorname{arg min}} \left\{ \eta_t \langle \mathbf{x}, \nabla f_t(\mathbf{x}_t) \rangle + \frac{\operatorname{KL}(\mathbf{x} \ \mathbf{x}_t)}{\operatorname{KL}(\mathbf{x} \ \mathbf{x}_t)} \right\}$	$\sum_{i=1}^{N} x_i \log x_i$	$\sqrt{\frac{\ln N}{T}}$	$\mathcal{O}(\sqrt{T\log N})$

 \Rightarrow We will give the (unified) regret analysis in the next lecture.

Advanced Optimization (Fall 2023)

Why is PEA useful?

• Prediction with Expert Advice is essentially a **meta-algorithm** for combining different experts, and the "expert" can be interpreted as any learning model with a particular kind of expertise.

• It is used in a variety of algorithmic design.

THEORY OF COMPUTING, Volume 8 (2012), pp. 121-164 www.theoryofcomputing.org

RESEARCH SURVEY

The Multiplicative Weights Update Method: A Meta-Algorithm and Applications

Sanjeev Arora* Elad Hazan

Satyen Kale

Received: July 22, 2008; revised: July 2, 2011; published: May 1, 2012.

Abstract: Algorithms in varied fields use the idea of maintaining a distribution over a certain set and use the multiplicative update rule to iteratively change these weights. Their analyses are usually very similar and rely on an exponential potential function.

In this survey we present a simple meta-algorithm that unifies many of these disparate algorithms and derives them as simple instantiations of the meta-algorithm. We feel that since this meta-algorithm and its analysis are so simple, and its applications so broad, it should be a standard part of algorithms courses, like "divide and conquer."

ACM Classification: G.1.6

AMS Classification: 68Q25

Key words and phrases: algorithms, game theory, machine learning

1 Introduction

The Multiplicative Weights (MW) method is a simple idea which has been repeatedly discovered in fields as diverse as Machine Learning, Optimization, and Game Theory. The setting for this algorithm is the following. A decision maker has a choice of n decisions, and needs to repeatedly make a decision and obtain an associated payoff. The decision maker's goal, in the long run, is to achieve a total payoff which is comparable to the payoff of that fixed decision that maximizes the total payoff with the benefit of

*This project was supported by David and Lucite Packard Fellowship and NSF grants MSPA-MCS 0528414 and CCR-0205594

@ 2012 Sanjeev Arora, Elad Hazan and Satyen Kale G Licensed under a Creative Commons Attribution License

DOI: 10.4086/toc.2012.v008a006

- Applications
 - Learning a linear classifier: the Winnow algorithm Solving zero-sum games approximately Plotkin, Shmoys, Tardos framework for packing/covering LPs Approximating multicommodity flow problems O(log n)-approximation for many NP-hard problems Learning theory and boosting Hard-core sets and the XOR Lemma Hannan's algorithm and multiplicative weights Online convex optimization Other applications Design of competitive online algorithms

The multiplicative weights update method: a meta-algorithm and applications. S Arora, E Hazan, S Kale. Theory of Computing, 2012

More Results on PEA

Prediction, Learning and Games. Nicolò Cesa-Bianchi and Gabor Lugosi. Cambridge University Press, 2006.

Nicolò Cesa-Bianchi

Gabor Lugosi

Advanced Optimization (Fall 2023)

Summary

Advanced Optimization (Fall 2023)