Prediction, Learning and Games - Chapter 4
Randomized prediction

Walid Krichene

November 14, 2013

Experts framework

Recall

- On iteration t: experts reveal their advice $f_{i,t}$

Experts framework

Recall

- On iteration t: experts reveal their advice $f_{i,t}$
- forecaster makes a decision $\hat{p}_t = \sum_{i=1}^{N} w_{i,t} f_{i,t}$

Experts framework

Recall

- On iteration t: experts reveal their advice $f_{i,t}$
- forecaster makes a decision $\hat{p}_t = \sum_{i=1}^{N} w_{i,t} f_{i,t}$
- the losses are revealed $\ell(f_{i,t}, y_t)$ and $\ell(\hat{p}_t, y_t)$

Experts framework

Recall

- On iteration t: experts reveal their advice $f_{i,t}$
- forecaster makes a decision $\hat{p}_t = \sum_{i=1}^{N} w_{i,t} f_{i,t}$
- the losses are revealed $\ell(f_{i,t}, y_t)$ and $\ell(\hat{p}_t, y_t)$
- forecaster updates weights $w_{i,t+1}$

Experts framework

Recall

- On iteration t: experts reveal their advice $f_{i,t}$
- forecaster makes a decision $\hat{p}_t = \sum_{i=1}^{N} w_{i,t} f_{i,t}$
- the losses are revealed $\ell(f_{i,t}, y_t)$ and $\ell(\hat{p}_t, y_t)$
- forecaster updates weights $w_{i,t+1}$
- $\ell(\cdot, y_t)$ is convex

Experts framework

Recall

- On iteration t: experts reveal their advice $f_{i,t}$
- forecaster makes a decision $\hat{p}_t = \sum_{i=1}^{N} w_{i,t} f_{i,t}$
- the losses are revealed $\ell(f_{i,t}, y_t)$ and $\ell(\hat{p}_t, y_t)$
- forecaster updates weights $w_{i,t+1}$
- $\ell(\cdot, y_t)$ is convex

Definition

Regret

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

Goal: $\frac{R_T}{T} = o(1)$

Multiplicative weight algorithms

Definition

Hedge algorithm

$$w_{i,t+1} \propto w_{i,t} \exp(-\gamma \ell(f_{i,t}, y_t))$$

Multiplicative weight algorithms

Definition

Hedge algorithm

$$w_{i,t+1} \propto w_{i,t} \exp(-\gamma \ell(f_{i,t}, y_t))$$

- Average regret $\frac{R_T}{T} \leq \frac{\ln N}{\gamma T} + \frac{\gamma}{8}$

Multiplicative weight algorithms

Definition

Hedge algorithm

$$w_{i,t+1} \propto w_{i,t} \exp(-\gamma \ell(f_{i,t}, y_t))$$

- Average regret $\frac{R_T}{T} \leq \frac{\ln N}{\gamma T} + \frac{\gamma}{8}$
- taking $\gamma_t = O(\sqrt{\frac{\ln N}{T}})$ yields $\frac{R_T}{T} \leq O(\frac{\ln N}{T}) + O(\frac{\ln N}{T})$

Multiplicative weight algorithms

Definition

Hedge algorithm

$$w_{i,t+1} \propto w_{i,t} \exp(-\gamma \ell(f_{i,t}, y_t))$$

- Average regret $\frac{R_T}{T} \leq \frac{\ln N}{\gamma T} + \frac{\gamma}{8}$
- taking $\gamma_t = O(\sqrt{\frac{\ln N}{T}})$ yields $\frac{R_T}{T} \leq O(\frac{\ln N}{T}) + O(\frac{\ln N}{T})$
- small losses: $\frac{R_T}{T} \leq \frac{1}{T} \left(\frac{\gamma}{1-e^{-\gamma}} - 1 \right) L_i^* + \frac{1}{T} \frac{\ln N}{1-e^{-\gamma}}$

Motivation for randomization

- What if the decision set is non-convex? E.g. $\{0, 1\}$
- Also: any deterministic algorithm can incur $\Omega(n)$ loss.
- solution (to both): Randomize

Setting

- Actions $\{1, 2, \ldots, N\} = [N]$

Setting

- Actions $\{1, 2, \ldots, N\} = [N]$
- forecaster maintains a probability distribution $(p_{i,t})_{i \in [N]}$

Setting

- Actions $\{1, 2, \ldots, N\} = [N]$
- forecaster maintains a probability distribution $(p_{i,t})_{i \in [N]}$
- randomly picks an action $I_t \sim p_t$

Setting

- Actions $\{1, 2, \ldots, N\} = [N]$
- forecaster maintains a probability distribution $(p_{i,t})_{i \in [N]}$
- randomly picks an action $I_t \sim p_t$
- losses are revealed $\ell(i, y_t)$

Setting

- Actions $\{1, 2, \ldots, N\} = [N]$
- forecaster maintains a probability distribution $(p_{i,t})_{i \in [N]}$
- randomly picks an action $I_t \sim p_t$
- losses are revealed $\ell(i, y_t)$
- sequence $y_1, \ldots, y_T$ can be fixed a priori (oblivious opponent) or can depend on player's decisions.

Regret

Definition

Expected loss (conditioned on past plays) $\bar{\ell}(p_t, Y_t) = \sum_{i=1}^{N} \ell(i, Y_t) p_{i,t}$

Definition

Expected regret

$$\bar{R}_{i,T} = \sum_{t=1}^{T} \bar{\ell}(p_t, Y_t) - \ell(i, Y_t)$$

Regret

Regret

- Note: p_{t+1} only depends on p_t and $\ell(i, Y_t)$ (not on forecaster randomization)

Regret

- Note: p_{t+1} only depends on p_t and $\ell(i, Y_t)$ (not on forecaster randomization)
- Strategy: bound expected regret

$$\frac{\bar{R}_T}{T} \leq B(T)$$

then with high probability ($\geq 1 - \delta$)

$$\frac{R_T}{T} \leq B(T) + \sqrt{\frac{-\ln \delta}{T}}$$

Back to experts

Can think of this as

- Strategy space is simplex on actions: Δ^N

Back to experts

Can think of this as

- Strategy space is simplex on actions: Δ^N
- Every expert has constant advice: $f_{i,t}$ is the i-th vertex of the simplex

Back to experts

Can think of this as

- Strategy space is simplex on actions: Δ^N
- Every expert has constant advice: $f_{i,t}$ is the i-th vertex of the simplex
- decision of forecaster is convex combination of vertices

Back to experts

Can think of this as

- Strategy space is simplex on actions: Δ^N
- Every expert has constant advice: $f_{i,t}$ is the i-th vertex of the simplex
- decision of forecaster is convex combination of vertices
- loss function is $\ell(\cdot, Y_t)$ is linear: (expected) loss incurred by forecaster is

$$\sum_i p_{t,i}\ell(f_{i,t}, Y_t) = \ell(\sum_i p_{t,i} f_{i,t}, Y_t) = \ell(\hat{p}_t, Y_t)$$

write the expected loss as $\langle \ell_t, p_t \rangle$ where $\ell_{t,i} = \ell(f_{i,t}, Y_t)$.

Hedge algorithm

Regularized greedy algorithm

Hedge is solution to

$$p_{t+1} = \arg\min_{p \in \Delta^N} \langle \ell_t, p \rangle + \frac{1}{\gamma} D_{KL}(p || p_t)$$

where D_{KL} is the Kullback-Leibler divergence $D_{KL}(p, q) = \sum_i \ln \frac{p_i}{q_i} p_i$

Hedge algorithm

Regularized greedy algorithm

Hedge is solution to

$$p_{t+1} = \arg\min_{p \in \Delta^{\boldsymbol{N}}} \langle \ell_t, p \rangle + \frac{1}{\gamma} D_{KL}(p || p_t)$$

where D_{KL} is the Kullback-Leibler divergence $D_{KL}(p, q) = \sum_i \ln \frac{p_i}{q_i} p_i$

Also

$$p_{t+1} = \arg\min_{p \in \Delta^{\boldsymbol{N}}} \langle L_t, p \rangle - \frac{1}{\gamma} H(p)$$

where H is the entropy $H(p) = -\sum_i p_i \ln p_i$

- Connection with stochastic optimization (last week)
- Greedy strategy called fictitious play

Follow the perturbed leader

- Regularized problem called: follow the perturbed leader

Follow the perturbed leader

- Regularized problem called: follow the perturbed leader
- They prove it in a more general case

Theorem (Theorem 4.2)

$$I_t = \arg\min_i L_{i,t-1} + Z_{i,t}$$

Then

$$\frac{R_T}{T} \leq \frac{1}{T}\left(\mathrm{E}\max_i Z_{i,1} + \mathrm{E}\max_i -Z_{i,1} + \sum_i \int F_t(z)(f_Z(z) - f_Z(z - \ell_t))\right)$$

where $F_t(z) = \ell(i_t(z), y_t)$ *and* f_Z *is the density of* Z.

Internal regret

- regret analyzed so far: external regret. Compare your cumulative loss to cumulative loss of a single expert / action

Internal regret

- regret analyzed so far: external regret. Compare your cumulative loss to cumulative loss of a single expert / action
- internal regret: Swap all instances of action i with j

Definition (Internal regret)

$$R_{i,j,T} = \sum_{t=1}^{T} p_{i,t}(\ell(i, Y_t) - \ell(j, Y_t))$$

Internal regret

- regret analyzed so far: external regret. Compare your cumulative loss to cumulative loss of a single expert / action
- internal regret: Swap all instances of action i with j

Definition (Internal regret)

$$R_{i,j,T} = \sum_{t=1}^{T} p_{i,t}(\ell(i, Y_t) - \ell(j, Y_t))$$

internal regret

$$\max_{i,j} R_{i,j,T}$$

Internal regret

- In a sense, stronger than external regret:

Internal regret

- In a sense, stronger than external regret:

$$\sum_i R_{i,j,T} = \sum_{t=1}^{T} \sum_i p_{t,i}(\ell_{t,i} - \ell_{t,j})$$
$$= \sum_{t=1}^{T} \langle p_t, \ell_t \rangle - \ell_{t,j}$$
$$= R_{j,T}$$

so if $\max_j R_{j,T} \leq N \max_{i,j} R_{i,j,T}$

Internal regret

- In a sense, stronger than external regret:

$$\sum_i R_{i,j,T} = \sum_{t=1}^{T} \sum_i p_{t,i}(\ell_{t,i} - \ell_{t,j})$$
$$= \sum_{t=1}^{T} \langle p_t, \ell_t \rangle - \ell_{t,j}$$
$$= R_{j,T}$$

so if $\max_j R_{j,T} \leq N \max_{i,j} R_{i,j,T}$

- Weighted average forecaster has large internal regret. $R_T = \Omega(T)$

Internal regret

- In a sense, stronger than external regret:

$$\sum_i R_{i,j,T} = \sum_{t=1}^{T} \sum_i p_{t,i}(\ell_{t,i} - \ell_{t,j})$$
$$= \sum_{t=1}^{T} \langle p_t, \ell_t \rangle - \ell_{t,j}$$
$$= R_{j,T}$$

  so if $\max_j R_{j,T} \leq N \max_{i,j} R_{i,j,T}$
- Weighted average forecaster has large internal regret. $R_T = \Omega(T)$
- But, can adapt to have small internal regret

Minimizing internal regret

Minimizing internal regret

- define modified strategies $p_t^{i \to j}$

$$p_{t,i}^{i \to j} \leftarrow 0$$
$$p_{t,j}^{i \to j} \leftarrow p_{t,i} + p_{t,j}$$

Minimizing internal regret

Minimizing internal regret

- define modified strategies $p_t^{i \to j}$

$$p_{t,i}^{i \to j} \leftarrow 0$$
$$p_{t,j}^{i \to j} \leftarrow p_{t,i} + p_{t,j}$$

- Apply an algorithm that minimizes external regret to a set of new experts $\{i \to j, i \neq j\}$. $O(N^2)$ experts for the new algorithm.
- An action is $p_t^{i \to j}$.
- Results in a sequence of meta-strategies $\mu_t \in \Delta^{N^2}$

Minimizing internal regret

Minimizing internal regret

- define modified strategies $p_t^{i\to j}$

$$p_{t,i}^{i\to j} \leftarrow 0$$
$$p_{t,j}^{i\to j} \leftarrow p_{t,i} + p_{t,j}$$

- Apply an algorithm that minimizes external regret to a set of new experts $\{i \to j, i \neq j\}$. $O(N^2)$ experts for the new algorithm.
- An action is $p_t^{i\to j}$.
- Results in a sequence of meta-strategies $\mu_t \in \Delta^{N^2}$
- The fixed point $p_t = \sum_{(i,j)} p_t^{i\to j} \mu_{t,(i,j)}$ minimizes internal regret

Minimizing internal regret

Minimizing internal regret

- define modified strategies $p_t^{i\to j}$

$$p_{t,i}^{i\to j} \leftarrow 0$$
$$p_{t,j}^{i\to j} \leftarrow p_{t,i} + p_{t,j}$$

- Apply an algorithm that minimizes external regret to a set of new experts $\{i \to j, i \neq j\}$. $O(N^2)$ experts for the new algorithm.
- An action is $p_t^{i\to j}$.
- Results in a sequence of meta-strategies $\mu_t \in \Delta^{N^2}$
- The fixed point $p_t = \sum_{(i,j)} p_t^{i\to j} \mu_{t,(i,j)}$ minimizes internal regret
- can be computed using Gaussian elimination. (write $p_t = A(\mu_t) p_t$)

Generalized regret

- experts react to the prediction, $f_{i,t}(I_t)$
- also, activation function $A_{i,t}(k)$

Generalized regret

- experts react to the prediction, $f_{i,t}(I_t)$
- also, activation function $A_{i,t}(k)$

Definition (Generalized regret)

$$r_{i,t} = \sum_{k=1}^{N} p_{t,k} A_{i,t}(k)(\ell(k, Y_t) - \ell(f_{i,t}(k), Y_t))$$

Generalized regret

- experts react to the prediction, $f_{i,t}(I_t)$
- also, activation function $A_{i,t}(k)$

Definition (Generalized regret)

$$r_{i,t} = \sum_{k=1}^{N} p_{t,k} A_{i,t}(k)(\ell(k, Y_t) - \ell(f_{i,t}(k), Y_t))$$

- External regret is a special case
  Set $A_{i,t}(k) = 1$ identically, and $f_{i,t}(k) = i$

Generalized regret

- experts react to the prediction, $f_{i,t}(I_t)$
- also, activation function $A_{i,t}(k)$

Definition (Generalized regret)

$$r_{i,t} = \sum_{k=1}^{N} p_{t,k} A_{i,t}(k)(\ell(k, Y_t) - \ell(f_{i,t}(k), Y_t))$$

- External regret is a special case
  Set $A_{i,t}(k) = 1$ identically, and $f_{i,t}(k) = i$
- Internal regret is a special case
  Consider experts $\{(i, j), i \neq j)\}$. Set

$$f_{(i,j),t}(k) = \begin{cases} k & \text{if } k \neq i \\ j & \text{otherwise} \end{cases}$$

Generalized regret

- Given a potential ϕ, there exists a forecaster that satisfies the Blackwell condition. (Theorem 4.3)

Generalized regret

- Given a potential ϕ, there exists a forecaster that satisfies the Blackwell condition. (Theorem 4.3)
- Recall

Blackwell condition

$$\langle r_t, \nabla\phi(R_{t-1})\rangle \leq 0$$

Generalized regret

- Given a potential ϕ, there exists a forecaster that satisfies the Blackwell condition. (Theorem 4.3)
- Recall

Blackwell condition

$$\langle r_t, \nabla\phi(R_{t-1})\rangle \leq 0$$

- From the Blackwell condition one can derive a bound on regret

Next week

- Chapter 5 (Efficient forecasting for special classes of Experts)
- or Chapter 6 (Limited information: multi-armed bandit versions)