Information incomplète et regret interne en prédiction de suites individuelles

Gilles Stoltz

Plan de l'exposé

1 – Le modèle de la prédiction de suites individuelles

Plan de l'exposé

1 – Le modèle de la prédiction de suites individuelles

2 – Améliorations de la borne générale $\sqrt{n}$ sur le regret en des bornes dépendant des données

Plan de l'exposé

1 – Le modèle de la prédiction de suites individuelles

2 – Améliorations de la borne générale $\sqrt{n}$ sur le regret en des bornes dépendant des données

3 – Un exemple de prédiction avec information incomplète : les jeux avec signaux

Plan de l'exposé

Individual Sequences

The EWA strategy

Incomplete information

Internal regret

1 – Le modèle de la prédiction de suites individuelles

2 – Améliorations de la borne générale $\sqrt{n}$ sur le regret en des bornes dépendant des données

3 – Un exemple de prédiction avec information incomplète : les jeux avec signaux

4 – Les fondements de la théorie de la minimisation du regret interne

Individual Sequences

Description and aim

A decision-maker has to predict a sequence $y_1, y_2, \ldots$ of elements from an outcome space $\mathcal{Y}$.

Description and aim

A decision-maker has to predict a sequence $y_1, y_2, \ldots$ of elements from an outcome space $\mathcal{Y}$.

His predictions $\widehat{p}_1, \widehat{p}_2, \ldots$ belong to a prediction space $\mathcal{X}$.

Description and aim

A decision-maker has to predict a sequence $y_1, y_2, \ldots$ of elements from an outcome space $\mathcal{Y}$.

His predictions $\widehat{p}_1, \widehat{p}_2, \ldots$ belong to a prediction space $\mathcal{X}$.

The observations

– (and the predictions) are made in a sequential fashion,

Description and aim

A decision-maker has to predict a sequence $y_1, y_2, \ldots$ of elements from an outcome space $\mathcal{Y}$.

His predictions $\widehat{p}_1, \widehat{p}_2, \ldots$ belong to a prediction space $\mathcal{X}$.

The observations

- (and the predictions) are made in a sequential fashion,
- rely on no stochastic model,

Description and aim

A decision-maker has to predict a sequence $y_1, y_2, \ldots$ of elements from an outcome space $\mathcal{Y}$.

His predictions $\widehat{p}_1, \widehat{p}_2, \ldots$ belong to a prediction space $\mathcal{X}$.

The observations

- (and the predictions) are made in a sequential fashion,
- rely on no stochastic model,

that is, the observations are given by individual sequences.

Description and aim

A decision-maker has to predict a sequence $y_1, y_2, \ldots$ of elements from an outcome space $\mathcal{Y}$.

His predictions $\widehat{p}_1, \widehat{p}_2, \ldots$ belong to a prediction space $\mathcal{X}$.

The observations

– (and the predictions) are made in a sequential fashion,

– rely on no stochastic model,

that is, the observations are given by individual sequences.

Our aim is, for all time rounds $t = 1, 2, \ldots$, to predict y_t based on $y_1^{t-1} = (y_1, \ldots, y_{t-1})$.

Description and aim

A decision-maker has to predict a sequence $y_1, y_2, \ldots$ of elements from an outcome space $\mathcal{Y}$.

His predictions $\widehat{p}_1, \widehat{p}_2, \ldots$ belong to a prediction space $\mathcal{X}$.

The observations

- (and the predictions) are made in a sequential fashion,
- rely on no stochastic model,

that is, the observations are given by individual sequences.

Our aim is, for all time rounds $t = 1, 2, \ldots$, to predict y_t based on $y_1^{t-1} = (y_1, \ldots, y_{t-1})$.

The observation $y_t \in \mathcal{Y}$ is then revealed and has to be compared to the prediction $\widehat{p}_t \in \mathcal{X}$.

Description and aim

A decision-maker has to predict a sequence $y_1, y_2, \ldots$ of elements from an outcome space $\mathcal{Y}$.

His predictions $\widehat{p}_1, \widehat{p}_2, \ldots$ belong to a prediction space $\mathcal{X}$.

The observations

- (and the predictions) are made in a sequential fashion,
- rely on no stochastic model,

that is, the observations are given by individual sequences.

Our aim is, for all time rounds $t = 1, 2, \ldots$, to predict y_t based on $y_1^{t-1} = (y_1, \ldots, y_{t-1})$.

The observation $y_t \in \mathcal{Y}$ is then revealed and has to be compared to the prediction $\widehat{p}_t \in \mathcal{X}$.

Example. Weather forecasting, $\mathcal{Y} = \{0, 1\}$ and $\mathcal{X} = [0, 1]$.

Prediction with expert advice

Finitely many reference forecasters, called experts, are usually considered. They are indexed by $j = 1, \ldots, N$.

Prediction with expert advice

Finitely many reference forecasters, called experts, are usually considered. They are indexed by $j = 1, \ldots, N$.

For all time rounds $t = 1, 2, \ldots$, they also form predictions $f_{j,t} = f_{j,t}(y_1^{t-1}) \in \mathcal{X}$, often referred to as advice.

Prediction with expert advice

Finitely many reference forecasters, called experts, are usually considered. They are indexed by $j = 1, \ldots, N$.

For all time rounds $t = 1, 2, \ldots$, they also form predictions $f_{j,t} = f_{j,t}(y_1^{t-1}) \in \mathcal{X}$, often referred to as advice.

We now predict y_t based on y_1^{t-1} and on the advice $f_{j,t}$.

Prediction with expert advice

Finitely many reference forecasters, called experts, are usually considered. They are indexed by $j = 1, \ldots, N$.

For all time rounds $t = 1, 2, \ldots$, they also form predictions $f_{j,t} = f_{j,t}(y_1^{t-1}) \in \mathcal{X}$, often referred to as advice.

We now predict y_t based on y_1^{t-1} and on the advice $f_{j,s}$, $s \leqslant t$.

Prediction with expert advice

Finitely many reference forecasters, called experts, are usually considered. They are indexed by $j = 1, \ldots, N$.

For all time rounds $t = 1, 2, \ldots$, they also form predictions $f_{j,t} = f_{j,t}(y_1^{t-1}) \in \mathcal{X}$, often referred to as advice.

We now predict y_t based on y_1^{t-1} and on the advice $f_{j,s}$, $s \leqslant t$.

Our goal is to perform almost as well as the best expert.

Note that the best expert may only be determined in hindsight, whereas we are asked to predict on-line.

Prediction with expert advice

Finitely many reference forecasters, called experts, are usually considered. They are indexed by $j = 1, \ldots, N$.

For all time rounds $t = 1, 2, \ldots$, they also form predictions $f_{j,t} = f_{j,t}(y_1^{t-1}) \in \mathcal{X}$, often referred to as advice.

We now predict y_t based on y_1^{t-1} and on the advice $f_{j,s}$, $s \leqslant t$.

Our goal is to perform almost as well as the best expert.

Note that the best expert may only be determined in hindsight, whereas we are asked to predict on-line.

To make the notion of best expert mathematically precise, we introduce a loss function $\ell : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$.

Prediction with expert advice

Finitely many reference forecasters, called experts, are usually considered. They are indexed by $j = 1, \ldots, N$.

For all time rounds $t = 1, 2, \ldots$, they also form predictions $f_{j,t} = f_{j,t}(y_1^{t-1}) \in \mathcal{X}$, often referred to as advice.

We now predict y_t based on y_1^{t-1} and on the advice $f_{j,s}$, $s \leqslant t$.

Our goal is to perform almost as well as the best expert.

Note that the best expert may only be determined in hindsight, whereas we are asked to predict on-line.

To make the notion of best expert mathematically precise, we introduce a loss function $\ell : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$.

Example. Weather forecasting, $\mathcal{Y} = \{0, 1\}$ and $\mathcal{X} = [0, 1]$,

$$\ell(x, y) = |y - x|$$

Prediction with expert advice

Finitely many reference forecasters, called experts, are usually considered. They are indexed by $j = 1, \ldots, N$.

For all time rounds $t = 1, 2, \ldots$, they also form predictions $f_{j,t} = f_{j,t}(y_1^{t-1}) \in \mathcal{X}$, often referred to as advice.

We now predict y_t based on y_1^{t-1} and on the advice $f_{j,s}$, $s \leqslant t$.

Our goal is to perform almost as well as the best expert.

Note that the best expert may only be determined in hindsight, whereas we are asked to predict on-line.

To make the notion of best expert mathematically precise, we introduce a loss function $\ell : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$.

The cumulative losses of the decision-maker and of expert j are

$$\widehat{L}_n = \widehat{L}_n(y_1^n) = \sum_{t=1}^{n} \ell(\widehat{p}_t, y_t) \text{ and } L_{j,n} = L_{j,n}(y_1^n) = \sum_{t=1}^{n} \ell(f_{j,t}, y_t).$$

External regret

The regret is defined as the difference of these cumulative losses,

$$R_{j,n} = R_{j,n}(y_1^n) = \widehat{L}_n - L_{j,n} = \sum_{t=1}^{n} \ell(\widehat{p}_t, y_t) - \sum_{t=1}^{n} \ell(f_{j,t}, y_t)$$

External regret

The regret is defined as the difference of these cumulative losses,

$$R_{j,n} = R_{j,n}(y_1^n) = \widehat{L}_n - L_{j,n} = \sum_{t=1}^{n} \ell(\widehat{p}_t, y_t) - \sum_{t=1}^{n} \ell(f_{j,t}, y_t)$$

Formally, we aim at constructing prediction stategies such that for all possible outcome sequences $y_1, y_2, \ldots$,

$$\frac{1}{n} \max_{j=1,\ldots,N} R_{j,n} = o(1)$$

External regret

The regret is defined as the difference of these cumulative losses,

$$R_{j,n} = R_{j,n}(y_1^n) = \widehat{L}_n - L_{j,n} = \sum_{t=1}^{n} \ell(\widehat{p}_t, y_t) - \sum_{t=1}^{n} \ell(f_{j,t}, y_t)$$

Formally, we aim at constructing prediction stategies such that for all possible outcome sequences $y_1, y_2, \ldots$,

$$\frac{1}{n} \max_{j=1,\ldots,N} R_{j,n} = o(1)$$

Due to the worst-case assessment of the strategies, in general no deterministic strategy will work.

External regret

The regret is defined as the difference of these cumulative losses,

$$R_{j,n} = R_{j,n}(y_1^n) = \widehat{L}_n - L_{j,n} = \sum_{t=1}^{n} \ell(\widehat{p}_t, y_t) - \sum_{t=1}^{n} \ell(f_{j,t}, y_t)$$

Formally, we aim at constructing prediction stategies such that for all possible outcome sequences $y_1, y_2, \ldots$,

$$\frac{1}{n} \max_{j=1,\ldots,N} R_{j,n} = o(1)$$

Due to the worst-case assessment of the strategies, in general no deterministic strategy will work.

(Worst-case assessments amount to playing against a malicious opponent who may read our mind).

External regret

The regret is defined as the difference of these cumulative losses,

$$R_{j,n} = R_{j,n}(y_1^n) = \widehat{L}_n - L_{j,n} = \sum_{t=1}^{n} \ell(\widehat{p}_t, y_t) - \sum_{t=1}^{n} \ell(f_{j,t}, y_t)$$

Formally, we aim at constructing prediction stategies such that for all possible outcome sequences $y_1, y_2, \ldots$,

$$\frac{1}{n} \max_{j=1,\ldots,N} R_{j,n} = o(1) \quad \text{a.s.}$$

Due to the worst-case assessment of the strategies, in general no deterministic strategy will work.

Thus, we assume that the decision-maker is allowed to use randomized strategies.

Randomized prediction

A strategy is a sequence of functions

Randomized prediction

A strategy is a sequence of functions

– that, to the past outcomes y_1^{t-1} and to the past and present expert advice $f_{j,s}$, $s \leqslant t$,

Randomized prediction

A strategy is a sequence of functions

- that, to the past outcomes y_1^{t-1} and to the past and present expert advice $f_{j,s}$, $s \leqslant t$,
- assign a probability distribution $\boldsymbol{p}_t = (p_{1,t}, \ldots, p_{N,t})$ over the set of experts.

Randomized prediction

A strategy is a sequence of functions

– that, to the past outcomes y_1^{t-1} and to the past and present expert advice $f_{j,s}$, $s \leqslant t$,

– assign a probability distribution $\boldsymbol{p}_t = (p_{1,t}, \ldots, p_{N,t})$ over the set of experts.

The formed prediction is given by drawing an expert index I_t according to $\boldsymbol{p}_t$,

Randomized prediction

A strategy is a sequence of functions

- that, to the past outcomes y_1^{t-1} and to the past and present expert advice $f_{j,s}$, $s \leqslant t$,
- assign a probability distribution $\boldsymbol{p}_t = (p_{1,t}, \ldots, p_{N,t})$ over the set of experts.

The formed prediction is given by drawing an expert index I_t according to $\boldsymbol{p}_t$,

$$\widehat{p}_t = f_{I_t,t}$$

Randomized prediction

A strategy is a sequence of functions

- that, to the past outcomes y_1^{t-1} and to the past and present expert advice $f_{j,s}$, $s \leqslant t$,
- assign a probability distribution $\boldsymbol{p}_t = (p_{1,t}, \ldots, p_{N,t})$ over the set of experts.

The formed prediction is given by drawing an expert index I_t according to $\boldsymbol{p}_t$,

$$\widehat{p}_t = f_{I_t,t}$$

The outcomes may be chosen by an opponent player (who reacts to our predictions).

Randomized prediction

A strategy is a sequence of functions

- that, to the past outcomes y_1^{t-1} and to the past and present expert advice $f_{j,s}$, $s \leqslant t$,
- assign a probability distribution $\boldsymbol{p}_t = (p_{1,t}, \ldots, p_{N,t})$ over the set of experts.

The formed prediction is given by drawing an expert index I_t according to $\boldsymbol{p}_t$,

$$\widehat{p}_t = f_{I_t,t}$$

The outcomes may be chosen by an opponent player (who reacts to our predictions).

The sequential prediction problem may be cast as a (zero-sum) repeated game between a decision-maker and an opponent player.

Sequential prediction as a repeated game

Parameters: prediction space $\mathcal{X}$, outcome space $\mathcal{Y}$, N experts, number of game rounds n ($n = \infty$ is allowed).

Sequential prediction as a repeated game

Parameters: prediction space $\mathcal{X}$, outcome space $\mathcal{Y}$, N experts, number of game rounds n ($n = \infty$ is allowed).

For each round $t = 1, 2, \ldots, n$,

Sequential prediction as a repeated game

Parameters: prediction space $\mathcal{X}$, outcome space $\mathcal{Y}$, N experts, number of game rounds n ($n = \infty$ is allowed).

For each round $t = 1, 2, \ldots, n$,

– the opponent player publicly chooses the experts' predictions $f_{1,t}, \ldots, f_{N,t} \in \mathcal{X}$, and the decision-maker has an immediate access to them;

Sequential prediction as a repeated game

Parameters: prediction space $\mathcal{X}$, outcome space $\mathcal{Y}$, N experts, number of game rounds n ($n = \infty$ is allowed).

For each round $t = 1, 2, \ldots, n$,

- the opponent player publicly chooses the experts' predictions $f_{1,t}, \ldots, f_{N,t} \in \mathcal{X}$, and the decision-maker has an immediate access to them;
- the decision-maker chooses a probability distribution $\boldsymbol{p}_t = (p_{1,t}, \ldots, p_{N,t})$ over the set of experts, privately draws an expert I_t at random according to $\boldsymbol{p}_t$, and privately predicts as $\widehat{p}_t = f_{I_t,t}$;

Sequential prediction as a repeated game

Parameters: prediction space $\mathcal{X}$, outcome space $\mathcal{Y}$, N experts, number of game rounds n ($n = \infty$ is allowed).

For each round $t = 1, 2, \ldots, n$,

- the opponent player publicly chooses the experts' predictions $f_{1,t}, \ldots, f_{N,t} \in \mathcal{X}$, and the decision-maker has an immediate access to them;
- the decision-maker chooses a probability distribution $\boldsymbol{p}_t = (p_{1,t}, \ldots, p_{N,t})$ over the set of experts, privately draws an expert I_t at random according to $\boldsymbol{p}_t$, and privately predicts as $\widehat{p}_t = f_{I_t,t}$;
- simultaneously, the opponent player privately chooses the outcome $y_t \in \mathcal{Y}$;

Sequential prediction as a repeated game

Parameters: prediction space $\mathcal{X}$, outcome space $\mathcal{Y}$, N experts, number of game rounds n ($n = \infty$ is allowed).

For each round $t = 1, 2, \ldots, n$,

– the opponent player publicly chooses the experts' predictions $f_{1,t}, \ldots, f_{N,t} \in \mathcal{X}$, and the decision-maker has an immediate access to them;

– the decision-maker chooses a probability distribution $\boldsymbol{p}_t = (p_{1,t}, \ldots, p_{N,t})$ over the set of experts, privately draws an expert I_t at random according to $\boldsymbol{p}_t$, and privately predicts as $\widehat{p}_t = f_{I_t,t}$;

– simultaneously, the opponent player privately chooses the outcome $y_t \in \mathcal{Y}$;

– the outcome y_t and the prediction $\widehat{p}_t$ are made public, and the losses may be computed.

Sequential prediction as a repeated game

Parameters: prediction space $\mathcal{X}$, outcome space $\mathcal{Y}$, N experts, number of game rounds n ($n = \infty$ is allowed).

For each round $t = 1, 2, \ldots, n$,

- the opponent player publicly chooses the experts' predictions $f_{1,t}, \ldots, f_{N,t} \in \mathcal{X}$, and the decision-maker has an immediate access to them;
- the decision-maker chooses a probability distribution $\boldsymbol{p}_t = (p_{1,t}, \ldots, p_{N,t})$ over the set of experts, privately draws an expert I_t at random according to $\boldsymbol{p}_t$, and privately predicts as $\widehat{p}_t = f_{I_t,t}$;
- simultaneously, the opponent player privately chooses the outcome $y_t \in \mathcal{Y}$;
- the outcome y_t and the prediction $\widehat{p}_t$ are made public, and the losses may be computed.

Exponentially Weighted Average Predictor

Expected regret

Recall that we aim at minimizing the regret

$$\widehat{L}_n - \min_{j=1,\ldots,N} L_{j,n} = \sum_{t=1}^{n} \ell(f_{I_t,t}, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t)$$

Expected regret

Recall that we aim at minimizing the regret

$$\widehat{L}_n - \min_{j=1,\ldots,N} L_{j,n} = \sum_{t=1}^{n} \ell(f_{I_t,t}, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t)$$

Denote by $\mathbb{E}_t$ the conditional expectation with respect to time rounds 1 to $t-1$,

$$\mathbb{E}_t \left[\ell(f_{I_t,t}, y_t) \right] = \sum_{i=1}^{N} p_{i,t} \ell(f_{i,t}, y_t) = \ell(\boldsymbol{p}_t, y_t)$$

Expected regret

Recall that we aim at minimizing the regret

$$\widehat{L}_n - \min_{j=1,\dots,N} L_{j,n} = \sum_{t=1}^{n} \ell(f_{I_t,t}, y_t) - \min_{j=1,\dots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t)$$

Denote by $\mathbb{E}_t$ the conditional expectation with respect to time rounds 1 to $t-1$,

$$\mathbb{E}_t \left[\ell(f_{I_t,t}, y_t) \right] = \sum_{i=1}^{N} p_{i,t} \ell(f_{i,t}, y_t) = \ell(\boldsymbol{p}_t, y_t)$$

By martingale convergence, under a boundedness assumption,

$$\widehat{L}_n - \overline{L}_n = \sum_{t=1}^{n} \ell(f_{I_t,t}, y_t) - \sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t) = o_{\mathbb{P}}(n)$$

Expected regret

Recall that we aim at minimizing the regret

$$\widehat{L}_n - \min_{j=1,\dots,N} L_{j,n} = \sum_{t=1}^{n} \ell(f_{I_t,t}, y_t) - \min_{j=1,\dots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t)$$

Denote by $\mathbb{E}_t$ the conditional expectation with respect to time rounds 1 to $t-1$,

$$\mathbb{E}_t \left[\ell(f_{I_t,t}, y_t) \right] = \sum_{i=1}^{N} p_{i,t} \ell(f_{i,t}, y_t) = \ell(\boldsymbol{p}_t, y_t)$$

By martingale convergence, under a boundedness assumption,

$$\widehat{L}_n - \overline{L}_n = \sum_{t=1}^{n} \ell(f_{I_t,t}, y_t) - \sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t) = o_{\mathbb{P}}(n)$$

We may thus focus first on the expected regret,

$$\overline{L}_n - \min_{j=1,\dots,N} L_{j,n} = \sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t) - \min_{j=1,\dots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t)$$

Exponential reweighting

The idea is to assign a higher probability to better-performing experts.

Exponential reweighting

The idea is to assign a higher probability to better-performing experts. A popular choice is given by exponential reweightings,

$$p_{i,t} = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \ell(f_{i,s}, y_s)\right)}{\sum_{j=1}^{N} \left(-\eta \sum_{s=1}^{t-1} \ell(f_{j,s}, y_s)\right)} \qquad i = 1, \ldots, N$$

where $\eta > 0$ is a parameter to be tuned.

Exponential reweighting

The idea is to assign a higher probability to better-performing experts. A popular choice is given by exponential reweightings,

$$p_{i,t} = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \ell(f_{i,s}, y_s)\right)}{\sum_{j=1}^{N} \left(-\eta \sum_{s=1}^{t-1} \ell(f_{j,s}, y_s)\right)} \qquad i = 1, \ldots, N$$

where $\eta > 0$ is a parameter to be tuned.

This forecaster, called exponentially weighted average predictor [henceforth referred to as EWA], was introduced by

- Vovk '90,
- Littlestone and Warmuth '94.

See also

- Cesa-Bianchi, Freund, Helmbold, Haussler, Schapire, and Warmuth '97,
- Cesa-Bianchi and Lugosi '99.

Exponential reweighting

The idea is to assign a higher probability to better-performing experts. A popular choice is given by exponential reweightings,

$$p_{i,t} = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \ell(f_{i,s}, y_s)\right)}{\sum_{j=1}^{N} \left(-\eta \sum_{s=1}^{t-1} \ell(f_{j,s}, y_s)\right)} \qquad i = 1, \ldots, N$$

where $\eta > 0$ is a parameter to be tuned.

Assume that the loss function takes bounded values, say $\ell : \mathcal{X} \times \mathcal{Y} \to [0, B]$.

Exponential reweighting

The idea is to assign a higher probability to better-performing experts. A popular choice is given by exponential reweightings,

$$p_{i,t} = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \ell(f_{i,s}, y_s)\right)}{\sum_{j=1}^{N} \left(-\eta \sum_{s=1}^{t-1} \ell(f_{j,s}, y_s)\right)} \qquad i = 1, \ldots, N$$

where $\eta > 0$ is a parameter to be tuned.

Assume that the loss function takes bounded values, say $\ell : \mathcal{X} \times \mathcal{Y} \to [0, B]$. It is now a folk theorem that for all strategies of the opponent player, the expected regret is bounded as

Exponential reweighting

The idea is to assign a higher probability to better-performing experts. A popular choice is given by exponential reweightings,

$$p_{i,t} = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \ell(f_{i,s}, y_s)\right)}{\sum_{j=1}^{N} \left(-\eta \sum_{s=1}^{t-1} \ell(f_{j,s}, y_s)\right)} \qquad i = 1, \ldots, N$$

where $\eta > 0$ is a parameter to be tuned.

Assume that the loss function takes bounded values, say $\ell : \mathcal{X} \times \mathcal{Y} \to [0, B]$. It is now a folk theorem that for all strategies of the opponent player, the expected regret is bounded as

$$\sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t) \leqslant \frac{\ln N}{\eta} + \frac{\eta n}{8} B^2$$

Exponential reweighting

The idea is to assign a higher probability to better-performing experts. A popular choice is given by exponential reweightings,

$$p_{i,t} = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \ell(f_{i,s}, y_s)\right)}{\sum_{j=1}^{N} \left(-\eta \sum_{s=1}^{t-1} \ell(f_{j,s}, y_s)\right)} \qquad i = 1, \ldots, N$$

where $\eta > 0$ is a parameter to be tuned.

Assume that the loss function takes bounded values, say $\ell : \mathcal{X} \times \mathcal{Y} \to [0, B]$. It is now a folk theorem that for all strategies of the opponent player, the expected regret is bounded as

$$\sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t) \leqslant \frac{\ln N}{\eta} + \frac{\eta n}{8} B^2 = B\sqrt{\frac{n}{2} \ln N}$$

with $\eta = B^{-1}\sqrt{8 \ln N / n}$.

Exponential reweighting

The idea is to assign a higher probability to better-performing experts. A popular choice is given by exponential reweightings,

$$p_{i,t} = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \ell(f_{i,s}, y_s)\right)}{\sum_{j=1}^{N} \left(-\eta \sum_{s=1}^{t-1} \ell(f_{j,s}, y_s)\right)} \qquad i = 1, \ldots, N$$

where $\eta > 0$ is a parameter to be tuned.

Assume that the loss function takes bounded values, say $\ell : \mathcal{X} \times \mathcal{Y} \to [0, B]$. It is now a folk theorem that for all strategies of the opponent player, the expected regret is bounded as

$$\sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t) \leqslant \frac{\ln N}{\eta} + \frac{\eta n}{8} B^2 = B\sqrt{\frac{n}{2} \ln N}$$

with $\eta = B^{-1}\sqrt{8 \ln N / n}$.

A first possible goal is to improve the general (optimal) $\sqrt{n}$ bound (and to deal with the tuning of η).

General analysis of EWA (1)

Freund and Schapire '97 replaced the $B\sqrt{n}$ factor by $\sqrt{B L_n^*}$, where

$$L_n^* = \min_{j=1,\ldots,N} L_{j,n} = \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t)$$

General analysis of EWA (1)

Freund and Schapire '97 replaced the $B\sqrt{n}$ factor by $\sqrt{BL_n^*}$, where

$$L_n^* = \min_{j=1,\ldots,N} L_{j,n} = \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t)$$

Auer, Cesa-Bianchi, and Gentile '02 replaced the fixed tuning parameter η by time-varying parameters $\eta_t > 0$ to deal with the dependency on the horizon n.

General analysis of EWA (1)

We combine the underlying methods and show a general analysis of EWA, that yields a forecaster which

General analysis of EWA (1)

We combine the underlying methods and show a general analysis of EWA, that yields a forecaster which

– requires no previous knowledge of n and B,

General analysis of EWA (1)

We combine the underlying methods and show a general analysis of EWA, that yields a forecaster which

- requires no previous knowledge of n and B,
- improves both the $B\sqrt{n}$ and $\sqrt{BL_n^*}$ bounds.

General analysis of EWA (1)

We combine the underlying methods and show a general analysis of EWA, that yields a forecaster which

- requires no previous knowledge of n and B,
- improves both the $B\sqrt{n}$ and $\sqrt{BL_n^*}$ bounds.

The adaptive version of EWA uses

$$p_{i,t} = \frac{\exp\left(-\eta_t \sum_{s=1}^{t-1} \ell(f_{i,s}, y_s)\right)}{\sum_{j=1}^{N}\left(-\eta_t \sum_{s=1}^{t-1} \ell(f_{j,s}, y_s)\right)} \qquad i = 1, \ldots, N$$

General analysis of EWA (1)

We combine the underlying methods and show a general analysis of EWA, that yields a forecaster which

- requires no previous knowledge of n and B,
- improves both the $B\sqrt{n}$ and $\sqrt{BL_n^*}$ bounds.

The adaptive version of EWA uses

$$p_{i,t} = \frac{\exp\left(-\eta_t \sum_{s=1}^{t-1} \ell(f_{i,s}, y_s)\right)}{\sum_{j=1}^{N} \left(-\eta_t \sum_{s=1}^{t-1} \ell(f_{j,s}, y_s)\right)} \qquad i = 1, \ldots, N$$

where the $\eta_t > 0$ will be defined by the analysis.

General analysis of EWA (1)

We combine the underlying methods and show a general analysis of EWA, that yields a forecaster which

- requires no previous knowledge of n and B,
- improves both the $B\sqrt{n}$ and $\sqrt{BL_n^*}$ bounds.

The adaptive version of EWA uses

$$p_{i,t} = \frac{\exp\left(-\eta_t \sum_{s=1}^{t-1} \ell(f_{i,s}, y_s)\right)}{\sum_{j=1}^{N} \left(-\eta_t \sum_{s=1}^{t-1} \ell(f_{j,s}, y_s)\right)} \qquad i = 1, \ldots, N$$

where the $\eta_t > 0$ will be defined by the analysis.

For all parameters $\eta_1, \eta_2, \ldots > 0$, for all strategies of the opponent player, the expected regret is less than

General analysis of EWA (1)

We combine the underlying methods and show a general analysis of EWA, that yields a forecaster which

- requires no previous knowledge of n and B,
- improves both the $B\sqrt{n}$ and $\sqrt{BL_n^*}$ bounds.

The adaptive version of EWA uses

$$p_{i,t} = \frac{\exp\left(-\eta_t \sum_{s=1}^{t-1} \ell(f_{i,s}, y_s)\right)}{\sum_{j=1}^{N} \left(-\eta_t \sum_{s=1}^{t-1} \ell(f_{j,s}, y_s)\right)} \qquad i = 1, \ldots, N$$

where the $\eta_t > 0$ will be defined by the analysis.

For all parameters $\eta_1, \eta_2, \ldots > 0$, for all strategies of the opponent player, the expected regret is less than

$$\sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t) \leqslant \left(\frac{2}{\eta_{n+1}} - \frac{1}{\eta_1}\right) \ln N + \sum_{t=1}^{n} \Phi_t$$

General analysis of EWA (2)

The bound is $\displaystyle \left(\frac{2}{\eta_{n+1}} - \frac{1}{\eta_1} \right) \ln N + \sum_{t=1}^{n} \Phi_t$, where

$$\Phi_t = \frac{1}{\eta_t} \ln \left(\sum_{i=1}^{N} p_{i,t} e^{\eta_t (\ell(\boldsymbol{p}_t, y_t) - \ell(f_{i,t}, y_t))} \right)$$

General analysis of EWA (2)

The bound is $\displaystyle \left(\frac{2}{\eta_{n+1}} - \frac{1}{\eta_1} \right) \ln N + \sum_{t=1}^{n} \Phi_t$, where

$$\Phi_t = \frac{1}{\eta_t} \ln \left(\sum_{i=1}^{N} p_{i,t} e^{\eta_t (\ell(\boldsymbol{p}_t, y_t) - \ell(f_{i,t}, y_t))} \right)$$

Hoeffding's inequality ensures that

$$\Phi_t \leqslant B^2 \eta_t / 8$$

General analysis of EWA (2)

The bound is $\left(\dfrac{2}{\eta_{n+1}} - \dfrac{1}{\eta_1} \right) \ln N + \sum_{t=1}^{n} \Phi_t$, where

$$\Phi_t = \frac{1}{\eta_t} \ln \left(\sum_{i=1}^{N} p_{i,t} e^{\eta_t (\ell(\boldsymbol{p}_t, y_t) - \ell(f_{i,t}, y_t))} \right)$$

Hoeffding's inequality ensures that

$$\Phi_t \leqslant B^2 \eta_t / 8$$

so that the (time-adaptive) choice $\eta_t = B^{-1} \sqrt{8 \ln N / (t-1)}$ ensures that

$$\sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t) \leqslant \square \, B \sqrt{n \ln N}$$

General analysis of EWA (2)

The bound is $\left(\dfrac{2}{\eta_{n+1}} - \dfrac{1}{\eta_1}\right) \ln N + \displaystyle\sum_{t=1}^{n} \Phi_t$, where

$$\Phi_t = \frac{1}{\eta_t} \ln \left(\sum_{i=1}^{N} p_{i,t} e^{\eta_t (\ell(\boldsymbol{p}_t, y_t) - \ell(f_{i,t}, y_t))} \right)$$

The general analysis relies on $\Phi_t \leqslant \eta_t v_t$, where

$$v_t = \mathrm{Var}_t\ \ell\left(f_{I_t,t}, y_t\right) = \sum_{i \leqslant N} p_{i,t} \ell(f_{i,t}, y_t)^2 - \left(\sum_{i \leqslant N} p_{i,t} \ell(f_{i,t}, y_t)\right)^2$$

General analysis of EWA (2)

The bound is $\left(\dfrac{2}{\eta_{n+1}} - \dfrac{1}{\eta_1} \right) \ln N + \displaystyle\sum_{t=1}^{n} \Phi_t$, where

$$\Phi_t = \frac{1}{\eta_t} \ln \left(\sum_{i=1}^{N} p_{i,t} e^{\eta_t (\ell(\boldsymbol{p}_t, y_t) - \ell(f_{i,t}, y_t))} \right)$$

The general analysis relies on $\Phi_t \leqslant \eta_t v_t$, where

$$v_t = \mathrm{Var}_t \; \ell\left(f_{I_t,t}, y_t\right) = \sum_{i \leqslant N} p_{i,t} \ell(f_{i,t}, y_t)^2 - \left(\sum_{i \leqslant N} p_{i,t} \ell(f_{i,t}, y_t) \right)^2$$

Using $\eta_t = \min \left\{ \dfrac{1}{2M_{t-1}}, \square \sqrt{\dfrac{\ln N}{V_{t-1}}} \right\}$, where

General analysis of EWA (2)

The bound is $\displaystyle \left(\frac{2}{\eta_{n+1}} - \frac{1}{\eta_1} \right) \ln N + \sum_{t=1}^{n} \Phi_t$, where

$$\Phi_t = \frac{1}{\eta_t} \ln \left(\sum_{i=1}^{N} p_{i,t} e^{\eta_t (\ell(\boldsymbol{p}_t, y_t) - \ell(f_{i,t}, y_t))} \right)$$

The general analysis relies on $\Phi_t \leqslant \eta_t v_t$, where

$$v_t = \mathrm{Var}_t\ \ell\left(f_{I_t,t}, y_t\right) = \sum_{i \leqslant N} p_{i,t} \ell(f_{i,t}, y_t)^2 - \left(\sum_{i \leqslant N} p_{i,t} \ell(f_{i,t}, y_t) \right)^2$$

Using $\displaystyle \eta_t = \min \left\{ \frac{1}{2M_{t-1}}, \Box \sqrt{\frac{\ln N}{V_{t-1}}} \right\}$, where

- M_{t-1} is roughly $\displaystyle \max_{i \leqslant N,\ s \leqslant t-1} \ell(f_{i,s}, y_s)$,
- $V_{t-1} = v_1 + \ldots + v_{t-1}$, we get

General analysis of EWA (2)

The bound is $\left(\frac{2}{\eta_{n+1}} - \frac{1}{\eta_1}\right) \ln N + \sum_{t=1}^{n} \Phi_t$, where

$$\Phi_t = \frac{1}{\eta_t} \ln \left(\sum_{i=1}^{N} p_{i,t} e^{\eta_t (\ell(\boldsymbol{p}_t, y_t) - \ell(f_{i,t}, y_t))} \right)$$

The general analysis relies on $\Phi_t \leqslant \eta_t v_t$, where

$$v_t = \mathrm{Var}_t \; \ell\left(f_{I_t,t}, y_t\right) = \sum_{i \leqslant N} p_{i,t} \ell(f_{i,t}, y_t)^2 - \left(\sum_{i \leqslant N} p_{i,t} \ell(f_{i,t}, y_t) \right)^2$$

Using $\eta_t = \min \left\{ \frac{1}{2M_{t-1}}, \square \sqrt{\frac{\ln N}{V_{t-1}}} \right\}$, where

– M_{t-1} is roughly $\max_{i \leqslant N,\ s \leqslant t-1} \ell(f_{i,s}, y_s)$,

– $V_{t-1} = v_1 + \ldots + v_{t-1}$, we get

$$\overline{L}_n - L_n^* = \sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t) \leqslant \square \sqrt{V_n \ln N} + \square\, B \ln N$$

Bounds that hold with high probability

So far, we have bounded the expected regret as

$$\overline{L}_n - L_n^* = \sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t) \leqslant \Box\, B\sqrt{n \ln N}$$

or

$$\overline{L}_n - L_n^* \leqslant \Box \sqrt{V_n \ln N} + \Box\, B \ln N$$

Bounds that hold with high probability

So far, we have bounded the expected regret as

$$\overline{L}_n - L_n^* = \sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t) \leqslant \square\, B\sqrt{n \ln N}$$

or

$$\overline{L}_n - L_n^* \leqslant \square \sqrt{V_n \ln N} + \square\, B \ln N$$

(Since $V_n \leqslant B\,\overline{L}_n$, the second bound implies a $\square \sqrt{B L_n^* \ln N}$ one, as well as many other interesting ones.)

Bounds that hold with high probability

So far, we have bounded the expected regret as

$$\overline{L}_n - L_n^* = \sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t) \leqslant \Box\, B\sqrt{n \ln N}$$

or

$$\overline{L}_n - L_n^* \leqslant \Box \sqrt{V_n \ln N} + \Box\, B \ln N$$

Theses bounds also hold with overwhelming probability for the regret $\widehat{L}_n - L_n^*$, since, with probability $1-\delta$,

$$\widehat{L}_n - \overline{L}_n = \sum_{t=1}^{n} \ell(f_{I_t,t}, y_t) - \sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t)$$

Individual Sequences

The EWA strategy
- Expected regret
- Exponential reweighting
- General analysis of EWA (1)
- General analysis of EWA (2)
- Non-expected bounds

Incomplete information

Internal regret

Bounds that hold with high probability

So far, we have bounded the expected regret as

$$\overline{L}_n - L_n^* = \sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t) \leqslant \square\, B\sqrt{n \ln N}$$

or

$$\overline{L}_n - L_n^* \leqslant \square \sqrt{V_n \ln N} + \square\, B \ln N$$

Theses bounds also hold with overwhelming probability for the regret $\widehat{L}_n - L_n^*$, since, with probability $1-\delta$,

$$\widehat{L}_n - \overline{L}_n = \sum_{t=1}^{n} \ell(f_{I_t,t}, y_t) - \sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t)$$

- $\leqslant \square\, B\sqrt{n \ln(1/\delta)}$ by Hoeffding's inequality,

Bounds that hold with high probability

So far, we have bounded the expected regret as

$$\overline{L}_n - L_n^* = \sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t) \leqslant \square\, B\sqrt{n \ln N}$$

or

$$\overline{L}_n - L_n^* \leqslant \square\, \sqrt{V_n \ln N} + \square\, B \ln N$$

Theses bounds also hold with overwhelming probability for the regret $\widehat{L}_n - L_n^*$, since, with probability $1-\delta$,

$$\widehat{L}_n - \overline{L}_n = \sum_{t=1}^{n} \ell(f_{I_t,t}, y_t) - \sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t)$$

- $\leqslant \square\, B\sqrt{n \ln(1/\delta)}$ by Hoeffding's inequality,
- $\leqslant \square\, \sqrt{V_n \ln(1/\delta)} + B \ln(1/\delta)$ by Bernstein's inequality.

Bounds that hold with high probability

So far, we have bounded the expected regret as

$$\overline{L}_n - L_n^* = \sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t) \leqslant \Box\, B\sqrt{n \ln N}$$

or

$$\overline{L}_n - L_n^* \leqslant \Box\, \sqrt{V_n \ln N} + \Box\, B \ln N$$

So, we have bounded the non-expected regret, with probability $1-\delta$, as

$$\widehat{L}_n - L_n^* = \sum_{t=1}^{n} \ell(f_{I_t,t}, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t) \leqslant \Box\, B\sqrt{n \ln(N/\delta)}$$

or

$$\widehat{L}_n - L_n^* \leqslant \Box\, \sqrt{V_n \ln(N/\delta)} + \Box\, B \ln(N/\delta)$$

Bounds that hold with high probability

So far, we have bounded the expected regret as

$$\overline{L}_n - L_n^* = \sum_{t=1}^{n} \ell(\boldsymbol{p}_t, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t) \leqslant \Box\, B\sqrt{n \ln N}$$

or

$$\overline{L}_n - L_n^* \leqslant \Box \sqrt{V_n \ln N} + \Box\, B \ln N$$

So, we have bounded the non-expected regret, with probability $1-\delta$, as

$$\widehat{L}_n - L_n^* = \sum_{t=1}^{n} \ell(f_{I_t,t}, y_t) - \min_{j=1,\ldots,N} \sum_{t=1}^{n} \ell(f_{j,t}, y_t) \leqslant \Box\, B\sqrt{n \ln(N/\delta)}$$

or

$$\widehat{L}_n - L_n^* \leqslant \Box \sqrt{V_n \ln(N/\delta)} + \Box\, B \ln(N/\delta)$$

To get refined bounds on the regret, we would first need martingale inequalities sharper than Bernstein's inequality.

Prediction with partial monitoring

On-line pricing (1)

Gábor's celebrated T-shirt selling example!

On-line pricing (1)

A vendor sells T-shirts on the Internet. Customers connect one by one to his Web site.

On-line pricing (1)

A vendor sells T-shirts on the Internet. Customers connect one by one to his Web site.

To customer number t, the vendor offers the T-shirt at a price $I_t \in \mathcal{X} = [0, 1]$.

On-line pricing (1)

A vendor sells T-shirts on the Internet. Customers connect one by one to his Web site.

To customer number t, the vendor offers the T-shirt at a price $I_t \in \mathcal{X} = [0, 1]$.

Each customer has in mind a maximum price $y_t \in \mathcal{Y} = [0, 1]$ he is willing to pay – but does not tell it to the vendor.

On-line pricing (1)

A vendor sells T-shirts on the Internet. Customers connect one by one to his Web site.

To customer number t, the vendor offers the T-shirt at a price $I_t \in \mathcal{X} = [0, 1]$.

Each customer has in mind a maximum price $y_t \in \mathcal{Y} = [0, 1]$ he is willing to pay – but does not tell it to the vendor.

When $y_t \geqslant I_t$, the product is bought, and the vendor suffers a loss (of earnings) $y_t - I_t$.

On-line pricing (1)

A vendor sells T-shirts on the Internet. Customers connect one by one to his Web site.

To customer number t, the vendor offers the T-shirt at a price $I_t \in \mathcal{X} = [0, 1]$.

Each customer has in mind a maximum price $y_t \in \mathcal{Y} = [0, 1]$ he is willing to pay – but does not tell it to the vendor.

When $y_t \geqslant I_t$, the product is bought, and the vendor suffers a loss (of earnings) $y_t - I_t$.

When $y_t < I_t$, the product is not bought, and the vendor's loss is a fixed $c \in [0, 1]$ (representing all his charges).

On-line pricing (1)

A vendor sells T-shirts on the Internet. Customers connect one by one to his Web site.

To customer number t, the vendor offers the T-shirt at a price $I_t \in \mathcal{X} = [0,1]$.

Each customer has in mind a maximum price $y_t \in \mathcal{Y} = [0,1]$ he is willing to pay – but does not tell it to the vendor.

When $y_t \geqslant I_t$, the product is bought, and the vendor suffers a loss (of earnings) $y_t - I_t$.

When $y_t < I_t$, the product is not bought, and the vendor's loss is a fixed $c \in [0,1]$ (representing all his charges).

The vendor's loss function is

$$\ell(I_t, y_t) = (y_t - I_t)\,\mathbb{I}_{y_t \geqslant I_t} + c\,\mathbb{I}_{y_t < I_t}$$

and he only gets the feedback

$$h(I_t, y_t) = \mathbb{I}_{y_t \geqslant I_t}$$

On-line pricing (2)

The values y_t are chosen by an arbitrary opponent.

On-line pricing (2)

The values y_t are chosen by an arbitrary opponent.

We compare to rivals using strategies defined by prior market research.

On-line pricing (2)

The values y_t are chosen by an arbitrary opponent.

We compare to rivals using strategies defined by prior market research.

Each market research indicates a constant price $q \in [0, 1]$, and we thus compare our cumulative loss to the one of the best constant prices, that is, the regret equals

On-line pricing (2)

The values y_t are chosen by an arbitrary opponent.

We compare to rivals using strategies defined by prior market research.

Each market research indicates a constant price $q \in [0,1]$, and we thus compare our cumulative loss to the one of the best constant prices, that is, the regret equals

$$\sum_{t=1}^{n} \ell(I_t, y_t) - \min_{q \in [0,1]} \sum_{t=1}^{n} \ell(q, y_t)$$

General model

Denote by $\mathcal{X}$ a finite set of allowed actions, and by $\mathcal{Y}$ a finite set of outcomes.

General model

Denote by $\mathcal{X}$ a finite set of allowed actions, and by $\mathcal{Y}$ a finite set of outcomes.

We introduce two functions,

General model

Denote by $\mathcal{X}$ a finite set of allowed actions, and by $\mathcal{Y}$ a finite set of outcomes.

We introduce two functions,

– a loss function $\ell : \mathcal{X} \times \mathcal{Y} \to [0, 1]$,

General model

Denote by $\mathcal{X}$ a finite set of allowed actions, and by $\mathcal{Y}$ a finite set of outcomes.

We introduce two functions,

- a loss function $\ell : \mathcal{X} \times \mathcal{Y} \to [0, 1]$,
- a feedback function $h : \mathcal{X} \times \mathcal{Y} \to \mathcal{S}$, where $\mathcal{S} \subset [-1, 1]$ is a finite set of signals.

General model

Denote by $\mathcal{X}$ a finite set of allowed actions, and by $\mathcal{Y}$ a finite set of outcomes.

We introduce two functions,

- a loss function $\ell : \mathcal{X} \times \mathcal{Y} \to [0, 1]$,
- a feedback function $h : \mathcal{X} \times \mathcal{Y} \to \mathcal{S}$, where $\mathcal{S} \subset [-1, 1]$ is a finite set of signals.

After choosing $I_t \in \mathcal{X}$, the decision-maker only observes the feedback $h(I_t, y_t)$, and suffers a loss $\ell(I_t, y_t)$.

General model

Denote by $\mathcal{X}$ a finite set of allowed actions, and by $\mathcal{Y}$ a finite set of outcomes.

We introduce two functions,

- a loss function $\ell : \mathcal{X} \times \mathcal{Y} \to [0,1]$,
- a feedback function $h : \mathcal{X} \times \mathcal{Y} \to \mathcal{S}$, where $\mathcal{S} \subset [-1,1]$ is a finite set of signals.

After choosing $I_t \in \mathcal{X}$, the decision-maker only observes the feedback $h(I_t, y_t)$, and suffers a loss $\ell(I_t, y_t)$.

The forecaster's goal is to predict in a way such that his regret

$$\sum_{t=1}^{n} \ell(I_t, y_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^{n} \ell(x, y_t) = o(n) \quad \text{a.s.}$$

independently of the opponent player's strategy.

Game of prediction with partial monitoring

The problem of sequential prediction under partial monitoring may be cast as a (zero-sum) repeated game between the decision-maker and an opponent player.

Game of prediction with partial monitoring

The problem of sequential prediction under partial monitoring may be cast as a (zero-sum) repeated game between the decision-maker and an opponent player.

For each round $t = 1, 2, \ldots,$

Game of prediction with partial monitoring

The problem of sequential prediction under partial monitoring may be cast as a (zero-sum) repeated game between the decision-maker and an opponent player.

For each round $t = 1, 2, \ldots,$

- the opponent player chooses the next outcome $y_t \in \mathcal{Y}$ without revealing it;

Game of prediction with partial monitoring

The problem of sequential prediction under partial monitoring may be cast as a (zero-sum) repeated game between the decision-maker and an opponent player.

For each round $t = 1, 2, \ldots,$

- the opponent player chooses the next outcome $y_t \in \mathcal{Y}$ without revealing it;
- the decision-maker chooses a probability distribution $\boldsymbol{p}_t$ over $\mathcal{X}$, and (privately) draws an action I_t according to $\boldsymbol{p}_t$;

Game of prediction with partial monitoring

The problem of sequential prediction under partial monitoring may be cast as a (zero-sum) repeated game between the decision-maker and an opponent player.

For each round $t = 1, 2, \ldots$,

- the opponent player chooses the next outcome $y_t \in \mathcal{Y}$ without revealing it;
- the decision-maker chooses a probability distribution $\boldsymbol{p}_t$ over $\mathcal{X}$, and (privately) draws an action I_t according to $\boldsymbol{p}_t$;
- the decision-maker incurs loss $\ell(I_t, y_t)$ and each $x \in \mathcal{X}$ incurs loss $\ell(x, y_t)$, but none of these values is revealed;

Game of prediction with partial monitoring

The problem of sequential prediction under partial monitoring may be cast as a (zero-sum) repeated game between the decision-maker and an opponent player.

For each round $t = 1, 2, \ldots$,

- the opponent player chooses the next outcome $y_t \in \mathcal{Y}$ without revealing it;
- the decision-maker chooses a probability distribution $\boldsymbol{p}_t$ over $\mathcal{X}$, and (privately) draws an action I_t according to $\boldsymbol{p}_t$;
- the decision-maker incurs loss $\ell(I_t, y_t)$ and each $x \in \mathcal{X}$ incurs loss $\ell(x, y_t)$, but none of these values is revealed;
- only the feedback $h(I_t, y_t)$ is revealed to the decision-maker.

Game of prediction with partial monitoring

The problem of sequential prediction under partial monitoring may be cast as a (zero-sum) repeated game between the decision-maker and an opponent player.

For each round $t = 1, 2, \ldots$,

- the opponent player chooses the next outcome $y_t \in \mathcal{Y}$ without revealing it;
- the decision-maker chooses a probability distribution $\boldsymbol{p}_t$ over $\mathcal{X}$, and (privately) draws an action I_t according to $\boldsymbol{p}_t$;
- the decision-maker incurs loss $\ell(I_t, y_t)$ and each $x \in \mathcal{X}$ incurs loss $\ell(x, y_t)$, but none of these values is revealed;
- only the feedback $h(I_t, y_t)$ is revealed to the decision-maker.

The decision-maker's goal is to predict in a way such that his regret

$$\sum_{t=1}^{n} \ell(I_t, y_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^{n} \ell(x, y_t) = o(n) \quad \text{a.s.}$$

Examples

Multi-armed bandit problem.
The only information the forecaster receives is his own loss,

$$h = \ell$$

(See Auer, Cesa-Bianchi, Freund, and Schapire '02.)

Examples

Multi-armed bandit problem.
The only information the forecaster receives is his own loss,

$$h = \ell$$

(See Auer, Cesa-Bianchi, Freund, and Schapire '02.)

Apple tasting. $\mathcal{X} = \mathcal{Y} = \{0, 1\}$,

$$\boldsymbol{L} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \qquad \boldsymbol{H} = \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix}$$

Examples

Multi-armed bandit problem.
The only information the forecaster receives is his own loss,

$$h = \ell$$

(See Auer, Cesa-Bianchi, Freund, and Schapire '02.)

Apple tasting. $\mathcal{X} = \mathcal{Y} = \{0, 1\}$,

$$\boldsymbol{L} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \qquad \boldsymbol{H} = \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix}$$

The forecaster only receives feedback when he chooses the first action. (See Helmbold and Panizza '97.)

Examples

Multi-armed bandit problem.
The only information the forecaster receives is his own loss,

$$h = \ell$$

(See Auer, Cesa-Bianchi, Freund, and Schapire ’02.)

Apple tasting. $\mathcal{X} = \mathcal{Y} = \{0, 1\}$,

$$\boldsymbol{L} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \qquad \boldsymbol{H} = \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix}$$

The forecaster only receives feedback when he chooses the first action. (See Helmbold and Panizza ’97.)

We now consider the matrices

$$\boldsymbol{L} = [\ell(x, y)]_{(x,y)} \quad \text{and} \quad \boldsymbol{H} = [h(x, y)]_{(x,y)}$$

instead of the functions ℓ and h.

A general predictor (1)

The strategy is inspired by Piccolboni and Schindelhauer '01.

A general predictor (1)

The strategy is inspired by Piccolboni and Schindelhauer '01.

It relies on the crucial structural assumption that the losses may be reconstructed from the feedback, that is,

A general predictor (1)

The strategy is inspired by Piccolboni and Schindelhauer '01.

It relies on the crucial structural assumption that the losses may be reconstructed from the feedback, that is,

there exists a matrix $\boldsymbol{K} = [k(x,z)]$ such that $\boldsymbol{L} = \boldsymbol{K}\,\boldsymbol{H}$

A general predictor (1)

The strategy is inspired by Piccolboni and Schindelhauer '01.

It relies on the crucial structural assumption that the losses may be reconstructed from the feedback, that is,

there exists a matrix $\boldsymbol{K} = [k(x,z)]$ such that $\boldsymbol{L} = \boldsymbol{K}\,\boldsymbol{H}$

Thus,

$$\ell(x,y) = \sum_{z \in \mathcal{X}} k(x,z) h(z,y)$$

A general predictor (1)

The strategy is inspired by Piccolboni and Schindelhauer '01.

It relies on the crucial structural assumption that the losses may be reconstructed from the feedback, that is,

there exists a matrix $\boldsymbol{K} = [k(x,z)]$ such that $\quad \boldsymbol{L} = \boldsymbol{K}\,\boldsymbol{H}$

Thus, $$\ell(x,y) = \sum_{z \in \mathcal{X}} k(x,z) h(z,y)$$

Then we may estimate the losses, for all $x \in \mathcal{X}$, by

$$\widetilde{\ell}(x,y_t) = \frac{k(x,I_t)h(I_t,y_t)}{p_{I_t,t}}$$

A general predictor (1)

The strategy is inspired by Piccolboni and Schindelhauer '01.

It relies on the crucial structural assumption that the losses may be reconstructed from the feedback, that is,

there exists a matrix $\boldsymbol{K} = [k(x,z)]$ such that $\boldsymbol{L} = \boldsymbol{K}\,\boldsymbol{H}$

Thus,
$$\ell(x,y) = \sum_{z \in \mathcal{X}} k(x,z)h(z,y)$$

Then we may estimate the losses, for all $x \in \mathcal{X}$, by

$$\widetilde{\ell}(x,y_t) = \frac{k(x,I_t)h(I_t,y_t)}{p_{I_t,t}}$$

These estimates are unbiased.

A general predictor (1)

The strategy is inspired by Piccolboni and Schindelhauer '01.

It relies on the crucial structural assumption that the losses may be reconstructed from the feedback, that is,

there exists a matrix $\boldsymbol{K} = [k(x,z)]$ such that $\boldsymbol{L} = \boldsymbol{K}\,\boldsymbol{H}$

Thus,
$$\ell(x,y) = \sum_{z \in \mathcal{X}} k(x,z) h(z,y)$$

Then we may estimate the losses, for all $x \in \mathcal{X}$, by

$$\widetilde{\ell}(x,y_t) = \frac{k(x,I_t)h(I_t,y_t)}{p_{I_t,t}}$$

These estimates are unbiased. Since I_t has law $\boldsymbol{p}_t$,

A general predictor (1)

The strategy is inspired by Piccolboni and Schindelhauer ’01.

It relies on the crucial structural assumption that the losses may be reconstructed from the feedback, that is,

there exists a matrix $\boldsymbol{K} = [k(x,z)]$ such that $\boldsymbol{L} = \boldsymbol{K}\,\boldsymbol{H}$

Thus, $$\ell(x,y) = \sum_{z \in \mathcal{X}} k(x,z) h(z,y)$$

Then we may estimate the losses, for all $x \in \mathcal{X}$, by

$$\widetilde{\ell}(x,y_t) = \frac{k(x,I_t)h(I_t,y_t)}{p_{I_t,t}}$$

These estimates are unbiased. Since I_t has law $\boldsymbol{p}_t$,

$$\mathbb{E}_t\left[\widetilde{\ell}(x,y_t)\right] = \sum_{z \in \mathcal{X}} p_{z,t} \frac{k(x,z)h(z,y_t)}{p_{z,t}} = \ell(x,y_t)$$

A general predictor (2)

Parameters. L, H (and K) are known.

A general predictor (2)

Parameters. L, H (and K) are known.

Initialization. $\widetilde{L}_{x,0} = 0$ for all $x \in \mathcal{X}$.

A general predictor (2)

Parameters. $\boldsymbol{L}$, $\boldsymbol{H}$ (and $\boldsymbol{K}$) are known.

Initialization. $\widetilde{L}_{x,0} = 0$ for all $x \in \mathcal{X}$.

For each round $t = 1, 2, \ldots$,

A general predictor (2)

Parameters. $\boldsymbol{L}$, $\boldsymbol{H}$ (and $\boldsymbol{K}$) are known.

Initialization. $\widetilde{L}_{x,0} = 0$ for all $x \in \mathcal{X}$.

For each round $t = 1, 2, \ldots$,

- let $\eta_t \sim t^{-2/3}$ and $\gamma_t \sim t^{-1/3}$;

A general predictor (2)

Parameters. $\boldsymbol{L}$, $\boldsymbol{H}$ (and $\boldsymbol{K}$) are known.

Initialization. $\widetilde{L}_{x,0} = 0$ for all $x \in \mathcal{X}$.

For each round $t = 1, 2, \ldots$,

- let $\eta_t \sim t^{-2/3}$ and $\gamma_t \sim t^{-1/3}$;
- draw an action I_t from $\mathcal{X}$ at random, according to the distribution $\boldsymbol{p}_t$ defined by

$$p_{x,t} = (1 - \gamma_t) \frac{e^{-\eta_t \widetilde{L}_{x,t-1}}}{\sum_{z \in \mathcal{X}} e^{-\eta_t \widetilde{L}_{z,t-1}}} + \frac{\gamma_t}{|\mathcal{X}|} \; ;$$

A general predictor (2)

Parameters. $\boldsymbol{L}$, $\boldsymbol{H}$ (and $\boldsymbol{K}$) are known.

Initialization. $\widetilde{L}_{x,0} = 0$ for all $x \in \mathcal{X}$.

For each round $t = 1, 2, \ldots$,

- let $\eta_t \sim t^{-2/3}$ and $\gamma_t \sim t^{-1/3}$;
- draw an action I_t from $\mathcal{X}$ at random, according to the distribution $\boldsymbol{p}_t$ defined by

$$p_{x,t} = (1-\gamma_t) \frac{e^{-\eta_t \widetilde{L}_{x,t-1}}}{\sum_{z \in \mathcal{X}} e^{-\eta_t \widetilde{L}_{z,t-1}}} + \frac{\gamma_t}{|\mathcal{X}|} ;$$

- let $\widetilde{L}_{x,t} = \widetilde{L}_{x,t-1} + \widetilde{\ell}(x, y_t)$ for all $x \in \mathcal{X}$,

A general predictor (2)

Parameters. $\boldsymbol{L}$, $\boldsymbol{H}$ (and $\boldsymbol{K}$) are known.

Initialization. $\widetilde{L}_{x,0} = 0$ for all $x \in \mathcal{X}$.

For each round $t = 1, 2, \ldots$,

- let $\eta_t \sim t^{-2/3}$ and $\gamma_t \sim t^{-1/3}$;
- draw an action I_t from $\mathcal{X}$ at random, according to the distribution $\boldsymbol{p}_t$ defined by

$$p_{x,t} = (1 - \gamma_t) \frac{e^{-\eta_t \widetilde{L}_{x,t-1}}}{\sum_{z \in \mathcal{X}} e^{-\eta_t \widetilde{L}_{z,t-1}}} + \frac{\gamma_t}{|\mathcal{X}|} \; ;$$

- let $\widetilde{L}_{x,t} = \widetilde{L}_{x,t-1} + \widetilde{\ell}(x, y_t)$ for all $x \in \mathcal{X}$, where

$$\widetilde{\ell}(x, y_t) = \frac{k(x, I_t) h(I_t, y_t)}{p_{I_t,t}}$$

Performance bound

Theorem. For all strategies of the opponent,

Performance bound

Theorem. For all strategies of the opponent, for all $n \geqslant 1$ and $\delta \in [0,1]$, and with probability at least $1-\delta$,

$$\sum_{t=1}^{n} \ell(I_t, y_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^{n} \ell(x, y_t) \leqslant \Box\, n^{2/3} \sqrt{\ln \frac{1}{\delta}}$$

Performance bound

Theorem. For all strategies of the opponent, for all $n \geqslant 1$ and $\delta \in [0,1]$, and with probability at least $1-\delta$,

$$\sum_{t=1}^{n} \ell(I_t, y_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^{n} \ell(x, y_t) \leqslant \square\, n^{2/3} \sqrt{\ln \frac{1}{\delta}}$$

Application. For the on-line pricing problem, we partition $[0,1]$ and use a doubling trick to bound the regret by $\square\, n^{4/5}$.

Performance bound

Theorem. For all strategies of the opponent, for all $n \geqslant 1$ and $\delta \in [0, 1]$, and with probability at least $1 - \delta$,

$$\sum_{t=1}^{n} \ell(I_t, y_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^{n} \ell(x, y_t) \leqslant \Box\, n^{2/3} \sqrt{\ln \frac{1}{\delta}}$$

Application. For the on-line pricing problem, we partition $[0, 1]$ and use a doubling trick to bound the regret by $\Box\, n^{4/5}$.

Proof.

Performance bound

Theorem. For all strategies of the opponent, for all $n \geqslant 1$ and $\delta \in [0,1]$, and with probability at least $1-\delta$,

$$\sum_{t=1}^{n} \ell(I_t, y_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^{n} \ell(x, y_t) \leqslant \Box\, n^{2/3} \sqrt{\ln \frac{1}{\delta}}$$

Application. For the on-line pricing problem, we partition $[0,1]$ and use a doubling trick to bound the regret by $\Box\, n^{4/5}$.

Proof. Recall that the variance terms are crucial. Now,

Performance bound

Theorem. For all strategies of the opponent, for all $n \geqslant 1$ and $\delta \in [0,1]$, and with probability at least $1-\delta$,

$$\sum_{t=1}^{n} \ell(I_t, y_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^{n} \ell(x, y_t) \leqslant \square\, n^{2/3} \sqrt{\ln \frac{1}{\delta}}$$

Application. For the on-line pricing problem, we partition $[0,1]$ and use a doubling trick to bound the regret by $\square\, n^{4/5}$.

Proof. Recall that the variance terms are crucial. Now,

$$\mathbb{E}_t \left[\widetilde{\ell}(x, y_t)^2 \right] = \sum_{z \in \mathcal{X}} p_{z,t} \left(\frac{k(x,z) h(z, y_t)}{p_{z,t}} \right)^2 \leqslant \frac{\left(|\mathcal{X}| \, \|K\|_\infty \right)^2}{\gamma_t} \sim t^{1/3}$$

Performance bound

Theorem. For all strategies of the opponent, for all $n \geqslant 1$ and $\delta \in [0,1]$, and with probability at least $1-\delta$,

$$\sum_{t=1}^{n} \ell(I_t, y_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^{n} \ell(x, y_t) \leqslant \square\, n^{2/3} \sqrt{\ln \frac{1}{\delta}}$$

Application. For the on-line pricing problem, we partition $[0,1]$ and use a doubling trick to bound the regret by $\square\, n^{4/5}$.

Proof. Recall that the variance terms are crucial. Now,

$$\mathbb{E}_t \left[\widetilde{\ell}(x, y_t)^2 \right] = \sum_{z \in \mathcal{X}} p_{z,t} \left(\frac{k(x,z)h(z,y_t)}{p_{z,t}} \right)^2 \leqslant \frac{\left(|\mathcal{X}|\, \|K\|_\infty\right)^2}{\gamma_t} \sim t^{1/3}$$

so that $\sqrt{V_n} \leqslant \square\, n^{2/3}$

Performance bound

Theorem. For all strategies of the opponent, for all $n \geqslant 1$ and $\delta \in [0,1]$, and with probability at least $1-\delta$,

$$\sum_{t=1}^{n} \ell(I_t, y_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^{n} \ell(x, y_t) \leqslant \Box\, n^{2/3} \sqrt{\ln \frac{1}{\delta}}$$

Application. For the on-line pricing problem, we partition $[0,1]$ and use a doubling trick to bound the regret by $\Box\, n^{4/5}$.

Proof. In particular, Bernstein's inequality for martingale difference sequences show that

$$\left(\sum_{t=1}^{n} \ell(I_t, y_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^{n} \ell(x, y_t) \right) = \left(\sum_{t=1}^{n} \widetilde{\ell}(\boldsymbol{p}_t, y_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^{n} \widetilde{\ell}(x, y_t) \right) + O_{\mathbb{P}}(n^{2/3})$$

Performance bound

Theorem. For all strategies of the opponent, for all $n \geqslant 1$ and $\delta \in [0, 1]$, and with probability at least $1 - \delta$,

$$\sum_{t=1}^{n} \ell(I_t, y_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^{n} \ell(x, y_t) \leqslant \square\, n^{2/3} \sqrt{\ln \frac{1}{\delta}}$$

Application. For the on-line pricing problem, we partition $[0, 1]$ and use a doubling trick to bound the regret by $\square\, n^{4/5}$.

Proof. Also, by the second-order prediction techniques,

$$\sum_{t=1}^{n} \widetilde{\ell}(\boldsymbol{p}_t, y_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^{n} \widetilde{\ell}(x, y_t) = O(n^{2/3})$$

thus concluding the proof.

Optimality of the general strategy

Consider the problem of label efficient prediction.

Optimality of the general strategy

Consider the problem of label efficient prediction. That is,

$$\boldsymbol{L} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \\ 1 & 0 \end{bmatrix}, \qquad \boldsymbol{H} = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 1 \end{bmatrix}$$

($\boldsymbol{K}$ such that $\boldsymbol{L} = \boldsymbol{K}\,\boldsymbol{H}$ indeed exists).

Optimality of the general strategy

Consider the problem of label efficient prediction. That is,

$$\boldsymbol{L} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \\ 1 & 0 \end{bmatrix}, \qquad \boldsymbol{H} = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 1 \end{bmatrix}$$

($\boldsymbol{K}$ such that $\boldsymbol{L} = \boldsymbol{K}\,\boldsymbol{H}$ indeed exists).

Only the first action is informative, but it incurs a maximal loss.

Optimality of the general strategy

Consider the problem of label efficient prediction.

Theorem. For any $n \geqslant 8$ and any (randomized) strategy, there exists a sequence $y_1, \ldots, y_n$ of outcomes such that

Optimality of the general strategy

Consider the problem of label efficient prediction.

Theorem. For any $n \geqslant 8$ and any (randomized) strategy, there exists a sequence $y_1, \ldots, y_n$ of outcomes such that

$$\mathbb{E}_{\mathbb{A}} \left[\sum_{t=1}^{n} \ell(I_t, y_t) \right] - \min_{i=1,2,3} \sum_{t=1}^{n} \ell(i, y_t) \geqslant \frac{n^{2/3}}{5} ,$$

where $\mathbb{E}_{\mathbb{A}}$ denotes expectation with respect to the auxiliary randomization used by the forecaster.

Optimality of the general strategy

Consider the problem of label efficient prediction.

Theorem. For any $n \geqslant 8$ and any (randomized) strategy, there exists a sequence $y_1, \ldots, y_n$ of outcomes such that

$$\mathbb{E}_{\mathbb{A}} \left[\sum_{t=1}^{n} \ell(I_t, y_t) \right] - \min_{i=1,2,3} \sum_{t=1}^{n} \ell(i, y_t) \geqslant \frac{n^{2/3}}{5} ,$$

where $\mathbb{E}_{\mathbb{A}}$ denotes expectation with respect to the auxiliary randomization used by the forecaster.

Proof. Use a randomization over the outcomes, as well as some information-theoretic techniques, namely, Pinsker's inequality, the chain rule for relative entropies.

Optimality of the general strategy

Consider the problem of label efficient prediction.

Theorem. For any $n \geqslant 8$ and any (randomized) strategy, there exists a sequence $y_1, \ldots, y_n$ of outcomes such that

$$\mathbb{E}_{\mathbb{A}}\left[\sum_{t=1}^{n} \ell(I_t, y_t)\right] - \min_{i=1,2,3} \sum_{t=1}^{n} \ell(i, y_t) \geqslant \frac{n^{2/3}}{5} ,$$

where $\mathbb{E}_{\mathbb{A}}$ denotes expectation with respect to the auxiliary randomization used by the forecaster.

A simple modification of an argument introduced by Piccolboni and Schindelhauer ’01 leads to the following alternative.

Optimality of the general strategy

Consider the problem of label efficient prediction.

Theorem. For any $n \geqslant 8$ and any (randomized) strategy, there exists a sequence $y_1, \ldots, y_n$ of outcomes such that

$$\mathbb{E}_{\mathbb{A}} \left[\sum_{t=1}^{n} \ell(I_t, y_t) \right] - \min_{i=1,2,3} \sum_{t=1}^{n} \ell(i, y_t) \geqslant \frac{n^{2/3}}{5} ,$$

where $\mathbb{E}_{\mathbb{A}}$ denotes expectation with respect to the auxiliary randomization used by the forecaster.

Theorem. Every prediction problem $(\boldsymbol{H}, \boldsymbol{L})$

- either reduces to a not more difficult problem $(\boldsymbol{H}', \boldsymbol{L}')$, with $\boldsymbol{L}' = \boldsymbol{K}'\boldsymbol{H}'$, so that the regret is bounded by $O_{\mathbb{P}}(n^{2/3})$ by the general forecaster,

Optimality of the general strategy

Consider the problem of label efficient prediction.

Theorem. For any $n \geqslant 8$ and any (randomized) strategy, there exists a sequence $y_1, \ldots, y_n$ of outcomes such that

$$\mathbb{E}_{\mathbb{A}}\left[\sum_{t=1}^{n} \ell(I_t, y_t)\right] - \min_{i=1,2,3} \sum_{t=1}^{n} \ell(i, y_t) \geqslant \frac{n^{2/3}}{5} ,$$

where $\mathbb{E}_{\mathbb{A}}$ denotes expectation with respect to the auxiliary randomization used by the forecaster.

Theorem. Every prediction problem $(\boldsymbol{H}, \boldsymbol{L})$

– either reduces to a not more difficult problem $(\boldsymbol{H}', \boldsymbol{L}')$, with $\boldsymbol{L}' = \boldsymbol{K}'\boldsymbol{H}'$, so that the regret is bounded by $O_{\mathbb{P}}(n^{2/3})$ by the general forecaster,

– or is such that no strategy is able to ensure a regret uniformly $o(n)$.

Optimality of the general strategy

Consider the problem of label efficient prediction.

Theorem. For any $n \geqslant 8$ and any (randomized) strategy, there exists a sequence $y_1, \ldots, y_n$ of outcomes such that

$$\mathbb{E}_{\mathbb{A}} \left[\sum_{t=1}^{n} \ell(I_t, y_t) \right] - \min_{i=1,2,3} \sum_{t=1}^{n} \ell(i, y_t) \geqslant \frac{n^{2/3}}{5} ,$$

where $\mathbb{E}_{\mathbb{A}}$ denotes expectation with respect to the auxiliary randomization used by the forecaster.

So, in general, the proposed forecaster and the $n^{2/3}$ rate for the regrets are optimal.

Optimality of the general strategy

Consider the problem of label efficient prediction.

Theorem. For any $n \geqslant 8$ and any (randomized) strategy, there exists a sequence $y_1, \ldots, y_n$ of outcomes such that

$$\mathbb{E}_{\mathbb{A}}\left[\sum_{t=1}^{n} \ell(I_t, y_t)\right] - \min_{i=1,2,3} \sum_{t=1}^{n} \ell(i, y_t) \geqslant \frac{n^{2/3}}{5} \,,$$

where $\mathbb{E}_{\mathbb{A}}$ denotes expectation with respect to the auxiliary randomization used by the forecaster.

So, in general, the proposed forecaster and the $n^{2/3}$ rate for the regrets are optimal.

In some special cases (full information, bandit prediction), the convergence rates may however be of the faster order of $\sqrt{n}$.

Internal regret

A refined notion of regret

Internal regret, a notion introduced by Foster and Vohra ’98, is concerned with consistent modifications of our prediction strategy.

A refined notion of regret

Internal regret, a notion introduced by Foster and Vohra ’98, is concerned with consistent modifications of our prediction strategy.

Each of these modifications is parametrized by a departure function $\Phi : \{1, \ldots, N\} \to \{1, \ldots, N\}$,

A refined notion of regret

Internal regret, a notion introduced by Foster and Vohra '98, is concerned with consistent modifications of our prediction strategy.

Each of these modifications is parametrized by a departure function $\Phi : \{1, \ldots, N\} \to \{1, \ldots, N\}$, and we ask that none of them is significantly better than the original strategy.

A refined notion of regret

Internal regret, a notion introduced by Foster and Vohra '98, is concerned with consistent modifications of our prediction strategy.

Each of these modifications is parametrized by a departure function $\Phi : \{1, \ldots, N\} \to \{1, \ldots, N\}$, and we ask that none of them is significantly better than the original strategy.

Formally, a strategy is said to suffer no internal regret whenever

$$\sum_{t=1}^{n} \ell\left(f_{I_t,t}, y_t\right) - \min_{\Phi} \sum_{t=1}^{n} \ell\left(f_{\Phi(I_t),t}, y_t\right) = o(n) \quad \text{a.s.}$$

A refined notion of regret

Internal regret, a notion introduced by Foster and Vohra '98, is concerned with consistent modifications of our prediction strategy.

Each of these modifications is parametrized by a departure function $\Phi : \{1, \ldots, N\} \to \{1, \ldots, N\}$, and we ask that none of them is significantly better than the original strategy.

Formally, a strategy is said to suffer no internal regret whenever

$$\sum_{t=1}^{n} \ell\left(f_{I_t,t}, y_t\right) - \min_{\Phi} \sum_{t=1}^{n} \ell\left(f_{\Phi(I_t),t}, y_t\right) = o(n) \quad \text{a.s.}$$

Note that (external) regret corresponds to the departures $\Phi \equiv j$ for some j.

A refined notion of regret

Internal regret, a notion introduced by Foster and Vohra '98, is concerned with consistent modifications of our prediction strategy.

Each of these modifications is parametrized by a departure function $\Phi : \{1, \ldots, N\} \to \{1, \ldots, N\}$, and we ask that none of them is significantly better than the original strategy.

Formally, a strategy is said to suffer no internal regret whenever

$$\sum_{t=1}^{n} \ell\left(f_{I_t,t}, y_t\right) - \min_{\Phi} \sum_{t=1}^{n} \ell\left(f_{\Phi(I_t),t}, y_t\right) = o(n) \quad \text{a.s.}$$

Note that (external) regret corresponds to the departures $\Phi \equiv j$ for some j.

No-internal-regret strategies are so good (or bad!) that they may not be improved easily. (No internal regret corresponds to an extremum of performance.)

A refined notion of regret

Internal regret, a notion introduced by Foster and Vohra '98, is concerned with consistent modifications of our prediction strategy.

Each of these modifications is parametrized by a departure function $\Phi : \{1, \dots, N\} \to \{1, \dots, N\}$, and we ask that none of them is significantly better than the original strategy.

Formally, a strategy is said to suffer no internal regret whenever

$$\sum_{t=1}^{n} \ell\left(f_{I_t,t}, y_t\right) - \min_{\Phi} \sum_{t=1}^{n} \ell\left(f_{\Phi(I_t),t}, y_t\right) = o(n) \quad \text{a.s.}$$

The notion of internal regret has been shown to be useful in the theory of equilibria in repeated games.

A refined notion of regret

Internal regret, a notion introduced by Foster and Vohra '98, is concerned with consistent modifications of our prediction strategy.

Each of these modifications is parametrized by a departure function $\Phi : \{1, \ldots, N\} \to \{1, \ldots, N\}$, and we ask that none of them is significantly better than the original strategy.

Formally, a strategy is said to suffer no internal regret whenever

$$\sum_{t=1}^{n} \ell\left(f_{I_t,t}, y_t\right) - \min_{\Phi} \sum_{t=1}^{n} \ell\left(f_{\Phi(I_t),t}, y_t\right) = o(n) \quad \text{a.s.}$$

The notion of internal regret has been shown to be useful in the theory of equilibria in repeated games.

Foster and Vohra '97, '99 showed that if all players of a finite game play strategies that suffer no internal regret, then the empirical frequencies of played profiles converge to the set of correlated equilibria.

A refined notion of regret

Internal regret, a notion introduced by Foster and Vohra ’98, is concerned with consistent modifications of our prediction strategy.

Each of these modifications is parametrized by a departure function $\Phi : \{1, \ldots, N\} \to \{1, \ldots, N\}$, and we ask that none of them is significantly better than the original strategy.

Formally, a strategy is said to suffer no internal regret whenever

$$\sum_{t=1}^{n} \ell\left(f_{I_t,t}, y_t\right) - \min_{\Phi} \sum_{t=1}^{n} \ell\left(f_{\Phi(I_t),t}, y_t\right) = o(n) \quad \text{a.s.}$$

The notion of internal regret has been shown to be useful in the theory of equilibria in repeated games.

See also Fudenberg and Levine ’99, Hart and Mas-Colell ’99.

A refined notion of regret

Internal regret, a notion introduced by Foster and Vohra ’98, is concerned with consistent modifications of our prediction strategy.

Each of these modifications is parametrized by a departure function $\Phi : \{1, \ldots, N\} \to \{1, \ldots, N\}$, and we ask that none of them is significantly better than the original strategy.

Formally, a strategy is said to suffer no internal regret whenever

$$\sum_{t=1}^{n} \ell\left(f_{I_t,t}, y_t\right) - \min_{\Phi} \sum_{t=1}^{n} \ell\left(f_{\Phi(I_t),t}, y_t\right) = o(n) \quad \text{a.s.}$$

Once again, we may focus on expected internal regret,

$$\sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{\varphi} \sum_{t=1}^{n} \ell\left(\varphi(\boldsymbol{p}_t), y_t\right)$$

A refined notion of regret

Internal regret, a notion introduced by Foster and Vohra ’98, is concerned with consistent modifications of our prediction strategy.

Each of these modifications is parametrized by a departure function $\Phi : \{1, \ldots, N\} \to \{1, \ldots, N\}$, and we ask that none of them is significantly better than the original strategy.

Formally, a strategy is said to suffer no internal regret whenever

$$\sum_{t=1}^{n} \ell\left(f_{I_t,t}, y_t\right) - \min_{\Phi} \sum_{t=1}^{n} \ell\left(f_{\Phi(I_t),t}, y_t\right) = o(n) \quad \text{a.s.}$$

Once again, we may focus on expected internal regret,

$$\sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{\varphi} \sum_{t=1}^{n} \ell\left(\varphi(\boldsymbol{p}_t), y_t\right)$$

where the minimum is over all linear functions φ from the simplex of order N into itself.

Link between external and internal regret

To minimize expected internal regret,

$$\sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{\varphi} \sum_{t=1}^{n} \ell\left(\varphi(\boldsymbol{p}_t), y_t\right)$$

Link between external and internal regret

To minimize expected internal regret,

$$\sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{\varphi} \sum_{t=1}^{n} \ell\left(\varphi(\boldsymbol{p}_t), y_t\right)$$

it suffices to consider only the "extremal" φ given by $\varphi(\boldsymbol{p}) = \boldsymbol{p}^{i \to j}$, where

Link between external and internal regret

To minimize expected internal regret,

$$\sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{\varphi} \sum_{t=1}^{n} \ell\left(\varphi(\boldsymbol{p}_t), y_t\right)$$

it suffices to consider only the “extremal” φ given by $\varphi(\boldsymbol{p}) = \boldsymbol{p}^{i \to j}$, where

$$p_k^{i \to j} = \begin{cases} 0 & \text{if } k = i \\ p_i + p_j & \text{if } k = j \\ p_k & \text{otherwise} \end{cases}$$

Link between external and internal regret

To minimize expected internal regret,

$$\sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{\varphi} \sum_{t=1}^{n} \ell\left(\varphi(\boldsymbol{p}_t), y_t\right)$$

it suffices to minimize

$$\begin{aligned}
\max_{i \neq j} R_n^{i \to j} &= \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{i \neq j} \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t^{i \to j}, y_t\right) \\
&= \max_{i \neq j} \sum_{t=1}^{n} p_{i,t}\left(\ell\left(f_{i,t}, y_t\right) - \ell\left(f_{j,t}, y_t\right)\right)
\end{aligned}$$

Link between external and internal regret

To minimize expected internal regret,

$$\sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{\varphi} \sum_{t=1}^{n} \ell\left(\varphi(\boldsymbol{p}_t), y_t\right)$$

it suffices to minimize

$$\max_{i \neq j} R_n^{i \to j} = \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{i \neq j} \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t^{i \to j}, y_t\right)$$

$$= \max_{i \neq j} \sum_{t=1}^{n} p_{i,t}\left(\ell\left(f_{i,t}, y_t\right) - \ell\left(f_{j,t}, y_t\right)\right)$$

(Expected) external regret is less than N times (expected) internal regret,

Link between external and internal regret

To minimize expected internal regret,

$$\sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{\varphi} \sum_{t=1}^{n} \ell\left(\varphi(\boldsymbol{p}_t), y_t\right)$$

it suffices to minimize

$$\begin{aligned} \max_{i \neq j} R_n^{i \to j} &= \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{i \neq j} \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t^{i \to j}, y_t\right) \\ &= \max_{i \neq j} \sum_{t=1}^{n} p_{i,t}\left(\ell\left(f_{i,t}, y_t\right) - \ell\left(f_{j,t}, y_t\right)\right) \end{aligned}$$

(Expected) external regret is less than N times (expected) internal regret,

external regret $\leqslant N \times$ internal regret

Link between external and internal regret

To minimize expected internal regret,

$$\sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{\varphi} \sum_{t=1}^{n} \ell\left(\varphi(\boldsymbol{p}_t), y_t\right)$$

it suffices to minimize

$$\max_{i \neq j} R_n^{i \to j} = \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{i \neq j} \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t^{i \to j}, y_t\right)$$

$$= \max_{i \neq j} \sum_{t=1}^{n} p_{i,t}\left(\ell\left(f_{i,t}, y_t\right) - \ell\left(f_{j,t}, y_t\right)\right)$$

(Expected) external regret is less than N times (expected) internal regret,

$$\max_{j=1,\ldots,N} \sum_{i=1}^{N} R_n^{i \to j} \leqslant N \times \text{internal regret}$$

Link between external and internal regret

To minimize expected internal regret,

$$\sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{\varphi} \sum_{t=1}^{n} \ell\left(\varphi(\boldsymbol{p}_t), y_t\right)$$

it suffices to minimize

$$\max_{i \neq j} R_n^{i \to j} = \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{i \neq j} \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t^{i \to j}, y_t\right)$$

$$= \max_{i \neq j} \sum_{t=1}^{n} p_{i,t}\left(\ell\left(f_{i,t}, y_t\right) - \ell\left(f_{j,t}, y_t\right)\right)$$

(Expected) external regret is less than N times (expected) internal regret,

$$\max_{j=1,\ldots,N} \sum_{i=1}^{N} R_n^{i \to j} \leqslant N \times \max_{i \neq j} R_n^{i \to j}$$

Link between external and internal regret

To minimize expected internal regret,

$$\sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{\varphi} \sum_{t=1}^{n} \ell\left(\varphi(\boldsymbol{p}_t), y_t\right)$$

it suffices to minimize

$$\begin{aligned} \max_{i \neq j} R_n^{i \to j} &= \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{i \neq j} \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t^{i \to j}, y_t\right) \\ &= \max_{i \neq j} \sum_{t=1}^{n} p_{i,t}\left(\ell\left(f_{i,t}, y_t\right) - \ell\left(f_{j,t}, y_t\right)\right) \end{aligned}$$

(Expected) external regret is less than N times (expected) internal regret,

$$\max_{j=1,\ldots,N} \sum_{i=1}^{N} R_n^{i \to j} \leqslant N \times \max_{i \neq j} R_n^{i \to j}$$

Thus, no-internal-regret strategies also suffer no external regret.

Link between external and internal regret

To minimize expected internal regret,

$$\sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{\varphi} \sum_{t=1}^{n} \ell\left(\varphi(\boldsymbol{p}_t), y_t\right)$$

it suffices to minimize

$$\begin{aligned} \max_{i \neq j} R_n^{i \to j} &= \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t, y_t\right) - \min_{i \neq j} \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t^{i \to j}, y_t\right) \\ &= \max_{i \neq j} \sum_{t=1}^{n} p_{i,t} \left(\ell\left(f_{i,t}, y_t\right) - \ell\left(f_{j,t}, y_t\right)\right) \end{aligned}$$

(Expected) external regret is less than N times (expected) internal regret,

$$\max_{j=1,\ldots,N} \sum_{i=1}^{N} R_n^{i \to j} \leqslant N \times \max_{i \neq j} R_n^{i \to j}$$

But no-external-regret strategies may suffer a large internal regret, see EWA.

From external to internal regret

Specific no internal regret strategies have to be constructed.

From external to internal regret

By a conversion trick, such strategies may be obtained from no-external-regret ones.

From external to internal regret

By a conversion trick, such strategies may be obtained from no-external-regret ones.

Blum and Mansour '05 found independently another conversion procedure, that relies on the linearity of the expected losses $\ell(\boldsymbol{p}_t, y_t)$ in $\boldsymbol{p}_t$, but offers some computational advantages.

From external to internal regret

By a conversion trick, such strategies may be obtained from no-external-regret ones.

Define $N(N-1)$ fictitious experts by their losses at time instants $1 \leqslant s \leqslant t-1$, which equal $\ell(\boldsymbol{p}_s^{i \to j}, y_s)$, $i \neq j$.

From external to internal regret

By a conversion trick, such strategies may be obtained from no-external-regret ones.

Define $N(N-1)$ fictitious experts by their losses at time instants $1 \leqslant s \leqslant t-1$, which equal $\ell(\boldsymbol{p}_s^{i \to j}, y_s)$, $i \neq j$.

Choose a probability distribution Δ_t over the pairs $i \neq j$ by EWA on this pool of fictitious experts, that is,

From external to internal regret

By a conversion trick, such strategies may be obtained from no-external-regret ones.

Define $N(N-1)$ fictitious experts by their losses at time instants $1 \leqslant s \leqslant t-1$, which equal $\ell(\boldsymbol{p}_s^{i \to j}, y_s)$, $i \neq j$.

Choose a probability distribution Δ_t over the pairs $i \neq j$ by EWA on this pool of fictitious experts, that is,

$$\Delta_{(i,j),t} = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \ell\left(\boldsymbol{p}_s^{i \to j}, y_s\right)\right)}{\sum_{(k,l):k \neq l} \exp\left(-\eta \sum_{s=1}^{t-1} \ell\left(\boldsymbol{p}_s^{k \to l}, y_s\right)\right)}$$

From external to internal regret

By a conversion trick, such strategies may be obtained from no-external-regret ones.

Define $N(N-1)$ fictitious experts by their losses at time instants $1 \leqslant s \leqslant t-1$, which equal $\ell(\boldsymbol{p}_s^{i \to j}, y_s)$, $i \neq j$.

Choose a probability distribution Δ_t over the pairs $i \neq j$ by EWA on this pool of fictitious experts, that is,

$$\Delta_{(i,j),t} = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \ell\left(\boldsymbol{p}_s^{i \to j}, y_s\right)\right)}{\sum_{(k,l):k \neq l} \exp\left(-\eta \sum_{s=1}^{t-1} \ell\left(\boldsymbol{p}_s^{k \to l}, y_s\right)\right)}$$

with $\eta = B^{-1}\sqrt{8(\ln N(N-1))/n}$ (the losses are assumed to be bounded by B)

From external to internal regret

By a conversion trick, such strategies may be obtained from no-external-regret ones.

Define $N(N-1)$ fictitious experts by their losses at time instants $1 \leqslant s \leqslant t-1$, which equal $\ell(\boldsymbol{p}_s^{i \to j}, y_s)$, $i \neq j$.

Choose a probability distribution Δ_t over the pairs $i \neq j$ by EWA on this pool of fictitious experts, that is,

$$\Delta_{(i,j),t} = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \ell\left(\boldsymbol{p}_s^{i \to j}, y_s\right)\right)}{\sum_{(k,l):k \neq l} \exp\left(-\eta \sum_{s=1}^{t-1} \ell\left(\boldsymbol{p}_s^{k \to l}, y_s\right)\right)}$$

Then, the analysis of EWA ensures

From external to internal regret

By a conversion trick, such strategies may be obtained from no-external-regret ones.

Define $N(N-1)$ fictitious experts by their losses at time instants $1 \leqslant s \leqslant t-1$, which equal $\ell(\boldsymbol{p}_s^{i \to j}, y_s)$, $i \neq j$.

Choose a probability distribution Δ_t over the pairs $i \neq j$ by EWA on this pool of fictitious experts, that is,

$$\Delta_{(i,j),t} = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \ell\left(\boldsymbol{p}_s^{i \to j}, y_s\right)\right)}{\sum_{(k,l):k \neq l} \exp\left(-\eta \sum_{s=1}^{t-1} \ell\left(\boldsymbol{p}_s^{k \to l}, y_s\right)\right)}$$

Then, the analysis of EWA ensures

$$\sum_{t=1}^{n} \sum_{i \neq j} \Delta_{(i,j),t} \ell\left(\boldsymbol{p}_t^{i \to j}, y_t\right) \leqslant \min_{i \neq j} \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t^{i \to j}, y_t\right) + B\sqrt{n \ln N}$$

From external to internal regret

By a conversion trick, such strategies may be obtained from no-external-regret ones.

Define $N(N-1)$ fictitious experts by their losses at time instants $1 \leqslant s \leqslant t-1$, which equal $\ell(\boldsymbol{p}_s^{i \to j}, y_s)$, $i \neq j$.

Choose a probability distribution Δ_t over the pairs $i \neq j$ by EWA on this pool of fictitious experts, that is,

$$\Delta_{(i,j),t} = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \ell\left(\boldsymbol{p}_s^{i \to j}, y_s\right)\right)}{\sum_{(k,l):k \neq l} \exp\left(-\eta \sum_{s=1}^{t-1} \ell\left(\boldsymbol{p}_s^{k \to l}, y_s\right)\right)}$$

Then, the analysis of EWA ensures

$$\sum_{t=1}^{n} \ell\left(\sum_{i \neq j} \Delta_{(i,j),t}\, \boldsymbol{p}_t^{i \to j}, y_t\right) \leqslant \min_{i \neq j} \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t^{i \to j}, y_t\right) + B\sqrt{n \ln N}$$

From external to internal regret

By a conversion trick, such strategies may be obtained from no-external-regret ones.

Define $N(N-1)$ fictitious experts by their losses at time instants $1 \leqslant s \leqslant t-1$, which equal $\ell(\boldsymbol{p}_s^{i \to j}, y_s)$, $i \neq j$.

Choose a probability distribution Δ_t over the pairs $i \neq j$ by EWA on this pool of fictitious experts, that is,

$$\Delta_{(i,j),t} = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \ell\left(\boldsymbol{p}_s^{i \to j}, y_s\right)\right)}{\sum_{(k,l):k \neq l} \exp\left(-\eta \sum_{s=1}^{t-1} \ell\left(\boldsymbol{p}_s^{k \to l}, y_s\right)\right)}$$

Then, the analysis of EWA ensures

$$\sum_{t=1}^{n} \ell\left(\sum_{i \neq j} \Delta_{(i,j),t}\, \boldsymbol{p}_t^{i \to j}, y_t\right) \leqslant \min_{i \neq j} \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t^{i \to j}, y_t\right) + B\sqrt{n \ln N}$$

As soon as we have $\boldsymbol{p}_t$ in the left-hand side, this is a bound on internal regret.

From external to internal regret

By a conversion trick, such strategies may be obtained from no-external-regret ones.

Define $N(N-1)$ fictitious experts by their losses at time instants $1 \leqslant s \leqslant t-1$, which equal $\ell(\boldsymbol{p}_s^{i \to j}, y_s)$, $i \neq j$.

Choose a probability distribution Δ_t over the pairs $i \neq j$ by EWA on this pool of fictitious experts, that is,

$$\Delta_{(i,j),t} = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \ell\left(\boldsymbol{p}_s^{i \to j}, y_s\right)\right)}{\sum_{(k,l):k \neq l} \exp\left(-\eta \sum_{s=1}^{t-1} \ell\left(\boldsymbol{p}_s^{k \to l}, y_s\right)\right)}$$

Then, the analysis of EWA ensures

$$\sum_{t=1}^{n} \ell\left(\sum_{i \neq j} \Delta_{(i,j),t}\, \boldsymbol{p}_t^{i \to j}, y_t\right) \leqslant \min_{i \neq j} \sum_{t=1}^{n} \ell\left(\boldsymbol{p}_t^{i \to j}, y_t\right) + B\sqrt{n \ln N}$$

We choose $\boldsymbol{p}_t$ such that $\boldsymbol{p}_t = \sum_{(i,j):i \neq j} \Delta_{(i,j),t}\, \boldsymbol{p}_t^{i \to j}$.

Applications of the conversion trick

In on-line portfolio selection, we could define a notion of internal regret, and introduce algorithms that minimize it.

Applications of the conversion trick

In on-line portfolio selection, we could define a notion of internal regret, and introduce algorithms that minimize it.

For instance, we defined the B1EXP strategy, which is the no-internal-regret counterpart of the EG strategy of Helmbold, Schapire, Singer, and Warmuth '98.

Applications of the conversion trick

In on-line portfolio selection, we could define a notion of internal regret, and introduce algorithms that minimize it.

It turns out that internal regret is linked to an increased stability with respect to bad choices of the tuning parameters,

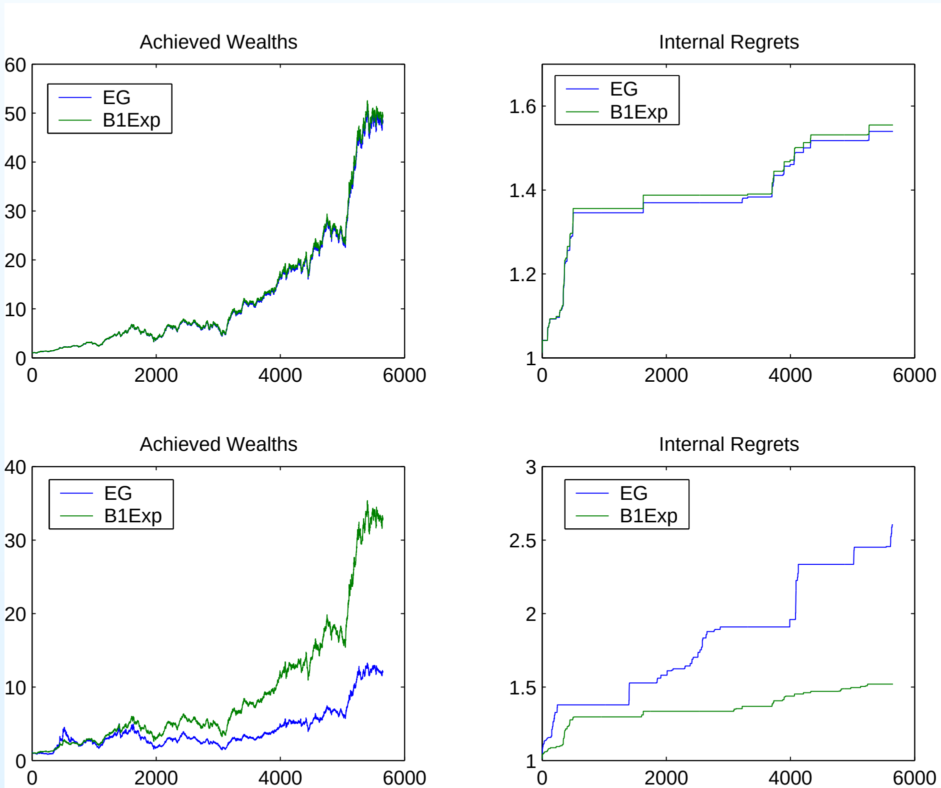

Applications of the conversion trick

In on-line portfolio selection, we could define a notion of internal regret, and introduce algorithms that minimize it.

As for repeated games, we could also extend Foster and Vohra's convergence result toward correlated equilibria to games with infinite (convex, compact) sets of strategies.