

CSE 541: Interactive Learning

Lecture 2: Squared loss, portfolios, and online convex optimization

Kevin Jamieson

Paul G. Allen School of Computer Science & Engineering
University of Washington

Monday, October 5, 2026



Last time, in one slide

Online learning with d actions

for $t = 1, 2, \dots, T$: player picks $p_t \in \Delta_d$ and plays $I_t \sim p_t$;

adversary reveals $\ell_t \in [0, 1]^d$; player pays $\ell_t(I_t)$

Classification: $d = |\mathcal{H}|$, $\ell_t(i) = \mathbf{1}\{h_i(x_t) \neq y_t\}$

Exponential weights

$w_1 \equiv 1$; $p_t(i) = w_t(i)/W_t$, $W_t = \sum_i w_t(i)$;

$w_{t+1}(i) = w_t(i) e^{-\eta \ell_t(i)}$

Theorem

For any $\ell_1, \dots, \ell_T \in [0, 1]^d$,

$\text{Regret}_T = \sum_t \langle p_t, \ell_t \rangle - \min_i \sum_t \ell_t(i) \leq \frac{\log d}{\eta} + \frac{\eta T}{8}$,

so $\eta = \sqrt{8 \log d / T}$ gives $\sqrt{T \log(d)/2}$.

The proof, in two lines. Track

$\log(W_{T+1}/W_1)$:

- *Upper:* $\frac{W_{t+1}}{W_t} = \mathbb{E}_{I \sim p_t}[e^{-\eta \ell_t(I)}]$, and Hoeffding's lemma gives $\log \mathbb{E}_{I \sim p_t}[e^{-\eta \ell_t(I)}] \leq -\eta \langle p_t, \ell_t \rangle + \frac{\eta^2}{8}$.
- *Lower:* $W_{T+1} \geq w_{T+1}(i) = e^{-\eta \sum_t \ell_t(i)}$ for every i .

Two messages from last time

Dropping IID cost nothing: same $\log d$, same $\sqrt{T}$ as ERM under IID.

The $\eta^2/8$ slack per round is exactly what becomes $\sqrt{T}$. Today: losses where that slack disappears.

Squared loss

Fast rates, in the batch and online worlds

Regression with a finite class

- $(X, Y) \sim \nu$ on $\mathcal{X} \times [0, 1]$; a finite class $\mathcal{F}$ of predictors $f : \mathcal{X} \rightarrow [0, 1]$.
- Risk and empirical risk: $L(f) = \mathbb{E}[(f(X) - Y)^2]$, $\hat{L}_n(f) = \frac{1}{n} \sum_{i=1}^n (f(x_i) - y_i)^2$.
- ERM $\hat{f} \in \arg \min_{f \in \mathcal{F}} \hat{L}_n(f)$; best in class $f_{\mathcal{F}}^* \in \arg \min_{f \in \mathcal{F}} L(f)$.

Regression with a finite class

- $(X, Y) \sim \nu$ on $\mathcal{X} \times [0, 1]$; a finite class $\mathcal{F}$ of predictors $f : \mathcal{X} \rightarrow [0, 1]$.
- Risk and empirical risk: $L(f) = \mathbb{E}[(f(X) - Y)^2]$, $\hat{L}_n(f) = \frac{1}{n} \sum_{i=1}^n (f(x_i) - y_i)^2$.
- ERM $\hat{f} \in \arg \min_{f \in \mathcal{F}} \hat{L}_n(f)$; best in class $f_{\mathcal{F}}^* \in \arg \min_{f \in \mathcal{F}} L(f)$.

Bayes predictor

$f^*(x) = \mathbb{E}[Y \mid X = x]$ minimizes L over *all* functions, and for any f ,
 $L(f) = L(f^*) + \mathbb{E}[(f(X) - f^*(X))^2]$ (the cross term vanishes by the tower property). In general $f^* \notin \mathcal{F}$.

Regression with a finite class

- $(X, Y) \sim \nu$ on $\mathcal{X} \times [0, 1]$; a finite class $\mathcal{F}$ of predictors $f : \mathcal{X} \rightarrow [0, 1]$.
- Risk and empirical risk: $L(f) = \mathbb{E}[(f(X) - Y)^2]$, $\hat{L}_n(f) = \frac{1}{n} \sum_{i=1}^n (f(x_i) - y_i)^2$.
- ERM $\hat{f} \in \arg \min_{f \in \mathcal{F}} \hat{L}_n(f)$; best in class $f_{\mathcal{F}}^* \in \arg \min_{f \in \mathcal{F}} L(f)$.

Bayes predictor

$f^*(x) = \mathbb{E}[Y \mid X = x]$ minimizes L over *all* functions, and for any f ,
 $L(f) = L(f^*) + \mathbb{E}[(f(X) - f^*(X))^2]$ (the cross term vanishes by the tower property). In general $f^* \notin \mathcal{F}$.

Slow rate. $(f(x) - y)^2 \in [0, 1]$, so Hoeffding plus the union bound gives, exactly as for classification,
 $L(\hat{f}) - L(f_{\mathcal{F}}^*) \leq \sqrt{2 \log(2|\mathcal{F}|/\delta)/n}$ w.p. $\geq 1 - \delta$. Only boundedness was used.

Regression with a finite class

- $(X, Y) \sim \nu$ on $\mathcal{X} \times [0, 1]$; a finite class $\mathcal{F}$ of predictors $f : \mathcal{X} \rightarrow [0, 1]$.
- Risk and empirical risk: $L(f) = \mathbb{E}[(f(X) - Y)^2]$, $\hat{L}_n(f) = \frac{1}{n} \sum_{i=1}^n (f(x_i) - y_i)^2$.
- ERM $\hat{f} \in \arg \min_{f \in \mathcal{F}} \hat{L}_n(f)$; best in class $f_{\mathcal{F}}^* \in \arg \min_{f \in \mathcal{F}} L(f)$.

Bayes predictor

$f^*(x) = \mathbb{E}[Y \mid X = x]$ minimizes L over *all* functions, and for any f ,
 $L(f) = L(f^*) + \mathbb{E}[(f(X) - f^*(X))^2]$ (the cross term vanishes by the tower property). In general $f^* \notin \mathcal{F}$.

Slow rate. $(f(x) - y)^2 \in [0, 1]$, so Hoeffding plus the union bound gives, exactly as for classification, $L(\hat{f}) - L(f_{\mathcal{F}}^*) \leq \sqrt{2 \log(2|\mathcal{F}|/\delta)/n}$ w.p. $\geq 1 - \delta$. Only boundedness was used.

Classification had a **fast rate** $\log |\mathcal{H}|/n$ when realizable; is there an analogue here? We need a concentration inequality that sees the *variance*.

Bernstein's inequality and self-bounding variance

Lemma (Bernstein's inequality)

Let $X_1, \dots, X_n$ be independent with $|X_i| \leq B$ and $\frac{1}{n} \sum_i \text{Var}(X_i) \leq \sigma^2$. With probability at least $1 - \delta$,

$$\left| \frac{1}{n} \sum_{i=1}^n (X_i - \mathbb{E}[X_i]) \right| \leq \sqrt{\frac{2\sigma^2 \log(2/\delta)}{n}} + \frac{2B \log(2/\delta)}{3n}.$$

Hoeffding's $(b - a)^2/4$ is replaced by the true variance σ^2 ; when σ^2 is small the $1/n$ term dominates.

Bernstein's inequality and self-bounding variance

Lemma (Bernstein's inequality)

Let $X_1, \dots, X_n$ be independent with $|X_i| \leq B$ and $\frac{1}{n} \sum_i \text{Var}(X_i) \leq \sigma^2$. With probability at least $1 - \delta$,

$$\left| \frac{1}{n} \sum_{i=1}^n (X_i - \mathbb{E}[X_i]) \right| \leq \sqrt{\frac{2\sigma^2 \log(2/\delta)}{n}} + \frac{2B \log(2/\delta)}{3n}.$$

Hoeffding's $(b-a)^2/4$ is replaced by the true variance σ^2 ; when σ^2 is small the $1/n$ term dominates.

Self-bounding variance

For $f : \mathcal{X} \rightarrow [0, 1]$ and $Y \in [0, 1]$,

$$\text{Var}[(f(X) - Y)^2] \leq \mathbb{E}[(f(X) - Y)^4] \leq L(f).$$

Proof: $z \in [0, 1] \Rightarrow z^2 \leq z$.

Consequence. With $\beta = \log(2|\mathcal{F}|/\delta)$, w.p. $\geq 1 - \delta$, for all $f \in \mathcal{F}$ at once,

$$|L(f) - \hat{L}_n(f)| \leq \sqrt{\frac{2L(f)\beta}{n}} + \frac{2\beta}{3n}.$$

Good predictors concentrate faster.

Bernstein's inequality and self-bounding variance

Lemma (Bernstein's inequality)

Let $X_1, \dots, X_n$ be independent with $|X_i| \leq B$ and $\frac{1}{n} \sum_i \text{Var}(X_i) \leq \sigma^2$. With probability at least $1 - \delta$,

$$\left| \frac{1}{n} \sum_{i=1}^n (X_i - \mathbb{E}[X_i]) \right| \leq \sqrt{\frac{2\sigma^2 \log(2/\delta)}{n}} + \frac{2B \log(2/\delta)}{3n}.$$

Hoeffding's $(b-a)^2/4$ is replaced by the true variance σ^2 ; when σ^2 is small the $1/n$ term dominates.

Self-bounding variance

For $f : \mathcal{X} \rightarrow [0, 1]$ and $Y \in [0, 1]$,

$$\text{Var}[(f(X) - Y)^2] \leq \mathbb{E}[(f(X) - Y)^4] \leq L(f).$$

Proof: $z \in [0, 1] \Rightarrow z^2 \leq z$.

Same mechanism as classification: $\text{Var}[\mathbf{1}\{h(X) \neq Y\}] \leq R(h)$, and $R(h_*) = 0$ is what made the realizable rate $\log |\mathcal{H}|/n$ possible.

Consequence. With $\beta = \log(2|\mathcal{F}|/\delta)$, w.p. $\geq 1 - \delta$, for all $f \in \mathcal{F}$ at once,

$$|L(f) - \hat{L}_n(f)| \leq \sqrt{\frac{2L(f)\beta}{n}} + \frac{2\beta}{3n}.$$

Good predictors concentrate faster.

Theorem (fast-rate ERM)

Let $\beta = \log(2|\mathcal{F}|/\delta)$. With probability at least $1 - \delta$,

$$L(\hat{f}) - L(f_{\mathcal{F}}^*) \leq 4\sqrt{\frac{2L(f_{\mathcal{F}}^*)\beta}{n}} + \frac{8\beta}{n}.$$

Fast rate for ERM with squared loss

Theorem (fast-rate ERM)

Let $\beta = \log(2|\mathcal{F}|/\delta)$. With probability at least $1 - \delta$,

$$L(\hat{f}) - L(f_{\mathcal{F}}^*) \leq 4\sqrt{\frac{2L(f_{\mathcal{F}}^*)\beta}{n}} + \frac{8\beta}{n}.$$

Proof idea (algebra in the notes). On the Bernstein event, apply the bound to $\hat{f}$ and to $f_{\mathcal{F}}^*$ and use $\hat{L}_n(\hat{f}) \leq \hat{L}_n(f_{\mathcal{F}}^*)$ to get $a \leq b + \alpha\sqrt{b} + \alpha\sqrt{a} + c$ with $a = L(\hat{f})$, $b = L(f_{\mathcal{F}}^*)$, $\alpha = \sqrt{2\beta/n}$, $c = \frac{4\beta}{3n}$. The unknown a is on both sides: solve the quadratic in $\sqrt{a}$, $\sqrt{a} \leq \sqrt{b} + \alpha + \sqrt{c}$, and square. ■

Theorem (fast-rate ERM)

Let $\beta = \log(2|\mathcal{F}|/\delta)$. With probability at least $1 - \delta$,

$$L(\hat{f}) - L(f_{\mathcal{F}}^*) \leq 4\sqrt{\frac{2L(f_{\mathcal{F}}^*)\beta}{n}} + \frac{8\beta}{n}.$$

Proof idea (algebra in the notes). On the Bernstein event, apply the bound to $\hat{f}$ and to $f_{\mathcal{F}}^*$ and use $\hat{L}_n(\hat{f}) \leq \hat{L}_n(f_{\mathcal{F}}^*)$ to get $a \leq b + \alpha\sqrt{b} + \alpha\sqrt{a} + c$ with $a = L(\hat{f})$, $b = L(f_{\mathcal{F}}^*)$, $\alpha = \sqrt{2\beta/n}$, $c = \frac{4\beta}{3n}$. The unknown a is on both sides: solve the quadratic in $\sqrt{a}$, $\sqrt{a} \leq \sqrt{b} + \alpha + \sqrt{c}$, and square. ■

Read it three ways

- $L(f_{\mathcal{F}}^*) = 0$: rate $O(\log |\mathcal{F}|/n)$, the realizable rate, even though $f^* = \mathbb{E}[Y | X]$ need not lie in $\mathcal{F}$.
- $L(f_{\mathcal{F}}^*) = \rho$ small: $O(\sqrt{\rho \log |\mathcal{F}|/n} + \log |\mathcal{F}|/n)$, between the two.
- $L(f_{\mathcal{F}}^*) = \Theta(1)$: back to the slow rate $O(\sqrt{\log |\mathcal{F}|/n})$.

Online regression with squared loss

Online regression

Input: finite $\mathcal{F}$ of $f : \mathcal{X} \rightarrow [0, 1]$

for $t = 1, 2, \dots, T$:

x_t arrives

Learner predicts $\hat{y}_t \in [0, 1]$

$y_t \in [0, 1]$ is revealed

Learner suffers $(\hat{y}_t - y_t)^2$

Goal

$$\text{Regret}_T = \sum_{t=1}^T (\hat{y}_t - y_t)^2 - \min_{f \in \mathcal{F}} \sum_{t=1}^T (f(x_t) - y_t)^2,$$

for every sequence (x_t, y_t) .

What we already know. $(f(x_t) - y_t)^2 \in [0, 1]$,
so exponential weights over $\mathcal{F}$ (sample $f_t \sim p_t$,
predict $f_t(x_t)$) has expected regret
 $\leq \sqrt{T \log(|\mathcal{F}|)/2}$.

Online regression with squared loss

Online regression

Input: finite $\mathcal{F}$ of $f : \mathcal{X} \rightarrow [0, 1]$

for $t = 1, 2, \dots, T$:

x_t arrives

Learner predicts $\hat{y}_t \in [0, 1]$

$y_t \in [0, 1]$ is revealed

Learner suffers $(\hat{y}_t - y_t)^2$

Goal

$$\text{Regret}_T = \sum_{t=1}^T (\hat{y}_t - y_t)^2 - \min_{f \in \mathcal{F}} \sum_{t=1}^T (f(x_t) - y_t)^2,$$

for every sequence (x_t, y_t) .

What we already know. $(f(x_t) - y_t)^2 \in [0, 1]$, so exponential weights over $\mathcal{F}$ (sample $f_t \sim p_t$, predict $f_t(x_t)$) has expected regret $\leq \sqrt{T \log(|\mathcal{F}|)/2}$.

But nothing forces $\hat{y}_t$ to be some $f(x_t)$. We may predict any number in $[0, 1]$, e.g. the **average** of the class under p_t . This is **improper learning**: the predictor need not belong to $\mathcal{F}$.

Averaged exponential weights

Initialize: $w_1(f) = 1$ for all $f \in \mathcal{F}$

for $t = 1, 2, \dots, T$:

Predict $\hat{y}_t = \sum_f w_t(f) f(x_t) / \sum_f w_t(f)$

Observe y_t ; $w_{t+1}(f) = w_t(f) e^{-\eta(f(x_t) - y_t)^2}$

Mixability: the average beats every mixture

Theorem (averaged exponential weights)

With $\eta = \frac{1}{2}$, for every sequence $(x_1, y_1), \dots, (x_T, y_T)$ with $y_t \in [0, 1]$: $\text{Regret}_T \leq 2 \ln |\mathcal{F}|$.

No $\sqrt{T}$; no dependence on T at all. Deterministic, and with no assumption on the data.

Mixability: the average beats every mixture

Theorem (averaged exponential weights)

With $\eta = \frac{1}{2}$, for every sequence $(x_1, y_1), \dots, (x_T, y_T)$ with $y_t \in [0, 1]$: $\text{Regret}_T \leq 2 \ln |\mathcal{F}|$.

No $\sqrt{T}$; no dependence on T at all. Deterministic, and with no assumption on the data.

Proof. Steps 1–2 are exactly as for exponential weights: with $W_t = \sum_f w_t(f)$ and $p_t = w_t/W_t$,

$$-\eta \sum_{t=1}^T (f^*(x_t) - y_t)^2 - \ln |\mathcal{F}| \leq \ln \frac{W_{T+1}}{W_1} = \sum_{t=1}^T \ln \mathbb{E}_{f \sim p_t} \left[e^{-\eta (f(x_t) - y_t)^2} \right].$$

Mixability: the average beats every mixture

Theorem (averaged exponential weights)

With $\eta = \frac{1}{2}$, for every sequence $(x_1, y_1), \dots, (x_T, y_T)$ with $y_t \in [0, 1]$: $\text{Regret}_T \leq 2 \ln |\mathcal{F}|$.

No $\sqrt{T}$; no dependence on T at all. Deterministic, and with no assumption on the data.

Proof. Steps 1–2 are exactly as for exponential weights: with $W_t = \sum_f w_t(f)$ and $p_t = w_t/W_t$,

$$-\eta \sum_{t=1}^T (f^*(x_t) - y_t)^2 - \ln |\mathcal{F}| \leq \ln \frac{W_{T+1}}{W_1} = \sum_{t=1}^T \ln \mathbb{E}_{f \sim p_t} \left[e^{-\eta(f(x_t) - y_t)^2} \right].$$

Step 3 (mixability). For classification, Hoeffding's lemma bounded each term by $-\eta \mathbb{E}_{p_t}[\ell_t] + \eta^2/8$, and the $\eta^2/8$ slack per round is what became $\sqrt{T}$. Here there is no slack: $g(z) = e^{-\eta z^2}$ has $g''(z) = (4\eta^2 z^2 - 2\eta) e^{-\eta z^2} \leq 0$ for $|z| \leq 1/\sqrt{2\eta}$, so for $\eta \leq \frac{1}{2}$, g is concave on $[-1, 1] \ni f(x_t) - y_t$, and Jensen gives

$$\mathbb{E}_{f \sim p_t} [g(f(x_t) - y_t)] \leq g(\mathbb{E}_{f \sim p_t} [f(x_t)] - y_t) = e^{-\eta(\hat{y}_t - y_t)^2}.$$

Mixability: the average beats every mixture

Theorem (averaged exponential weights)

With $\eta = \frac{1}{2}$, for every sequence $(x_1, y_1), \dots, (x_T, y_T)$ with $y_t \in [0, 1]$: $\text{Regret}_T \leq 2 \ln |\mathcal{F}|$.

No $\sqrt{T}$; no dependence on T at all. Deterministic, and with no assumption on the data.

Proof. Steps 1–2 are exactly as for exponential weights: with $W_t = \sum_f w_t(f)$ and $p_t = w_t/W_t$,

$$-\eta \sum_{t=1}^T (f^*(x_t) - y_t)^2 - \ln |\mathcal{F}| \leq \ln \frac{W_{T+1}}{W_1} = \sum_{t=1}^T \ln \mathbb{E}_{f \sim p_t} \left[e^{-\eta(f(x_t) - y_t)^2} \right].$$

Step 3 (mixability). For classification, Hoeffding's lemma bounded each term by $-\eta \mathbb{E}_{p_t}[\ell_t] + \eta^2/8$, and the $\eta^2/8$ slack per round is what became $\sqrt{T}$. Here there is no slack: $g(z) = e^{-\eta z^2}$ has $g''(z) = (4\eta^2 z^2 - 2\eta) e^{-\eta z^2} \leq 0$ for $|z| \leq 1/\sqrt{2\eta}$, so for $\eta \leq \frac{1}{2}$, g is **concave** on $[-1, 1] \ni f(x_t) - y_t$, and Jensen gives

$$\mathbb{E}_{f \sim p_t} [g(f(x_t) - y_t)] \leq g(\mathbb{E}_{f \sim p_t} [f(x_t)] - y_t) = e^{-\eta(\hat{y}_t - y_t)^2}.$$

Combine: $-\eta \sum_t (f^*(x_t) - y_t)^2 - \ln |\mathcal{F}| \leq -\eta \sum_t (\hat{y}_t - y_t)^2$, so $\text{Regret}_T \leq \ln |\mathcal{F}|/\eta = 2 \ln |\mathcal{F}|$. ■

Proper vs. improper, and the parallel with the batch world

	Classification, 0/1 loss	Regression, squared loss
Batch, slow rate	$\sqrt{\log \mathcal{H} /n}$ (Hoeffding)	$\sqrt{\log \mathcal{F} /n}$ (Hoeffding)
Batch, fast rate	$\log \mathcal{H} /n$ if $R(h_*) = 0$ (var. $\leq$ mean)	$\log \mathcal{F} /n$ if $L(f_{\mathcal{F}}^*) = 0$ (var. $\leq$ mean)
Online	$\sqrt{T \log \mathcal{H} /2}$, randomized, proper	$2 \ln \mathcal{F} $, deterministic, improper

Proper vs. improper, and the parallel with the batch world


	Classification, 0/1 loss	Regression, squared loss
Batch, slow rate	$\sqrt{\log \mathcal{H} /n}$ (Hoeffding)	$\sqrt{\log \mathcal{F} /n}$ (Hoeffding)
Batch, fast rate	$\log \mathcal{H} /n$ if $R(h_*) = 0$ (var. $\leq$ mean)	$\log \mathcal{F} /n$ if $L(f_{\mathcal{F}}^*) = 0$ (var. $\leq$ mean)
Online	$\sqrt{T \log \mathcal{H} /2}$, randomized, proper	$2 \ln \mathcal{F} $, deterministic, improper

- **Improper** = the prediction is not made by any single $f \in \mathcal{F}$. That freedom is what buys the T -free bound; a proper learner must randomize and pays $\sqrt{T}$ (the $\ln |\mathcal{F}|/\eta$ vs. $\eta T/8$ trade-off).
- Why not average for the 0/1 loss? $e^{-\eta \mathbf{1}\{\cdot\}}$ is not concave in anything useful: the 0/1 loss is *not mixable*, and any deterministic predictor can be made wrong every round.
- Curvature is the common thread: it gives the fast rate in the batch world (through the variance) and mixability in the online world. Next: a loss with curvature where the comparator is a continuum.

Online portfolio optimization

The loss decides who you are competing with


Which stock would you have picked?



Daily closes for eight large US stocks, Sep 29, 2025 to Sep 25, 2026, each normalized to \$1 on day 1.

- The leader changes hands over and over; the year's winner was mid-pack for most of it.
- Halfway through, the eventual worst stock had been the best.

Which stock would you have picked?



Daily closes for eight large US stocks, Sep 29, 2025 to Sep 25, 2026, each normalized to \$1 on day 1.

- The leader changes hands over and over; the year's winner was mid-pack for most of it.
- Halfway through, the eventual worst stock had been the best.

Picking the best stock is hard *even in hindsight*, let alone in advance. Can we still be competitive with it, with no model of the market at all?

Investing without a model of the market

Setup. d stocks. On day t , stock i has *gross return* $r_t(i) = \text{price}_{t+1}(i)/\text{price}_t(i) \geq 0$.

A portfolio is $p_t \in \Delta_d$ (wealth fractions), rebalanced daily. Starting from wealth v_1 ,

$$v_{T+1} = v_1 \prod_{t=1}^T \langle p_t, r_t \rangle,$$

so the log-wealth is

$$\log(v_{T+1}/v_1) = \sum_{t=1}^T \log \langle p_t, r_t \rangle.$$

Investing without a model of the market

Setup. d stocks. On day t , stock i has *gross return* $r_t(i) = \text{price}_{t+1}(i)/\text{price}_t(i) \geq 0$.

A portfolio is $p_t \in \Delta_d$ (wealth fractions), rebalanced daily. Starting from wealth v_1 ,

$$v_{T+1} = v_1 \prod_{t=1}^T \langle p_t, r_t \rangle,$$

so the log-wealth is

$$\log(v_{T+1}/v_1) = \sum_{t=1}^T \log \langle p_t, r_t \rangle.$$

Classical (Markowitz): assume r_t IID with mean μ and covariance Σ estimated from history; solve $\min_{p \in \Delta_d} p^\top \Sigma p$ s.t. $p^\top \mu \geq \bar{r}$. IID fails exactly when it matters.

Investing without a model of the market

Setup. d stocks. On day t , stock i has gross return $r_t(i) = \text{price}_{t+1}(i)/\text{price}_t(i) \geq 0$.

A portfolio is $p_t \in \Delta_d$ (wealth fractions), rebalanced daily. Starting from wealth v_1 ,

$$v_{T+1} = v_1 \prod_{t=1}^T \langle p_t, r_t \rangle,$$

so the log-wealth is

$$\log(v_{T+1}/v_1) = \sum_{t=1}^T \log \langle p_t, r_t \rangle.$$

Classical (Markowitz): assume r_t IID with mean μ and covariance Σ estimated from history; solve $\min_{p \in \Delta_d} p^\top \Sigma p$ s.t. $p^\top \mu \geq \bar{r}$. IID fails exactly when it matters.

Online view

Treat $r_1, r_2, \dots$ as adversarial. Loss on day t : $\ell_t(p) = -\log \langle p, r_t \rangle$. Compete with a fixed comparator:

$$\max_{p \in \mathcal{C}} \sum_{t=1}^T \log \langle p, r_t \rangle - \sum_{t=1}^T \log \langle p_t, r_t \rangle.$$

Which comparator class $\mathcal{C}$?

- $\mathcal{C} = \{e_1, \dots, e_d\}$: the best single stock
- $\mathcal{C} = \Delta_d$: the best fixed (constantly rebalanced) portfolio

Competing with the best single stock

Theorem

Run exponential weights over the d stocks with $\eta = 1$ and losses $\ell_t(i) = -\log r_t(i)$, i.e., $w_{t+1}(i) = w_t(i) r_t(i)$ and $p_t(i) = w_t(i)/W_t$. Then for any return sequence,

$$\max_{i \in [d]} \sum_{t=1}^T \log r_t(i) - \sum_{t=1}^T \log \langle p_t, r_t \rangle \leq \log d.$$

Competing with the best single stock

Theorem

Run exponential weights over the d stocks with $\eta = 1$ and losses $\ell_t(i) = -\log r_t(i)$, i.e., $w_{t+1}(i) = w_t(i) r_t(i)$ and $p_t(i) = w_t(i)/W_t$. Then for any return sequence,

$$\max_{i \in [d]} \sum_{t=1}^T \log r_t(i) - \sum_{t=1}^T \log \langle p_t, r_t \rangle \leq \log d.$$

Proof. The losses are unbounded, but the potential argument goes through *exactly*, because $\langle p_t, r_t \rangle = \sum_i \frac{w_t(i)}{W_t} r_t(i) = \frac{W_{t+1}}{W_t}$. So

$$\sum_{t=1}^T \log \langle p_t, r_t \rangle = \log \frac{W_{T+1}}{W_1} \geq \log \frac{w_{T+1}(i^*)}{d} = \sum_{t=1}^T \log r_t(i^*) - \log d. \quad \blacksquare$$

The loss changes the comparator

Linear losses $p \mapsto \langle p, \ell_t \rangle$ (classification, experts):

$$\min_{p \in \Delta_d} \sum_t \langle p, \ell_t \rangle = \min_{i \in [d]} \sum_t \ell_t(i).$$

A linear function on the simplex is minimized at a vertex. Competing with the best mixture is the same as competing with the best expert.

Nothing gained.

Log losses $p \mapsto -\log \langle p, r_t \rangle$ are convex but not linear.

The loss changes the comparator

Linear losses $p \mapsto \langle p, \ell_t \rangle$ (classification, experts):

$$\min_{p \in \Delta_d} \sum_t \langle p, \ell_t \rangle = \min_{i \in [d]} \sum_t \ell_t(i).$$

A linear function on the simplex is minimized at a vertex. Competing with the best mixture is the same as competing with the best expert.

Nothing gained.

Log losses $p \mapsto -\log \langle p, r_t \rangle$ are convex but not linear.

Example: $d = 2$, alternating returns

$r_t = (2, \frac{1}{2})$ for odd t , $r_t = (\frac{1}{2}, 2)$ for even t .

- Each single stock: $\times 2$ then $\times \frac{1}{2}$. Wealth after T days: 1.
- Uniform $p = (\frac{1}{2}, \frac{1}{2})$: $\langle p, r_t \rangle = \frac{5}{4}$ every day. Wealth: $(5/4)^T$.

The best fixed *portfolio* beats the best stock by an exponential factor.

The loss changes the comparator

Linear losses $p \mapsto \langle p, \ell_t \rangle$ (classification, experts):

$$\min_{p \in \Delta_d} \sum_t \langle p, \ell_t \rangle = \min_{i \in [d]} \sum_t \ell_t(i).$$

A linear function on the simplex is minimized at a vertex. Competing with the best mixture is the same as competing with the best expert.

Nothing gained.

Log losses $p \mapsto -\log \langle p, r_t \rangle$ are convex but not linear.

Example: $d = 2$, alternating returns

$r_t = (2, \frac{1}{2})$ for odd t , $r_t = (\frac{1}{2}, 2)$ for even t .

- Each single stock: $\times 2$ then $\times \frac{1}{2}$. Wealth after T days: 1.
- Uniform $p = (\frac{1}{2}, \frac{1}{2})$: $\langle p, r_t \rangle = \frac{5}{4}$ every day. Wealth: $(5/4)^T$.

The best fixed *portfolio* beats the best stock by an exponential factor.

To compete with the best $p \in \Delta_d$ we need exponential weights over a *continuum* of hypotheses.

Exponential weights over a continuum

Continuous exponential weights (CEW)

Input: convex $\mathcal{A} \subseteq \mathbb{R}^d$, $\eta > 0$, base measure ν on $\mathcal{A}$
for $t = 1, 2, \dots, T$:

Sample $a_t \sim p_t$, $dp_t(a) \propto e^{-\eta \sum_{s < t} \ell_s(a)} d\nu(a)$

Observe ℓ_t ; suffer $\ell_t(a_t)$

Theorem (CEW regret)

If each $\ell_t : \mathcal{A} \rightarrow [0, 1]$ is convex and ν is Lebesgue, the expected regret $\mathbb{E}[\sum_t \ell_t(a_t)] - \min_{a \in \mathcal{A}} \sum_t \ell_t(a)$ is at most

$$\frac{d \log T}{\eta} + \frac{\eta T}{8} + 1 \stackrel{\eta = \sqrt{8d \log T / T}}{=} \sqrt{\frac{dT \log T}{2}} + 1.$$

Exponential weights over a continuum

Continuous exponential weights (CEW)

Input: convex $\mathcal{A} \subseteq \mathbb{R}^d$, $\eta > 0$, base measure ν on $\mathcal{A}$
for $t = 1, 2, \dots, T$:

Sample $a_t \sim p_t$, $dp_t(a) \propto e^{-\eta \sum_{s < t} \ell_s(a)} d\nu(a)$

Observe ℓ_t ; suffer $\ell_t(a_t)$

Theorem (CEW regret)

If each $\ell_t : \mathcal{A} \rightarrow [0, 1]$ is convex and ν is Lebesgue, the expected regret $\mathbb{E}[\sum_t \ell_t(a_t)] - \min_{a \in \mathcal{A}} \sum_t \ell_t(a)$ is at most

$$\frac{d \log T}{\eta} + \frac{\eta T}{8} + 1 \stackrel{\eta = \sqrt{8d \log T / T}}{=} \sqrt{\frac{dT \log T}{2}} + 1.$$

What changes in the proof. The upper bound on $\log(W_{T+1}/W_1)$ is identical. For the lower bound a single point a^* has measure zero, so integrate over the shrunken copy $\mathcal{A}^* = (1 - \frac{1}{T})a^* + \frac{1}{T}\mathcal{A}$: by convexity $\ell_t(a) \leq \ell_t(a^*) + \frac{1}{T}$ on $\mathcal{A}^*$, and $\nu(\mathcal{A}^*)/\nu(\mathcal{A}) = T^{-d}$. Hence

$$\log \frac{W_{T+1}}{W_1} \geq \log \frac{\nu(\mathcal{A}^*)}{\nu(\mathcal{A})} - \eta \sum_{t=1}^T (\ell_t(a^*) + \frac{1}{T}) = -d \log T - \eta \sum_{t=1}^T \ell_t(a^*) - \eta,$$

and combining with the upper bound as before gives the theorem: $\log(\text{number of actions})$ has become $d \log T$, the log-number of hypotheses at resolution $1/T$ in d dimensions. ■

Universal portfolio (Cover, 1991)

Take $\mathcal{A} = \Delta_d$, $\ell_t(p) = -\log\langle p, r_t \rangle$, $\eta = 1$: $dp_t(p) \propto \prod_{s < t} \langle p, r_s \rangle dp$, the wealth the rebalanced portfolio p would have earned so far.

Cover's algorithm

for $t = 1, 2, \dots, T$:

Play the posterior mean $\bar{p}_t = \frac{\int_{\Delta_d} p \prod_{s < t} \langle p, r_s \rangle dp}{\int_{\Delta_d} \prod_{s < t} \langle p, r_s \rangle dp}$

Observe r_t ; wealth multiplies by $\langle \bar{p}_t, r_t \rangle$

No sampling: $\mathbb{E}_{p \sim p_t} \langle p, r_t \rangle = \langle \bar{p}_t, r_t \rangle$ since the wealth factor is *linear* in p .

Corollary

For any return sequence,

$$\max_{p \in \Delta_d} \sum_t \log \langle p, r_t \rangle - \sum_t \log \langle \bar{p}_t, r_t \rangle \leq d \log T + 2.$$

Universal portfolio (Cover, 1991)

Take $\mathcal{A} = \Delta_d$, $\ell_t(p) = -\log\langle p, r_t \rangle$, $\eta = 1$: $dp_t(p) \propto \prod_{s < t} \langle p, r_s \rangle dp$, the wealth the rebalanced portfolio p would have earned so far.

Cover's algorithm

for $t = 1, 2, \dots, T$:

Play the posterior mean $\bar{p}_t = \frac{\int_{\Delta_d} p \prod_{s < t} \langle p, r_s \rangle dp}{\int_{\Delta_d} \prod_{s < t} \langle p, r_s \rangle dp}$

Observe r_t ; wealth multiplies by $\langle \bar{p}_t, r_t \rangle$

Corollary

For any return sequence,

$$\max_{p \in \Delta_d} \sum_t \log \langle p, r_t \rangle - \sum_t \log \langle \bar{p}_t, r_t \rangle \leq d \log T + 2.$$

No sampling: $\mathbb{E}_{p \sim p_t} \langle p, r_t \rangle = \langle \bar{p}_t, r_t \rangle$ since the wealth factor is *linear* in p .

Proof. $\langle \bar{p}_t, r_t \rangle = W_{t+1}/W_t$ exactly (linearity), so $\sum_t \log \langle \bar{p}_t, r_t \rangle = \log(W_{T+1}/W_1)$. On $\mathcal{A}^* = (1 - \frac{1}{T})p^* + \frac{1}{T}\Delta_d$ we have $\langle p, r_t \rangle \geq (1 - \frac{1}{T})\langle p^*, r_t \rangle$, so

$$\log \frac{W_{T+1}}{W_1} \geq -d \log T + \sum_{t=1}^T \log \langle p^*, r_t \rangle + T \log \left(1 - \frac{1}{T}\right) \geq \sum_{t=1}^T \log \langle p^*, r_t \rangle - d \log T - 2. \quad \blacksquare$$

Cost: an integral over Δ_d every round; exact only for tiny d , otherwise sampling.

Online convex optimization

Trading the 0/1 loss for tractability

Back to classification: we have no practical algorithm

- **Adversarial:** exponential weights keeps a weight for every $h \in \mathcal{H}$. For $\mathcal{H}$ = linear classifiers in $\mathbb{R}^d$, $|\mathcal{H}| = \infty$; even discretized it is $e^{\Theta(d)}$.
- **IID:** ERM for the 0/1 loss, $\arg \min_h \sum_s \mathbf{1}\{h(x_s) \neq y_s\}$, is NP-hard for hyperplanes.
- **Portfolios:** the comparator was a continuum Δ_d and the loss was not linear.

Back to classification: we have no practical algorithm

- **Adversarial:** exponential weights keeps a weight for every $h \in \mathcal{H}$. For $\mathcal{H} =$ linear classifiers in $\mathbb{R}^d$, $|\mathcal{H}| = \infty$; even discretized it is $e^{\Theta(d)}$.
- **IID:** ERM for the 0/1 loss, $\arg \min_h \sum_s \mathbf{1}\{h(x_s) \neq y_s\}$, is NP-hard for hyperplanes.
- **Portfolios:** the comparator was a continuum Δ_d and the loss was not linear.

Solution: relax the objective

Parameterize $h_w(x) = \text{sign}(\langle w, x \rangle)$ with $w \in \mathcal{K} \subseteq \mathbb{R}^d$ convex, and replace $\mathbf{1}\{h_w(x) \neq y\}$ by a **convex surrogate** in w , e.g. the logistic loss

$$\ell(w, (x, y)) = \log(1 + \exp(-y\langle w, x \rangle)) \quad \text{or the hinge loss } \max\{0, 1 - y\langle w, x \rangle\}.$$

Measure regret in the surrogate: $\text{Regret}_T = \sum_{t=1}^T \ell_t(w_t) - \min_{w \in \mathcal{K}} \sum_{t=1}^T \ell_t(w).$

Back to classification: we have no practical algorithm

- **Adversarial:** exponential weights keeps a weight for every $h \in \mathcal{H}$. For $\mathcal{H} =$ linear classifiers in $\mathbb{R}^d$, $|\mathcal{H}| = \infty$; even discretized it is $e^{\Theta(d)}$.
- **IID:** ERM for the 0/1 loss, $\arg \min_h \sum_s \mathbf{1}\{h(x_s) \neq y_s\}$, is NP-hard for hyperplanes.
- **Portfolios:** the comparator was a continuum Δ_d and the loss was not linear.

Solution: relax the objective

Parameterize $h_w(x) = \text{sign}(\langle w, x \rangle)$ with $w \in \mathcal{K} \subseteq \mathbb{R}^d$ convex, and replace $\mathbf{1}\{h_w(x) \neq y\}$ by a **convex surrogate** in w , e.g. the logistic loss

$$\ell(w, (x, y)) = \log(1 + \exp(-y\langle w, x \rangle)) \quad \text{or the hinge loss } \max\{0, 1 - y\langle w, x \rangle\}.$$

Measure regret in the surrogate: $\text{Regret}_T = \sum_{t=1}^T \ell_t(w_t) - \min_{w \in \mathcal{K}} \sum_{t=1}^T \ell_t(w).$

Both the portfolio problem and this one are instances of one protocol.

Online convex optimization (OCO)

Input: convex set $\mathcal{K} \subseteq \mathbb{R}^d$

for $t = 1, 2, \dots, T$:

Player picks $a_t \in \mathcal{K}$

Adversary reveals a convex $\ell_t : \mathcal{K} \rightarrow \mathbb{R}$

Player suffers $\ell_t(a_t)$

Special cases. Experts: $\mathcal{K} = \Delta_d$, loss

$p \mapsto \langle p, \ell_t \rangle$. Portfolios: $\mathcal{K} = \Delta_d$,

$\ell_t(p) = -\log \langle p, r_t \rangle$. Classification:

$\mathcal{K} = \{\|w\|_2 \leq B\}$, logistic loss.

Online convex optimization and online gradient descent

Online convex optimization (OCO)

Input: convex set $\mathcal{K} \subseteq \mathbb{R}^d$

for $t = 1, 2, \dots, T$:

Player picks $a_t \in \mathcal{K}$

Adversary reveals a convex $\ell_t : \mathcal{K} \rightarrow \mathbb{R}$

Player suffers $\ell_t(a_t)$

Special cases. Experts: $\mathcal{K} = \Delta_d$, loss $p \mapsto \langle p, \ell_t \rangle$. Portfolios: $\mathcal{K} = \Delta_d$, $\ell_t(p) = -\log \langle p, r_t \rangle$. Classification: $\mathcal{K} = \{\|w\|_2 \leq B\}$, logistic loss.

Online gradient descent (OGD)

Input: $\mathcal{K}$, step size $\eta > 0$; any $a_1 \in \mathcal{K}$

for $t = 1, 2, \dots, T$:

Play a_t ; observe ℓ_t ; let $g_t \in \partial \ell_t(a_t)$

$a_{t+1} = \Pi_{\mathcal{K}}(a_t - \eta g_t)$

where $\Pi_{\mathcal{K}}(v) = \arg \min_{a \in \mathcal{K}} \|a - v\|_2$

Why gradients suffice. Convexity gives a global linear lower bound from local information:

$$\ell_t(a) \geq \ell_t(a_t) + \langle g_t, a - a_t \rangle \quad \forall a \in \mathcal{K}.$$

One gradient and one projection per round, regardless of how “many” hypotheses there are.

Theorem (OGD)

Let $\mathcal{K}$ be convex with diameter $D = \sup_{a,b \in \mathcal{K}} \|a - b\|_2$, and suppose $\|g_t\|_2 \leq G$ for all t . For any $a^* \in \mathcal{K}$,

$$\sum_{t=1}^T \ell_t(a_t) - \ell_t(a^*) \leq \frac{D^2}{2\eta} + \frac{\eta G^2 T}{2} \stackrel{\eta = D/(G\sqrt{T})}{=} DG\sqrt{T}.$$

Theorem (OGD)

Let $\mathcal{K}$ be convex with diameter $D = \sup_{a,b \in \mathcal{K}} \|a - b\|_2$, and suppose $\|g_t\|_2 \leq G$ for all t . For any $a^* \in \mathcal{K}$,

$$\sum_{t=1}^T \ell_t(a_t) - \ell_t(a^*) \leq \frac{D^2}{2\eta} + \frac{\eta G^2 T}{2} \stackrel{\eta=D/(G\sqrt{T})}{=} DG\sqrt{T}.$$

Proof. Convexity: $\ell_t(a_t) - \ell_t(a^*) \leq \langle g_t, a_t - a^* \rangle$. Non-expansiveness of $\Pi_{\mathcal{K}}$:

$$\|a_{t+1} - a^*\|^2 \leq \|a_t - \eta g_t - a^*\|^2 = \|a_t - a^*\|^2 - 2\eta \langle g_t, a_t - a^* \rangle + \eta^2 \|g_t\|^2.$$

Rearrange, sum over t , and telescope the distances:

$$\sum_{t=1}^T \langle g_t, a_t - a^* \rangle \leq \frac{\|a_1 - a^*\|^2 - \|a_{T+1} - a^*\|^2}{2\eta} + \frac{\eta}{2} \sum_{t=1}^T \|g_t\|^2 \leq \frac{D^2}{2\eta} + \frac{\eta G^2 T}{2}. \quad \blacksquare$$

Online gradient descent: regret bound

Theorem (OGD)

Let $\mathcal{K}$ be convex with diameter $D = \sup_{a,b \in \mathcal{K}} \|a - b\|_2$, and suppose $\|g_t\|_2 \leq G$ for all t . For any $a^* \in \mathcal{K}$,

$$\sum_{t=1}^T \ell_t(a_t) - \ell_t(a^*) \leq \frac{D^2}{2\eta} + \frac{\eta G^2 T}{2} \stackrel{\eta=D/(G\sqrt{T})}{=} DG\sqrt{T}.$$

Proof. Convexity: $\ell_t(a_t) - \ell_t(a^*) \leq \langle g_t, a_t - a^* \rangle$. Non-expansiveness of $\Pi_{\mathcal{K}}$:

$$\|a_{t+1} - a^*\|^2 \leq \|a_t - \eta g_t - a^*\|^2 = \|a_t - a^*\|^2 - 2\eta \langle g_t, a_t - a^* \rangle + \eta^2 \|g_t\|^2.$$

Rearrange, sum over t , and telescope the distances:

$$\sum_{t=1}^T \langle g_t, a_t - a^* \rangle \leq \frac{\|a_1 - a^*\|^2 - \|a_{T+1} - a^*\|^2}{2\eta} + \frac{\eta}{2} \sum_{t=1}^T \|g_t\|^2 \leq \frac{D^2}{2\eta} + \frac{\eta G^2 T}{2}. \quad \blacksquare$$

Example (online logistic regression). $\|x_t\|_2 \leq 1$, $\mathcal{K} = \{w : \|w\|_2 \leq B\}$: $G \leq 1$, $D = 2B$, so regret $\leq 2B\sqrt{T}$ against an adversary, with $O(d)$ work per round.

Portfolios with OGD: stay away from the boundary

$\ell_t(p) = -\log\langle p, r_t \rangle$ is convex on Δ_d with $\nabla\ell_t(p) = -r_t/\langle p, r_t \rangle$, which blows up as $\langle p, r_t \rangle \rightarrow 0$. OGD needs $\|\nabla\ell_t\|_2 \leq G$, so restrict to

$$\mathcal{K}_\alpha = \{p \in \Delta_d : p_i \geq \alpha/d \ \forall i\}, \quad \text{on which} \quad \langle p, r_t \rangle \geq \frac{\alpha}{d} \sum_i r_t(i) \geq \frac{\alpha}{d} \|r_t\|_\infty,$$

$$\text{so} \quad \|\nabla\ell_t(p)\|_2 = \frac{\|r_t\|_2}{\langle p, r_t \rangle} \leq \frac{\sqrt{d} \|r_t\|_\infty}{(\alpha/d) \|r_t\|_\infty} = \frac{d^{3/2}}{\alpha}.$$

Portfolios with OGD: stay away from the boundary

$\ell_t(p) = -\log\langle p, r_t \rangle$ is convex on Δ_d with $\nabla\ell_t(p) = -r_t/\langle p, r_t \rangle$, which blows up as $\langle p, r_t \rangle \rightarrow 0$. OGD needs $\|\nabla\ell_t\|_2 \leq G$, so restrict to

$$\mathcal{K}_\alpha = \{p \in \Delta_d : p_i \geq \alpha/d \ \forall i\}, \quad \text{on which} \quad \langle p, r_t \rangle \geq \frac{\alpha}{d} \sum_i r_t(i) \geq \frac{\alpha}{d} \|r_t\|_\infty,$$

$$\text{so} \quad \|\nabla\ell_t(p)\|_2 = \frac{\|r_t\|_2}{\langle p, r_t \rangle} \leq \frac{\sqrt{d} \|r_t\|_\infty}{(\alpha/d) \|r_t\|_\infty} = \frac{d^{3/2}}{\alpha}.$$

Regret on $\mathcal{K}_\alpha$ ($D \leq \sqrt{2}$): OGD gives $\leq \sqrt{2} d^{3/2} \sqrt{T}/\alpha$.

Cost of shrinking the comparator: for $p^* \in \Delta_d$, $\tilde{p} = (1 - \alpha)p^* + \alpha \mathbf{1}/d \in \mathcal{K}_\alpha$ and $\log\langle \tilde{p}, r_t \rangle \geq \log\langle p^*, r_t \rangle + \log(1 - \alpha) \geq \log\langle p^*, r_t \rangle - 2\alpha$ for $\alpha \leq \frac{1}{2}$.

Total: $\sqrt{2} d^{3/2} \sqrt{T}/\alpha + 2\alpha T$. With $\alpha = d^{3/4} T^{-1/4}$, regret = $O(d^{3/4} T^{3/4})$.

	Regret	Per round
Universal portfolio	$d \log T$	integral on Δ_d
OGD on $\mathcal{K}_\alpha$	$d^{3/4} T^{3/4}$	$O(d)$

Tractability at a steep price in regret, because only convexity was used. Exploiting exp-concavity (Online Newton Step, Hazan–Agarwal–Kale 2007) gives $O(d \log T)$ in $\text{poly}(d)$ time per round.

Two lectures in one table

Setting	Assumption	Algorithm	Guarantee
Batch, realizable	IID	ERM	$R(\hat{h}) \leq \log(\mathcal{H} /\delta)/n$
Batch, agnostic	IID	ERM	$R(\hat{h}) - R(h_*) \leq \sqrt{2 \log(2 \mathcal{H} /\delta)/n}$
Online, realizable	IID	ERM each round	$O(\log \mathcal{H} \log T)$ mistakes
Online, realizable	none	Halving	$\log_2 \mathcal{H} $ mistakes
Online, agnostic	IID	ERM each round	$O(\sqrt{T \log \mathcal{H} })$ regret
Online, agnostic	none	Exponential weights	$\sqrt{T \log(\mathcal{H})/2}$ regret
Batch, squared loss	IID	ERM	$L(\hat{f}) - L(f_{\mathcal{F}}^*) \leq 4\sqrt{2L(f_{\mathcal{F}}^*)\beta/n} + 8\beta/n$
Online, squared loss	none	Averaged exp. weights	$2 \ln \mathcal{F} $ regret
Portfolio vs. best stock	none	Exp. weights ($\eta=1$)	$\log d$ log-wealth regret
Portfolio vs. best CRP	none	CEW ($\eta=1$, Cover)	$d \log T + 2$
Portfolio vs. best CRP	none	OGD on $\mathcal{K}_\alpha$	$O(d^{3/4}T^{3/4})$
Online convex	none	OGD	$DG\sqrt{T}$

Two lectures in one table

Setting	Assumption	Algorithm	Guarantee
Batch, realizable	IID	ERM	$R(\hat{h}) \leq \log(\mathcal{H} /\delta)/n$
Batch, agnostic	IID	ERM	$R(\hat{h}) - R(h_*) \leq \sqrt{2 \log(2 \mathcal{H} /\delta)/n}$
Online, realizable	IID	ERM each round	$O(\log \mathcal{H} \log T)$ mistakes
Online, realizable	none	Halving	$\log_2 \mathcal{H} $ mistakes
Online, agnostic	IID	ERM each round	$O(\sqrt{T \log \mathcal{H} })$ regret
Online, agnostic	none	Exponential weights	$\sqrt{T \log(\mathcal{H})/2}$ regret
Batch, squared loss	IID	ERM	$L(\hat{f}) - L(f_{\mathcal{F}}^*) \leq 4\sqrt{2L(f_{\mathcal{F}}^*)\beta/n} + 8\beta/n$
Online, squared loss	none	Averaged exp. weights	$2 \ln \mathcal{F} $ regret
Portfolio vs. best stock	none	Exp. weights ($\eta=1$)	$\log d$ log-wealth regret
Portfolio vs. best CRP	none	CEW ($\eta=1$, Cover)	$d \log T + 2$
Portfolio vs. best CRP	none	OGD on $\mathcal{K}_\alpha$	$O(d^{3/4}T^{3/4})$
Online convex	none	OGD	$DG\sqrt{T}$

The message

Dropping IID cost nothing in the rates. What you pay for is the **comparator** and **computation**.