HDPA Summer School
July 1st - 5th, 2019
ENS Paris

The five miracles of mirror descent

Sébastien Bubeck.

⤳ Robustness, potential-based, tracking
information geometry, adaptability.

Scribed by Claire Boyer

Chapter 1 The original miracles

I. Brief reminder of convexity

Let f be a cvx fct, $\forall x, y$

$$f(y) \geq f(x) + \nabla f(x) \cdot (y - x)$$

The difference between the left h.s & right h.s $=: D_f(y, x)$

Bregman divergence $\approx \frac{1}{2} \nabla^2 \Phi(x) [y - x ; y - x]$

Lemma 1: $y = \underset{x \in K}{\operatorname{argmin}} f(x) \iff -\nabla f(y) \in N_K(y)$

with $N_K(y) := \{ \theta \in \mathbb{R}^n, \ \forall x \in K \quad \theta \cdot (x - y) \leq 0 \}$

Proof: Say $-\nabla f(y) \notin N_K(y)$

$\iff$ $\exists$ a direction inside K, positively correlated with $-\nabla f(y)$

$\iff$ y is not a local (hence global) min.

Lemma 2: $K = \{x : \forall j \in [m] \quad a_j \cdot x \le b_j\}$

then $N_K(x) = \left\{ \sum_{j:\, a_j\cdot x = b_j} \lambda_j a_j \,,\ \lambda_j \ge 0 \ \ \forall j \in [m] \right\}$.

proof: $\boxed{\supseteq}$ $\displaystyle\sum_{j:\, a_j\cdot x = b_j} \lambda_j a_j \cdot (y - x) = \sum_{j:} \underbrace{\lambda_j}_{\ge 0} \underbrace{(a_j \cdot x - b_j)}_{\le 0} \le 0$.

$\boxed{\subseteq}$ Let $\theta \notin \{\sum \lambda_j a_j\}$

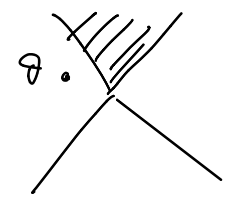

$\exists w$ such that $\begin{cases} w\cdot\theta = 1 \\ w \cdot \sum_j \lambda_j a_j < 1 \end{cases}$

positively correlated with θ

in fact < 1 must be ≤ 0.

so w is negatively correlated with any tight a_j.

$\Rightarrow x + \varepsilon w \in K$

εw is a direction inside K and pos. correlated with θ

$\Rightarrow \theta \notin N_K(x)$. $\square$

Lemma 3: $y = \mathrm{Proj}_K(x) = \arg\min_{z \in K} \|z - x\|_2$

then y is closer than x to <u>any</u> point in K.

proof: Apply Lemma 1 to $f(z) = \frac{1}{2}\|z - x\|_2^2$

$\nabla f(z) = z - x$

$\Rightarrow (y - x)\cdot(z - y) \ge 0 \quad \forall z \in K$

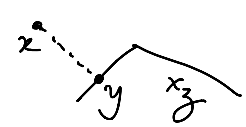

$$\|x-z\|^2 = \|x-y\|^2 + \|y-z\|^2 + 2\,(x-y)\cdot(y-z) \geq \|x-y\|^2 \qquad \square$$

II. Gradient descent

$$\boxed{x_{t+1} = x_t - \eta\, \nabla f(x_t)}$$

$\operatorname{argmin} f(x) =: x^*$ ×

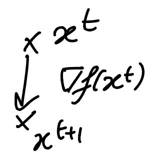

$$f(x_t) - f(x^*) \leq \nabla f(x_t)\cdot(x_t - x^*)$$

$-\nabla f(x_t)$ is positively correlated with $x^* - x_t$.

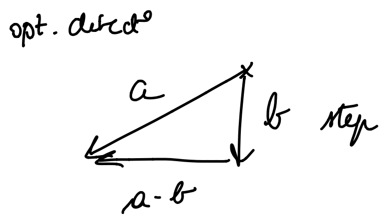

$$\|a\|^2 - \|a-b\|^2 = \underbrace{2\,a\cdot b}_{\text{pos. correlated}} - \underbrace{\|b\|^2}_{\text{"error term"}}$$

related to the gradient flow $\dfrac{d}{dt}x(t) = -\nabla f(x(t))$.

Gradient Descent Analysis

$$\|x^*-x_t\|^2 - \|x^*-x_{t+1}\|^2 = 2\eta\,(x^*-x_t)\cdot\nabla f(x_t) - \eta^2\|\nabla f(x_t)\|^2$$

$$\geq 2\eta\,(f(x_t)-f(x^*)) - \eta^2 L^2$$

f cvx and ∇f is L-Lipschitz

$$\sum_{t=1}^{T}\left(f(x_t)-f(x^*)\right) \leq \frac{\|x^*-x_1\|^2}{2\eta} + \frac{\eta}{2}TL^2$$

// no dependency on the dim. but a gradient has n bits of info.

Theorem: With optimal η $\exists t \in [T]$ such that

$$\underbrace{f(x_t) - f(x^*)}_{\text{smaller than } \frac{1}{T}\sum_{t=1}^{T}} \leq \frac{\|x^* - x_1\| L}{\sqrt{T}}$$

III – The first mini-miracle: constraints do not matter

$$x_{t+1} = P_K\left(x_t - \eta \nabla f(x_t)\right)$$

By Lemma 3, the same analysis holds: projection can only make things better.

IV – First miracle: robustness

Say we do a step with $g_t \in \mathbb{R}^n$ instead of $\nabla f(x_t)$

We proved

$$\sum_{t=1}^{T} g_t \cdot (x_t - x^*) = \underbrace{\overbrace{\frac{\|x^* - x_1\|^2 - \|x^* - x_{T+1}\|^2}{2\eta}}^{\text{range}}}_{\text{forgetting the initial condition}} + \overbrace{\frac{\eta}{2} \sum_{t=1}^{T} \|g_t\|_2^2}^{\text{variance term}}$$

Consider Stochastic Gradient Descent (SGD)

g_t is such that $\mathbb{E}[g_t \mid \text{past}] = \nabla f(x_t)$.

Theorem: $\mathbb{E}\left[f(x_t) - f(x^*)\right] \leq \frac{\|x^* - x_1\|_2 B}{\sqrt{T}}$

where $B := \max_t \mathbb{E}\|g_t\|_2^2$

[Robbins_Monro '51]

[Zinkevich 2003] Robustness to adversarial / biased gradients.

Say ε-fraction of the steps are corrupted (g_t is some arbitrary bounded vector)

$$\mathbb{E}\left[f(x_t) - f(x^*)\right] \leq \frac{\|x^* - x_1\|_2 B}{\sqrt{T}} + \varepsilon.$$

Regret against $f_1, \dots, f_T$ sequence of convex fcts

$$\frac{1}{T} \sum_{t=1}^{T} \left(f_t(x_t) - f_t(x^*)\right) \leq \frac{1}{\sqrt{T}}$$

V A first non-Euclidean setting: prediction with expert advice.

At each time step, the player algo picks an action $I_t \in [n]$ simultaneously an adversary picks $\ell_t \in \{0,1\}^n$.

Regret $$R_T(i) = \sum_{t=1}^{T} \left(\ell_t(I_t) - \ell_t(i) \right)$$

<u>Approach 1</u>: use Gradient Descent on $f_t(p) = \ell_t \cdot p$

with $K = \Delta_n = \{ p \in \mathbb{R}_+^n : \sum p_i = 1 \}$

$\nabla f_t(p) = \ell_t$ can be lip $\simeq \sqrt{n}$.

We get for any $q \in \Delta_n$ $$\sum_{t=1}^{T} \ell_t \cdot p_t - \ell_t \cdot q \leq \frac{1}{2\eta} + \frac{\eta}{2} Tn \leq \sqrt{Tn}.$$

this is suboptimal!

if $I_t \sim p_t$, $\mathbb{E} \sum \ell_t(I_t) = \mathbb{E} \sum \ell_t \cdot p_t$

OPT q is sparse: optimal rate is $\sqrt{T \log(n)}$

GD does not use the fact ~~that~~ OPT is sparse.

$\hookrightarrow$ we will modify GD so that it moves faster when far from a sparse point.

Multiplicative Weight U algorithm [Littlestone – Warmuth 89]

Fix $w_1 = (1, \dots, 1)$ at time t

for any i s.t. $\ell_t(i) = 1$

$$w_{t+1}(i) = (1 - \eta) w_t(i)$$ (reduce by η all the guys that are making a mistake)

otherwise $w_{t+1}(i) = w_t(i)$.

Play from $p_t = \frac{w_t}{\|w_t\|_1}$.

$$L_T = \sum_{t=1}^{T} l_t \cdot p_t$$

$$L^* = \min_{q \in \Delta_n} \sum_{t=1}^{T} l_t \cdot q$$

Theorem: $L_T \leq (1+\eta) L^* + \frac{\log(n)}{\eta}$

$\Rightarrow L_T - L^* \leq 2\sqrt{L^* \log(n)} \leq 2\sqrt{T \log(n)}$.

proof: track $\|w_{t+1}\|_1 = \|w_t\|_1 - \eta \underbrace{\sum_{i : l_t(i)=1} w_t(i)}_{l_t \cdot w_t}$

$w_{t+1}(i) \leq$

L^* "

for $t=T$ and $i=i^*$

$$= \|w_t\|_1 (1 - \eta\, l_t \cdot p_t)$$

$$\leq \|w_t\|_1 e^{-\eta l_t \cdot p_t} \leq n e^{-\eta L_t}$$

starting point.

$L^* < n e^{-\eta L_T}$

$L^* (\log (T \cdot \eta)$

Chapter 2 Mirror Descent and its second miracle –

Endow K with a Riemanian structure : $\langle \cdot, \cdot \rangle_x$ for any $x \in K$

before $\nabla f(x)$; $f(x+dx) \approx f(x) + \nabla f(x) \cdot dx$.

$\text{grad}_\eta f(x)$: $f(x+dx) \approx f(x) + \langle \text{grad}_\eta f(x), dx \rangle_x$

We will consider M such that $\langle \cdot, \cdot \rangle_x = \nabla^2 \phi(x) [\cdot, \cdot]$

Now $\text{grad}_\eta f(x) = \left[\nabla^2 \phi(x)\right]^{-1} \nabla f(x)$ for ϕ a fixed cvx fct.

$\rightsquigarrow$ preconditioning

__Idea__ : $x_{t+1} = x_t - \eta \nabla^2 \phi(x_t)^{-1} g_t$

__Question__ : What is the potential ?

__Miracle__ : $\partial \phi$ is a potential [Nemirovski '79]

I – Continuous time mirror descent

$g : \mathbb{R}_+ \to \mathbb{R}^n$ is a smooth vector field

$$\frac{d}{dt} x(t) = -\nabla^2 \phi(x(t))^{-1} \left(\eta g(t) + \partial(t)\right)$$

where $\partial(t) \in N_K(x(t))$, $x(t) \in K$

$$\iff x(t+dt) = \underset{x \in K}{\text{argmin}}\ \eta g(t) \cdot x + \frac{1}{2} \|x - x(t)\|^2_{x(t)}.$$

(Recall that $x_{t+1} = x_t - \eta g_t = \underset{x \in \mathbb{R}^n}{\text{argmin}}\ \eta g_t \cdot x + \frac{1}{2} \|x_t - x\|^2$)

Thm: $x(t)$ exists and is unique.

Analysis:

$$D\phi(y, x(t)) = \phi(y) - \phi(x(t)) - \nabla\phi(x(t)) \cdot (y - x(t))$$

$$\frac{d}{dt} D\phi = \cancel{-\nabla\phi(x(t))} \cdot \frac{d}{dt} x(t) - \nabla^2\phi(x(t)) \frac{d}{dt} x(t) \cdot (y - x(t)) \cancel{- \nabla\phi(x(t))} \cdot \left(-\frac{d}{dt} x(t)\right)$$

$$= (\eta\, g(t) + \lambda(t)) \cdot \underbrace{(y - x(t))}$$

$y \in K$ and $\lambda(t) \in N_K(x(t))$
$\Rightarrow \lambda(t) \cdot (y - x(t)) \leq 0$.

$$\leq \eta \underbrace{g(t) \cdot (y - x(t))}_{\text{how } y \text{ is subopt.}}$$

if you regret, you are learning and making the Bregman DV decrease.

$$\int_0^T g(t) \cdot (x(t) - y)\, dt \leq \frac{D\phi(y, x(0))}{\eta}$$

for any $y \in K$.

Here only the range term

Discrete time version:

discretize $\frac{d}{dt} x(t) = -\nabla^2\phi(x(t))^{-1} g(t)$

$$\Longleftrightarrow \nabla^2\phi(x(t)) \frac{d}{dt} x(t) = -g(t)$$

$$\underbrace{\qquad\qquad\qquad}_{\frac{d}{dt}\nabla\phi(x(t))}$$

$$\nabla\phi(x_{t+1}) = \underbrace{\nabla\phi}(x_t) - \eta\, g_t$$

map the primal point to a dual pt

↳ lives in the dual
it's a linear fct.

Chapter 3: Metrical task systems and the third miracle

MTS — Let X be a finite metric space

the player maintains a state $\rho_t \in X$

At time $t+1$, he receives a task fct $c_{t+1} : X \to \mathbb{R}_+$

and move to ρ_{t+1} (change the state to complete the task more easily)

$\hookrightarrow$ movement cost $= \mathrm{dist}(\rho_t, \rho_{t+1})$

$\hookrightarrow$ service cost $= c_{t+1}(\rho_{t+1})$

[Borodin, Linial, Saks '92]

Some applications

- Self-adjusting data structure

  ex: $X \equiv S_n$ set of permutations on n elts
  cost fct on the set of — " —

  3241.

- power-management, 3 states

	on	sleep	off
task 1	0	$+\infty$	$+\infty$
task 2	1	ε	0

(I want to use my comp.)

- strictly more general than prediction with expert advice -

Type of guarantees:

- competitive ratio $\equiv ALG/OPT$
- conjecture in online algo: for any metric space $|X|=n$ $\exists$ a $O(\log(n))$ competitive algorithm.

Today thm [B., Cohen, J. Lee, YT Lee '18]

they prove $O(\log^2(n))$

$\forall X$ $\Omega\left(\frac{\log(n)}{\log\log n}\right)$ is necessary.

I Warm-up: deterministic ~~strategy~~ on uniform metric

all dist $= 1$.

$O(n)$

Stay in a state until it accumulates a cost of 1.
Then mark the state and move to an unmarked state
Once everything is marked I paid $\leq 2n$
and OPT pays at least 1.
So I'm $2n$-competitive

$\Omega(n)$

$c_{t+1}(\rho) = +\infty \mathbb{1}_{\{\rho=\rho_t\}}$
over n time steps I pay n and opt pays less than 1
so $\Omega(n)$-competitive.

⇝ Randomized strategy for uniform metric

$O(n)$

Stay in a state until it accumulates a cost of 1.
Then mark the state and move to an random unmarked state

$f(n)$ = expected mvmt when n states are unmarked

$$f(n) = \frac{1}{n} + f(n-1) \Rightarrow f(n) = \log(n)$$

$\Omega(\log n)$

$C_{t+1}(p) = \infty$ at a random location
ALG cost over n times is 1 but it takes $n\log(n)$ steps to cover X.

Weighted-star case

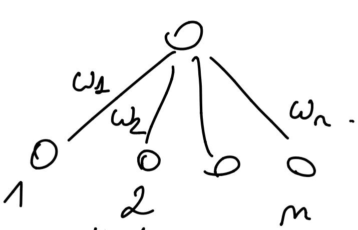

Open question: "weighted" coupon collector
For given weights $(w_i)_{1 \le i \le n}$ what is lower-bound?

II. Beyond uniform spaces: weighted stars

Continuous-time pb : instead of c_t, $t \in \mathbb{N} \to c(t)$ $t \in \mathbb{R}_+$

Lemma: Continuous-time guarantees transfer to discrete-time guarantees.

proof: $c_t \to c(t)$ is a water-filling

$c_t =$ 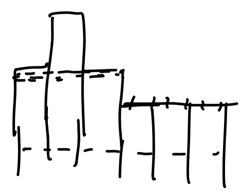

from $\rho(t)$ to ρ_t $\to$ smallest cost state visited by the CT alg.

$$\text{SC-DT} \leq \text{SC-CT} \quad \text{(because of the water-filling)}$$

$$\text{MC\_DT} \leq \text{MC\_CT}$$

$$\underbrace{\text{OPT\_CT}}_{\text{because more degrees of freedom}} \leq \text{OPT\_DT}$$

Continuous-space problem

instead of $\rho_t \in X$, we maintain $x_t \in \Delta_n$.

$$\operatorname{dist}(\rho_t, \rho_{t+1}) \rightsquigarrow W_1(x_t, x_{t+1}) := \inf_{\substack{(X,Y):\, X \sim x_t \\ Y \sim x_{t+1}}} \mathbb{E}\operatorname{dist}(X,Y)$$

Finally, assuming that $W_1(p,q) = \|p-q\|$

then $SC = \int c(t) \cdot x(t)\, dt$

$$MC = \int \left\| \frac{d}{dt} x(t) \right\| dt$$

The mirror-descent flow on $K = \Delta_n$

$$\begin{cases} \frac{d}{dt} x(t) = -\nabla^2 \phi(x(t))^{-1} \left(\eta c(t) + \lambda(t)\right) \\ \lambda(t) \in N_K(x(t)) \qquad x(t) \in K. \end{cases}$$

Recall Miracle 2: $\frac{d}{dt} D_\phi(y, x(t)) \leq \eta c(t) \cdot (y - x(t))$

Miracle 3: $\frac{d}{dt} D_\phi(y(t), x) \lessapprox$ movement of $y(t)$

$$= \|\dot{y}(t)\|$$

OPT moves to y is moving.

$\left(\phi(y) - \phi(x) - \nabla\phi(x) \cdot (y - x)\right)$ differentiate

$$\left(\nabla\phi(y(t)) - \nabla\phi(x)\right) \cdot \dot{y}(t) \leq 2L \|\dot{y}(t)\|$$

where $L = \sup_{z \in K} \|\nabla \phi(z)\|_*$.

We get $\frac{d}{dt} D_\phi\big(y(t), x(t)\big) \leq \eta\, c(t) \cdot \big(y(t) - x(t)\big) + 2L \|\dot{y}(t)\|$.
(when both are moving)

~~integrate~~

$$0 \underset{\substack{\uparrow \\ \text{if } x(0)=y(0)}}{\leq} D\phi\big(y(T), x(T)\big) - D\phi\big(y(0), x(0)\big) \leq \eta\big(SC\text{-}OPT - SC\text{-}ALG + 2L\, MC\_OPT\big)$$

Lemma: Take $\eta = 2L$ assume ~~that~~ $x(0) = y(0)$
then $SC\_ALG \leq OPT$

$$MC\_ALG = \|\dot{x}(t)\| = \Big\| \nabla^2\phi(x(t))^{-1} \big(\underbrace{\eta\, c(t)}_{\text{mvmt by the cost fct}} + \underbrace{\partial(t)}_{\text{mvmt by the constraint}}\big) \Big\|$$

What about MC_ALG?

↳ Weighted-star: metric space ≡ leaves of

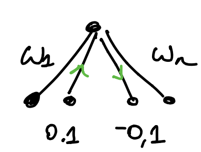

$\|h\| = \sum_{i=1}^{n} w_i |h_i|$ the norm is just a weighted ℓ^1-norm.

$\|g\|_* = \max_{i \in [n]} \frac{1}{w_i} |h_i|$

↳ Design of the mirror map:

Natural to first look at separable maps $\phi(x) = \sum_{i=1}^{n} \varphi_i(x_i)$

$\nabla^2 \phi(x)^{-1} = \mathrm{diag}\left(\frac{1}{\varphi_i''(x)}\right)$

the mvmt due to $c(t)$ the cost fct is

$$\|\nabla^2\Phi(x(t))^{-1}c(t)\| = \sum_{i=1}^{n} w_i \frac{c_i(t)}{\varphi_i''(x_i(t))}$$

We want $\underbrace{\sum w_i \frac{c_i}{\varphi_i''(x_i)}}_{\text{MC due to } c} \leq \underbrace{\sum_i c_i x_i}_{\text{SC}}$

then $MC \leq \eta\, SC \leq \eta\, OPT \implies \eta$ - competitivity.

$$\varphi_i''(s) = \frac{w_i}{s} \Leftarrow \varphi_i(s) = w_i s \log s$$

$$\Leftrightarrow \Phi(x) = \underbrace{\sum_{i=1}^{n} w_i\, x_i \log x_i}_{\text{pb: this is not Lipschitz}}.$$

pb: the dual norm of the gradient $\|\nabla\Phi(x)\|_*$ can be $+\infty$.

Fix goes back to Herbster & Warmuth 1998.

$$\Phi(x) = \sum w_i (x_i+\delta)\log(x_i+\delta)$$ the mirror map is the shifted entropy.

$$\implies \|\Phi(x)\|_* \leq \log(1/\delta)$$

$$\implies \log(1/\delta) \text{ - competitive}.$$

$$\sum w_i \frac{c_i}{\varphi_i''(x_i)} \leq \sum_i c_i (x_i+\delta)$$

$$\varphi_i''(s) = \frac{w_i}{s+\delta}$$

$\rightsquigarrow$ extra mvmt of $\delta\, c \cdot \mathbb{1}$

many reasons to expect that $\delta \simeq 1/n$ will be "fine"

i.e. $\displaystyle\int \delta\, c \cdot \mathbb{1} \lesssim \int c \cdot x$

(*)

~~proof of (*)~~ : consider $D\phi(\delta, x)$ the Bregman DV.

Assume that $K = \{x \geq 0,\ Ax \leq b\}$, $\delta \in K$ and $A\delta = b$.

$$\frac{d}{dt} D\phi(\delta, x(t)) = -\nabla^2\phi(x)\,\dot{x}(t) \cdot (\delta - x(t))$$

$$= (\eta\, c(t) + \lambda(t)) \cdot (\delta - x(t))$$

$$= \eta\, c(t) \cdot (\delta - x(t))$$

Now $\displaystyle\int \delta\, c \cdot \mathbb{1} \leq \int c \cdot x + \text{cst}$

this is true
assuming that the
constraint $x(t) \geq 0$ is inactive.

"reduced cost" if $x_i(t) = 0$ then $\hat{c}_i(t) = 0$
o.w. $\hat{c}_i(t) = c_i(t)$.

What is $\lambda(t)$? $K = \{x \in \mathbb{R}^n,\ \underbrace{x \geq 0}_{\xi \geq 0},\ \underbrace{\sum x_i = 1}_{\mu \in \mathbb{R}}\}$

$N_K(x) = \{-\xi - \mu \mathbb{1}\ \ \mu \in \mathbb{R}\ \ \xi \geq 0\ \ \xi_i > 0 \text{ when } x_i = 0\}$

?

$$\dot{x}_i(t) = -\frac{1}{\varphi_i''(x_i(t))}\left(\eta\, c_i(t) + \lambda_i(t)\right)$$

$$= -\frac{x_i+\delta}{w_i}\left(\eta\, c_i - \xi_i - \mu\right)$$

$$= -\frac{x_i+\delta}{w_i}\Big(\underbrace{\eta\, c_i - \xi_i}_{\text{reduced cost } \hat{c}_i} - \mu\Big)$$

μ: allows to keep a proba dist.

$\Rightarrow$ We see that necessarily $\mu \geq 0$

$\Rightarrow$ For the mvmt cost up to a factor 2 it is enough to look for $\|(\dot{x}(t))_-\|$ (the negative part)

Beyond stars : trees

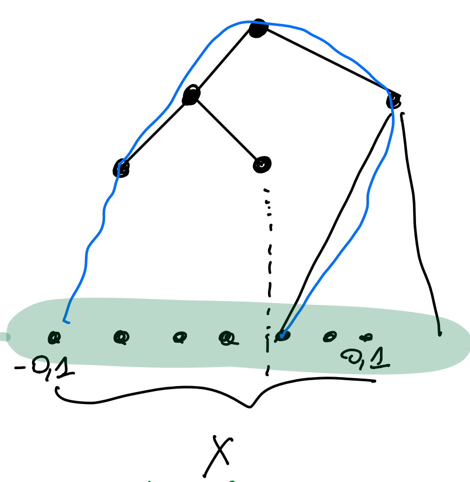

Simplex on the leaves

but inner var through the tree.

$h \in \mathbb{R}^L \to \hat{h} \in \mathbb{R}^V$

by $\hat{h}_u = \sum_{\ell \in u} h_\ell$

For any $u \in V$,

$$W_1(x, x+h) = \sum_u w_u |\hat{h}_u|$$

for trees

Now we work with $K = \left\{ x \in \mathbb{R}^V,\ x \geq 0,\ x_u = \sum_{v : l(v) = u} x_v \right\}$ with $x_{root} = 1$

and $\phi(x) = \sum_u w_u (x_u + \delta_u) \log(x_u + \delta_u)$

subtrees containing u as a leaf.

Thm: this algorithm is $\underbrace{\text{Depth}(G)}_{\text{from the Lagrangian analysis.}} \times \log n$-competitive

For the proof, some tricks:

1. instead of $x_e \geq 0$ only $x_e \geq 0$ on the leaves $l \in \mathcal{L}$.
2. in the def of K, enough to have $x_u \leq \sum_{v : l(v) = u} x_v$

Beyond trees: metric embeddings [Lecture 5]

$X \overset{D}{\hookrightarrow} Y$ X & Y 2 metric spaces

means that $\exists f : X \to Y$ s.t. $\forall (x,y) \in X$

$$1 \leq \frac{d_Y(f(x), f(y))}{d_X(x,y)} \leq D.$$

If this is the case and we have an α-competitive algo in Y

$$ALG_X \leq ALG_Y \leq \alpha\, OPT_Y \leq \alpha D \cdot OPT_X.$$

mvmt ~~that~~ I suffer in X.

(Pb) if $C_n \overset{D}{\hookrightarrow}$ Tree then $D = \Omega(n)$.
cycle of length n

⇝ Thm: (Rabinovich & Raz 98)

But Observation due to Karp [Alon, Karp, Peleg, West 90ish]
instead of considering a deterministic map φ, consider a random one.

randomized

$X \overset{D}{\hookrightarrow} Y$

a random map

means that $\exists f : X \to Y$ s.t. $\forall (x,y) \in X$

a.s.

only in $\mathbb{E}$

$$1 \leq \frac{d_Y(f(x), f(y))}{d_X(x,y)} \leq D.$$

Then:

$$\underbrace{\mathbb{E}[ALG_X]}_{} \leq \mathbb{E}[ALG_Y] \leq \mathbb{E}[\alpha \, OPT_Y] \leq \mathbb{E}\left[\alpha D \, \underbrace{OPT_X}_{}\right] = \alpha D \, OPT_X$$

this qtity depends on f!

OPT does not depend on f.

<u>As for the cycle?</u> If we cut C_n at a random edge

$$d_T(i,i+1) = \begin{cases} 1 & \text{w.p. } 1-1/n \\ n & \text{w.p. } \frac{1}{n} \end{cases}$$

then $\mathbb{E}\left[d_T(i,i+1) \right] = 2$. (Much better!)

<u>Thm:</u> [FRT 03, Bartal 96]

For any X, $\exists$ a $O(\log(n))$ ① randomized embedding into a tree of depth $\log(n)$ ②

For MTS, this gives $\log^3(n)$ ratio (①.② . the third one from yesterday)

But in fact these trees are special and instead of depth $\times \log(n)$) we can get $\log(n)$ on them.

First we prove a weaker version: $O(\log^2(n))$ distortion <u>and</u> only for metric spaces with $poly(n)$ aspect ratio.

proof:

Partition lemma: For any metric space

$\exists$ a random partition P_X such that all elements of the partition are small w/ diameter $\leq \Delta$ and such that $\forall x \in X$ and any $r \geq \dfrac{\Delta}{\mathrm{poly}(n)}$

$$\mathbb{P}\left(B(x,r) \text{ is cut by } P_X \right) \overset{(*)}{=} O\left(\frac{\log(n)\, r}{\Delta} \right)^{*}$$

small $\mathbb{P}$ that small balls are cut by the partition.

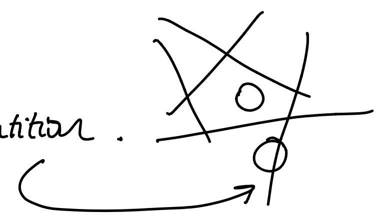

proof of the weaker thm using the lemma:

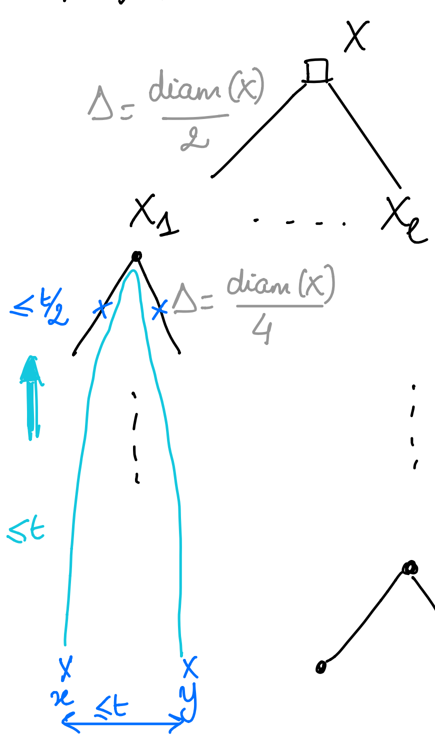

partition lemma on X with $\Delta = \mathrm{diam}(X)/2$.

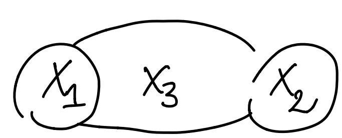

the depth is $O(\log n)$ because of the aspect ratio assumption.

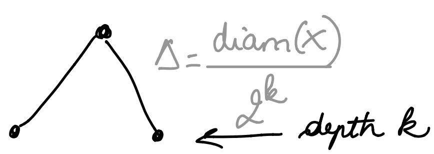

1) Non-contraction (a.s.) : $d_X(x,y) \leq d_T(x,y)$

[identity mapping]

Say that $d_T(x,y) \leq t$

$\Rightarrow$ x and y are in a partition element of diam $\leq t$

$\Rightarrow$ $d_X(x,y) \leq t$

for the metric in the original space.

2) Expected expansion : if x and y are cut at depth k and not $k-1$.

then

$$d_T(x,y) = 2 \sum_{\ell \geq k} \frac{\mathrm{diam}(X)}{2^\ell} = 4 \frac{\mathrm{diam}(X)}{2^k}$$

because of the 2 branches to go to x from y

this cancels with $\frac{1}{\Delta}$.

$$\Rightarrow \mathbb{E}\, d_T(x,y) \leq 4 \sum_k \underbrace{\mathbb{P}\big((x,y) \text{ is cut at depth } k\big)}_{} \cdot \frac{\mathrm{diam}(X)}{2^k}$$

the lemma gives $O\big(\log(n) \cdot d_X(x,y)\big) \cdot \frac{1}{\mathrm{diam}(X)/2^k}$

Coming from the sum.

$$\leq O\left(\log^2(n) \cdot d_X(x,y)\right)$$

$\blacksquare$

Proof of the partition lemma :

Assume we already have $X_1, \dots, X_{i-1} \subset X$

pick $x_i \notin X_1 \cup \dots \cup X_{i-1}$, pick a random radius $R_i \leq \Delta$

and set $X_i = B(x_i, R_i) \setminus \bigcup_{j<i} X_j$

Amazing trick : pick $R_i \sim Exp\left(\Theta\left(\frac{\log(n)}{\Delta}\right)\right)$

$\Rightarrow \mathbb{P}(R_i > \Delta) = \frac{1}{poly(n)}$

cdf $e^{-\lambda x} \mathbb{1}_{\{x \geq 0\}}$ of order Δ

On this event, just return all singletons.

Why $\circledast$ is true ?

$$\mathbb{P}(B(x,r) \text{ is cut by } X_i) \leq \mathbb{P}(Exp \leq r)$$
$$= 1 - e^{-\Theta\left(\frac{\log(n)}{\Delta}\right) r}$$

No need of a union bound to conclude, we just consider the X_i that matters for $B(x,r)$. ■

To get the strong version: Need an improved version of the lemma.

pick x_i at random.

$\cancel{r \geq \frac{\Delta}{poly(n)}}$

all equal / uniform in

$R = R_i \in \left[\frac{\Delta}{2}, \frac{\Delta}{4}\right]$

$$\mathbb{P}(B(x,r) \text{ is cut by } P_X)$$
$$\parallel$$
$$O\left(\frac{r}{\Delta} \log\left(\frac{|B(x,\Delta)|}{|B(x,\Delta/8)|}\right)\right)$$

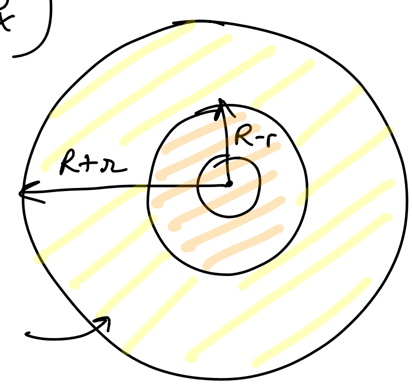

points that threaten the little ball.

$$\mathbb{P}(B(x,r) \text{ is not cut}) \geq \frac{|B(x,R-r)|}{|B(x,R+r)|}$$

if we pick a point there, it's okay since we contain the ball.

Finish with easy calculation, see James R. Lee blo. ■

Chapter 4 : The Fourth Miracle information geometry

$$\phi^*(\theta) = \max_{x \in \mathbb{R}^n} x \cdot \theta - \phi(x)$$

MD_ flow $\quad \dot{x}(t) = -\nabla^2\phi(x(t))^{-1}\left(\eta g(t) + \cancel{\lambda(t)}\right)$

$\lambda(t) \in N_K(x(t))$

$x(t) \in K$

For the moment, ignore the constraint

Miracle 2 : $\quad \frac{d}{dt} D_\phi(y, x(t)) = \eta g(t) \cdot (y - x(t))$

Miracle 3 : $\quad \frac{d}{dt} D_\phi(y(t), x) = L \|\dot{y}(t)\|$

Miracle 4 : $\quad \frac{d^2}{dt^2} D_\phi(y, x(t)) = -\eta g(t) \cdot \dot{x}(t)$

assume $g(t)$ is locally constant

$$= \eta g(t) \cdot \nabla^2\phi(x(t))^{-1}(\eta g(t))$$

Recall that Local norm $\|h\|_x^2 = h \cdot \nabla^2\phi(x) \cdot h$

dual norm $\|y\|_{x,*}^2 = y\left(\nabla^2\phi(x)\right)^{-1} y$.

$$= \eta^2 \|g(t)\|_{x,*}^2$$

With no constraint:

Continuous-time mirror descent: $\frac{d}{dt} \nabla\phi(x(t)) = -\eta g(t)$

Discrete-time mirror descent: $\nabla\phi(x_{t+1}) - \nabla\phi(x_t) = -\eta g_t$

With constraints:

Let's z_{t+1} be such that $\nabla\phi(z_{t+1}) = \nabla\phi(x_t) - \eta g_t$

$$x_{t+1} = \underset{x\in K}{\operatorname{argmin}} \; D\phi(x, z_{t+1})$$

Mirror Descent

What we expect for the regret $\sum_t g_t \cdot (x_t - y)$?

From miracle 2, this should be $\lesssim \frac{1}{\eta} \; D\phi(y, x_t) - D\phi(y, x_{t+1})$

From miracle 4, we expect $+\, \eta^2 \|g_t\|^2_{x_t, *}$.

Thus:

$$\boxed{\sum_{t=1}^{T} g_t \cdot (x_t - y) \leqslant \frac{D\phi(y, x(1))}{\eta} + \frac{1}{\eta} \sum_{t=1}^{T} D\phi^*\big(\overbrace{\nabla\phi(x_t) - \eta g_t}^{\nabla\phi(z_{t+1})}, \nabla\phi(x_t)\big)}$$

$$D\phi(x, y) \simeq \frac{1}{2}(x-y) \; \nabla^2\phi(y) \; (x-y)$$

$$D_{\phi^*}\big(\underbrace{\nabla\phi(y), \nabla\phi(x)}_{\text{dual points}}\big) = D\phi(\underbrace{x, y}_{\text{primal points}})$$

Denote $x^*(\theta)$ for the maximizer in ϕ^*

$$\nabla \phi^*(\theta) = x(\theta)$$

We also know at the optimum:

$$\underbrace{\theta - \nabla \phi(x(\theta))}_{\text{derivative at the optimum.}} = 0$$

$$\Rightarrow \nabla \phi(\nabla \phi^*(\theta)) = \theta$$

$$\Rightarrow \underbrace{\nabla^2 \phi^*(\theta)}_{} \nabla^2 \phi \left(\nabla \phi^*(\theta)\right) = I_n.$$

Hessian of the dual fct is the inverse of the hessian of the primal.

$\leadsto$ 2nd term of the thm OK

Proof: term by term:

$$\eta g_t \cdot (x_t - y) = \left(\nabla \phi(x_t) - \nabla \phi(z_{t+1})\right) \cdot (x_t - y)$$

I would like to make appear:

$$\text{Drop} = D\phi(y, x_t) - D\phi(y, x_{t+1}) = \phi(x_{t+1}) - \phi(x_t) - \nabla \phi(x_t) \cdot (y - x_t) + \nabla \phi(x_{t+1})(y - x_{t+1})$$

then $\eta g_t \cdot (x_t - y) = \nabla \phi(x_t) \cdot (x_t - y) - \nabla \phi(z_{t+1}) \cdot (x_t - x_{t+1}) - \nabla \phi(z_{t+1}) \cdot (x_{t+1} - y)$

the 1st order optimality condition of the projection x_{t+1} of z_{t+1}:

$$x_{t+1} = \underset{x \in K}{\arg\min}\, D\phi(x, z_{t+1})$$

then $\underbrace{\left[\nabla_x \left(D\phi(x_{t+1}, z_{t+1})\right)\right]}_{= \nabla \phi(x_{t+1}) - \nabla \phi(z_{t+1})} \cdot (y - x_{t+1}) \geq 0 \quad \forall y \in K.$

Therefore:

$$\eta\, g_t \cdot (x_t - y) = \nabla\phi(x_t)\cdot(x_t - y) - \nabla\phi(z_{t+1})\cdot(x_t - x_{t+1}) \underbrace{-\nabla\phi(z_{t+1})\cdot(x_{t+1} - y)}_{\leqslant \nabla\phi(x_{t+1})\cdot(x_{t+1}-y)}.$$

$$\leqslant \text{drop} + \phi(x_t) - \phi(x_{t+1}) - \nabla\phi(z_{t+1})\cdot(x_t - x_{t+1})$$

And,

$$D\phi^*(\nabla\phi(z_{t+1}), \nabla\phi(x_t)) = D\phi\,(x_t, z_{t+1})$$

$$= \phi(x_t) - \phi(x_{t+1}) - \nabla\phi(z_{t+1})\cdot(x_t - z_{t+1})$$

Then

$$\eta\, g_t \cdot (x_t - y) \leqslant \text{drop} + D\phi\,(x_t, z_{t+1}) + \underbrace{\phi(z_{t+1}) - \phi(x_{t+1}) - \nabla\phi(z_{t+1})\cdot(z_{t+1} - x_{t+1})}_{\text{by convexity } \leqslant 0}.$$

That's it! $\blacksquare$

Revisiting multiplicative weights update

$l_t \in [0,1]^n$, run Discrete-Time-Mirror Descent with $\phi(x) = \sum_{i=1}^{n} x_i \log x_i - x_i$.

$$\nabla\phi(x) = \log x_i$$

$$\nabla\phi(z_{t+1}) = \nabla\phi(x_t) - \eta\, l_t$$

$$\iff z_{t+1}(i) = x_t(i)\, e^{-\eta\, l_t(i)}.$$

To have a $\mathbb{P}$ dist, we project it using the Bregman DV

$$x_{t+1} = \underset{x \in K}{\operatorname{argmin}}\ \partial\phi(x, z_{t+1})$$

$$\Longleftrightarrow \nabla\phi(x_{t+1}) - \nabla\phi(z_{t+1}) = \lambda_t \in N_K(x_{t+1}).$$

$$= \mu_t \mathbb{1}$$

$$\Longleftrightarrow \log x_{t+1}(i) = \log z_{t+1}(i) + \log\mu_t$$

$\rightsquigarrow$ x_{t+1} is proportional to z_{t+1}.

$$\Longrightarrow x_{t+1} = \frac{z_{t+1}}{\|z_{t+1}\|_1}.$$

Looking at the result of the thm :

$$\sum_{t=1}^{T} g_t \cdot (x_t - y) \leqslant \underbrace{\frac{D_\phi(y, x(1))}{\eta}}_{\substack{\leqslant \Phi(y) - \Phi(x_1) \\ \text{if } x_1 = \underset{x \in K}{\operatorname{argmin}} \phi(x)}} + \frac{1}{\eta} \underbrace{\sum_{t=1}^{T} D_{\phi^*}(\overbrace{\nabla\phi(x_t) - \eta g_t}^{\nabla\phi(z_{t+1})}, \nabla\phi(x_t))}_{\frac{1}{2}\|\eta g_t\|^2_{\omega_t,*}}$$

for some ω_t such that $\nabla\phi(\omega_t) \in [\nabla\phi(x_t) - \eta g_t, \nabla\phi(x_t)]$

In particular if $\Phi(x) = \sum_{i=1}^{n} \varphi(x_i)$ (separable mirror map)

$$\|g_t\|^2_{\omega_t,*} = \sum_{i=1}^{n} \frac{g_t(i)^2}{\varphi''(\omega_t(i))}$$

$$\varphi'(\omega_t(i)) \in [\varphi'(x_t(i)) - \eta g_t(i), \varphi'(x_t(i))].$$

φ is convex so φ' is increasing.

so if $g_t(i) \geq 0$ then $\omega_t(i) \leqslant x_t(i)$

If φ'' is $\searrow$ so $\frac{1}{\varphi''} \nearrow$ with the $g_t(i) \geq 0$, then

$$\|g_t\|^2_{\omega_t,*} \leqslant \|g_t\|^2_{x_t,*}.$$

Corollary: If $g_t(i) \geq 0 \quad \forall t \quad \forall i$ and $\varphi'' \geq 0$

$\varphi''' \leqslant 0$

$$\boxed{\text{then } \sum_{t=1}^{T} g_t \cdot (x_t - y) \leqslant \frac{D\phi(y, x_1)}{\eta} + \frac{\eta}{2} \sum_{t=1}^{T} \|g_t\|^2_{x_t,*}.}$$

Application with $\varphi(s) = s\log(s)$: $\underbrace{\varphi''(s) = 1/s}_{\geq 0}$ $\underbrace{\varphi'''(s) = -1/s^2}_{\leqslant 0}$.

$\ell_t \in [0,1]^n$.

We get
$$\sum_{t=1}^{T} \ell_t \cdot (x_t - y) \leqslant \underbrace{\frac{-\sum x_1(i) \log x_1(i)}{\eta}}_{\leqslant \log(n)/\eta} + \underbrace{\frac{\eta}{2} \sum_{t=1}^{T} x_t(i) \ell_t(i)^2}_{(*)}$$

in particular $R_T \leqslant \frac{\log(n)}{\eta} + \frac{\eta}{2} T \leqslant \sqrt{\frac{T}{2} \log(n)}$

We get even better $L_T = \sum \ell_t \cdot x_t \qquad L^* = \sum \ell_t \cdot y$

$$L_T - L^* \leqslant \frac{\log(n)}{\eta} + \frac{\eta}{2} L_T$$

$$\iff L_T - L^* \leqslant \frac{1}{1-\eta/2} \left[\frac{\log(n)}{\eta} + \eta/2 \, L^* \right]$$

$$\Rightarrow \quad R_T \leqslant cst \sqrt{L^* \log(n)} + cst \log(n).$$

"ADAPTIVITY".

if an action has a small loss then we get a better regret

Multi - armed bandit

Prediction with Expert Advice (PwEA)
except that instead of observing l_t, you observe $l_t(I_t)$
[Robbins '51].

non-stochastic '90.

Q What are you doing MD on? What is g_t?

<u>Answer</u>: take $g_t(i) = \underbrace{\frac{l_t(i)}{x_t(i)} \mathbb{1}_{\{i=I_t\}}}_{\text{you got it only at the action that you play}}$ → reweighting

$$\mathbb{E}_{I_t \sim x_t}\left[g_t(i)\right] = l_t(i)$$

But the variance kills the game, and the weighted norm (*) saves us.

Q Regret analysis of multiplicative weights update with propensity score estimation.

The algorithm is called (Exp3) [Auer, Cesa Bianchi, Freund, Schapire 1996]

$$R_T = \left[\sum_{t=1}^{T} l_t \cdot (x_t - y)\right] = \mathbb{E}\left[\sum g_t \cdot (x_t - y)\right]$$

because of the unbiased estimator g_t.

$$\leqslant \frac{\log(n)}{\eta} + \frac{\eta}{2}\mathbb{E}\Big[\sum_{i,t} x_t(i)\underbrace{g_t(i)^2}_{\leqslant \frac{1}{x(i)^2}\mathbb{1}\{I_t=i\}}\Big]$$

because the loss $\leqslant 1$.

$$\leqslant \frac{\log(n)}{\eta} + \frac{\eta}{2} Tn$$

optimize in η.

$$\leqslant \sqrt{\frac{T}{2} n \log(n)}$$

Better mirror maps for bandit feedback

① L^* bound? We want $\frac{1}{\varphi''(s)} = s^2$ with $\varphi''' < 0$.

then variance term $= \mathbb{E}\sum_i x_t(i)^2 g_t(i)^2 \leqslant \mathbb{E}\Big[\sum_{i=1}^n x_t(i)^2 \ell_t(i)^2\Big]$

But what is the range?

$\varphi(s) = -\log(s)$ the mirror map is the log barrier.

pb: the range is going to be ∞.

So we get $\mathbb{E}\Big[\sum_{i=1}^T \ell_t \cdot (x_t - y)\Big] \leqslant \frac{\Phi(y)}{\eta} + \frac{\eta}{2} L_T$.

Idea: instead of $y = (1, 0, \dots 0)$, put $y = (1-n\delta, \delta, \dots, \delta)$

$\Longrightarrow R_T \leqslant \sqrt{L^* n \log(T)}$

we get a $\log(T)$ for every coordinate.

So far : entropy $\quad (\varphi'')^{-1} = x \quad$ range $= \log(n)$

	$(\varphi'')^{-1}$	Range /	Variance trade-off
Entropy	x	$\log(n)$	n
log-barrier	x^2	$n \log(T)$	$\ell_t \cdot x_t \leqslant 1 \Rightarrow R_T \leqslant \sqrt{n L^* \log(T)}$
<u>Idea</u> :	$x^{3/2}$?	$\sqrt{n}$	$\sqrt{n} \quad \Longrightarrow R_T \leqslant \sqrt{2nT}$.

Spoilers !

[Audibert, B. 2009]

[Audibert, Bubeck, Lugosi 2011] .

$$\varphi''(s) = \frac{1}{x^{3/2}} \qquad \varphi(s) = -\sqrt{s} \quad \rightsquigarrow \text{ Tsallis entropy}$$

$$\varphi'''(s) \leqslant 0 \qquad \qquad \leqslant 0$$

The thm gives :
$$R_T \leqslant \frac{-\Phi(x_1)}{\eta} + \frac{\eta}{2} \mathbb{E}\left[\sum_{t=1}^{T} x_{i,t}^{3/2} g_{i,t}^2\right]$$

$$\leqslant -\frac{\Phi(x_1)}{\eta} + \frac{\eta}{2} \sum_{t=1}^{T} \underbrace{\sum_{i=1}^{n} \sqrt{x_{i,t}}}_{\leqslant \sqrt{n} \text{ by Cauchy-Schwarz}}$$

$$-\Phi(x_1) \leqslant \sqrt{n} .$$

$$\leqslant \sqrt{2nT}$$

Linear Bandit

[Lecture 8]

Fix K a cvx body in $\mathbb{R}^n$.

At every time $t \geq 1$, the player selects $a_t \in K$

adv. selects $l_t \in \mathbb{R}^n$

the loss $:= l_t \cdot a_t$ Assume loss $\in [-1,1]$

- Bandit Feedback : only observe your loss -

- We want to apply Mirror Descent : (Q) What is g_t?
  ↳ gives $x_t \in K$ point suggested by MD.

We cannot choose a deterministic action because we cannot be predictable ⇝ we randomize.

Sampling scheme $p : K \to \mathcal{P}(K)$ with

$$Cov(p(x)) = \mathbb{E}_{a \sim p(x)} (a - \mu(x))(a - \mu(x))^T$$

where $\mu(x) = \mathbb{E}_{a \sim p(x)} [a]$

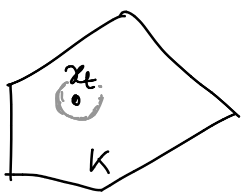

We restrict to p such that $\mu(x) = x$

I am going to play $a_t \sim p(x_t)$

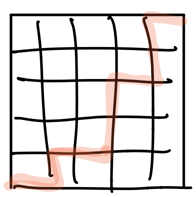

path $p \in \{0,1\}^m$

K = convex hull (p)

Answer: one-point linear regression, i.e.

$$g_t := \underbrace{Cov(x(t))^{-1} (a_t - x_t)}_{\in \mathbb{R}^n} \cdot \underbrace{a_t^T l_t}_{\in \mathbb{R}} \in \mathbb{R}^n$$

$$\mathbb{E}_{a_t \sim p(x_t)} g_t = l_t$$

Algorithm : $x_{t+1} = \underset{x \in K}{\operatorname{argmin}} \; D_\phi\left(x, \nabla\phi^*\left(\nabla\phi(x_t) - \eta g_t\right)\right)$

$a_{t+1} \sim p(x_{t+1})$

observe $l_{t+1} \cdot a_{t+1}$

build g_{t+1}

Regret $= \mathbb{E}\left[\sum_{t=1}^{T} g_t \cdot (x_t - y)\right] \leq \frac{\phi(y) - \phi(x_1)}{\eta} + \frac{\eta}{2} \mathbb{E}\left[\sum_{t=1}^{T} \|g_t\|^2_{x_t,*}\right]$ (proved yesterday without $\mathbb{E}$)

$||$

But Regret $= \mathbb{E}\left[\sum_{t=1}^{T} l_t \cdot (a_t - y)\right]$

$(a_t \leftarrow x_t$

$l_t \leftarrow g_t)$

Variance term :

$$\mathbb{E}\, \|g_t\|^2_{x_t,*} = \mathbb{E}\left[g_t^T \nabla^2\phi(x_t)^{-1} g_t\right]$$

$$= \mathbb{E}\Big[\underbrace{(a_t^T l_t)^2}_{\leq 1} \underbrace{(a_t - x_t)^T \Sigma_t^{-1} \nabla^2\phi(x_t)^{-1} \Sigma_t^{-1} (a_t - x_t)}_{Tr\left(\Sigma_t^{-1} \nabla^2\phi^{-1} \Sigma_t^{-1} (a_t - x_t)(a_t - x_t)^T\right)}\Big]$$

$$\overset{\circledast}{=} Tr\left(\Sigma_t^{-1} \nabla^2\phi(x_t)^{-1}\right)$$

since $a^T A a = Tr(A a a^T)$

Corollary: if $Tr\left(\Sigma_t^{-1} \nabla^2\phi(x_t)^{-1}\right) \leq V$ then

$R_T \leq \sqrt{V T \text{ range}}$ (up to w_t vs. x_t)

We would like $\nabla^2\Phi(x_t) \longrightarrow \infty$ as $x \to \partial K$.

Abernethy – Hazan – Rakhlin sampling scheme:

relation to interior-point method

The x_t vs w_t pb is about whether $\nabla^2\Phi(x_t) \simeq \nabla^2\Phi(w_t)$
(which is a well-conditioning pb)

Self-concordance: [Nesterov & Nemirovski 94]

$$\nabla^2\Phi(x) \simeq \nabla^2\Phi(x+h) \text{ for any } h \text{ such that } \|h\|_x < 1.$$

In particular, for a self-concordance barrier $\left(\Phi(x) \xrightarrow[x \to \partial K]{} +\infty\right)$

$$\underbrace{\{x+h : \|h\|_x < 1\}}_{\text{Dikin-ellipsoid}} \subset K$$

(because the Hessian is finite, so in the interior of the cvx body)

with Φ cvx K cvx.

Idea: $p(x) \equiv$ uniform on the boundary of Dikin ellipse at x

Take the ellipse $x^T \Sigma x \leq 1$ then $\Sigma^{-1/2} U$ with $U \sim \mathcal{U}(\mathbb{S}^{n-1})$

uniform on the boundary

$$x^T \Sigma x = U^T U \leq 1$$

$$x + \nabla^2\phi^{-1/2} U \sim p(x)$$

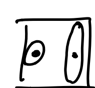

$$\Longrightarrow \operatorname{Cov}(x) = \mathbb{E}_{X \sim p(x)}\left[(X-x)(X-x)^T\right]$$

$$= \mathbb{E}\left[\nabla^2\Phi(x)^{-1/2} U U^T \nabla^2\Phi(x)^{-1/2}\right]$$

$$= \frac{1}{n} \nabla^2\Phi(x)^{-1}.$$

Conclusion: V for AHR sampling scheme is bounded by n^2

$\rightsquigarrow$ $n\sqrt{T\,\text{range}}$ bound.

Thm (Nesterov - Nemirovski) For all cvx body $\exists$ a self-concordant b with range $\leq n \log$ (dist of y to ∂K)

$\Rightarrow$ AHR $\leq n^{3/2}\sqrt{T\log T}$

After the break: $n\sqrt{T\log T}$

Deriving the bound $n\sqrt{T\log T}$ [Lecture 9]

Entropic barrier [Bubeck - Eldan '15]

P_θ is a distribution on K:

$$\frac{dP_\theta}{dx} \propto e^{\theta\cdot x}\,\mathbb{1}_{\{x\in K\}}$$

Claim: Duality between K and $\Theta = \mathbb{R}^n$

$\forall x \in \operatorname{int}(K)$ $\exists \theta$ such that

$$\operatorname{mean}(P_\theta) = x$$

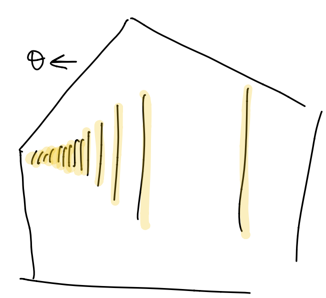

Proof: by the continuity of $\nabla\psi$.

Story: Let $\psi(\theta) = \log\left(\int e^{\theta\cdot x}\,dx\right)$

$$\nabla\psi(\theta) = \frac{\int x e^{\theta\cdot x}\,dx}{\int e^{\theta\cdot x}\,dx} = \operatorname{mean}(P_\theta).$$

$$\nabla^2 \Psi(\theta) = \frac{\int x x^T e^{\theta \cdot x}}{\int e^{\theta \cdot x} dx} - \frac{\left(\int x e^{\theta \cdot x} dx\right)^{\otimes 2}}{\left(\int e^{\theta \cdot x} dx\right)^2}$$

$= \mathrm{Cov}(P_\theta)$ which is psd. ($\Rightarrow \Psi$ cvx).

Def: Entropic barrier $\phi : K \to \mathbb{R}$ is $\phi \equiv \Psi^*$.

We have : $\nabla^2 \phi(x) = \nabla^2 \phi^* (\nabla \phi(x))^{-1} = \mathrm{Cov}\left(P_{\nabla \phi(x)}\right)^{-1}$

Take the sampling scheme $x \mapsto P_{\nabla \phi(x)}$ (sample from its "dual" param.)

↳ the variance is n.

↳ $\sqrt{nT \text{ range}}$.

Aside $\quad \theta = \nabla \phi(x)$

$\phi(x) = -H(P_\theta)$ with $H(p) = -\int p \log p$

proof $\quad \phi(x) = \Psi^*(x) = \theta \cdot x - \Psi(\theta)$ because θ is the dual param.

$$= \frac{\theta \cdot \int y e^{y \cdot \theta} dy}{\int e^{y \cdot \theta} dy} - \log\left(\int e^{\theta \cdot y} dy\right)$$

$$= \frac{\int e^{\theta \cdot y} \log\left(e^{\theta \cdot y} / Z\right) dy}{Z}$$

Chapter 5 : The Fifth Miracle : Adaptivity

Stochastic bandit : Multi-Armed Bandit w/ the assumption that (ℓ_t) are iid.

Let μ_i be the mean reward of action i.

A famous solution : Upper_Confidence Bound (UCB)
w/ in the reward model (so we want to maximize instead of minimize)

Let $T_{i,t}$ be the number of times that action i was played up to time t.

$$\hat{\mu}_{i,t} = \text{empirical mean}$$

$$\forall t \in [T] \quad \forall i \in [n] \quad \left| \hat{\mu}_{i,t} - \mu_i \right| \leq \sqrt{\frac{2\log(T)}{T_{i,t}}}$$

w.p. $\geq 1 - 1/T$.

Algorithm : play $I_t = \underset{i \in [n]}{\operatorname{argmax}} \hat{\mu}_{i,t} + \sqrt{\frac{2\log(T)}{T_{i,t}}}$ (optimistic)

Regret Analysis : $\Delta_i = \underbrace{\max_j \mu_j}_{\text{I wish to have played the max}} - \underbrace{\mu_i}_{\text{but I played } \mu_i}$

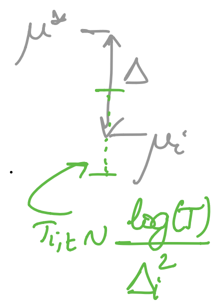

$$R_T = \sum_{i \neq i^*} \Delta_i \, \mathbb{E}\left[T_{i,t}\right]$$

$$\text{if } \hat{\mu}_{i,t} + \sqrt{\frac{2\log(T)}{T_{i,t}}} \leq \mu^* \leq \hat{\mu}_{i^*,t} + \sqrt{\frac{2\log(T)}{T_{i^*,t}}} \quad (*)$$

then i is never played anymore.

But $\hat{\mu}_{i,t} + \sqrt{\frac{2\log(T)}{T_{i,t}}} \leq \mu_i + 2\sqrt{\frac{2\log(T)}{T_{i,t}}}$

Then when $2\sqrt{\frac{2\log(T)}{T_{i,t}}} \leq \Delta_i \Rightarrow (*) \Rightarrow i$ never played anymore.

Thm: UCB has regret $\sum_{i \neq i^*} \frac{8\log(T)}{\Delta_i}$

Problem: Miracle 1 (robustness) does not apply at all to UCB ($\varepsilon = \frac{\log(T)}{T}$ corruptions sufficient to give a constant regret)

Question: Can you get best of both worlds?

Thm: [Bubeck - Slivkins '12]: Yes but with an ugly algorithm.

Thm: [Zimmert - Seldin '18] (++)

MD with Tsallis entropy ($\phi(x) = -\sum \sqrt{x_i}$) and time-dependent step size $\eta_t = \frac{1}{\sqrt{t}}$ achieves this

the magic is here and not in the step size.

Key-lemma: Any strategy such that the regret $R_T \leq c\, \mathbb{E}\left[\sum_{i \neq i^*, t} \sqrt{\frac{x_{i,t}}{t}}\right]$ satisfies best of both worlds properties

Proof: Adversarial bound $\sum_i \sqrt{x_{i,t}} \leq \sqrt{n \sum_i x_{i,t}}$

$$\Rightarrow R_T \leq c \sum_t \sqrt{\frac{n}{t}} \leq 2c\sqrt{nT}$$

Stochastic bound $\left(\sqrt{ab} \leq \frac{a+b}{2} \right)$

$$R_T \leq c\, \mathbb{E} \sum \sqrt{\frac{x_{i,t}}{t}} \leq \underbrace{\frac{1}{2} \sum \mathbb{E}\left[x_{i,t} \Delta_i\right]}_{\frac{1}{2} R_T} + \frac{1}{2} \sum \frac{c^2}{t \Delta_i}$$

$$\Rightarrow \frac{1}{2} R_T \leq \frac{1}{2} \sum_{i,t} \frac{c^2}{t \Delta_i} \leq \frac{1}{2} \sum_i \frac{c^2 \log(T)}{\Delta_i} .$$

Rk: UCB relies on concentration

MD does not care about concentration.

Open: replacing regular entropy by Tsallis entropy leads to adaptivity for multi-armed bandit.

What is the equivalent modification for the entropic barrier in linear bandit?