OMD:

"mirror map"

$$w_{t+1} = \nabla\psi^*_\Omega\left(\nabla\psi(w_t) - \eta \cdot g_t\right).$$

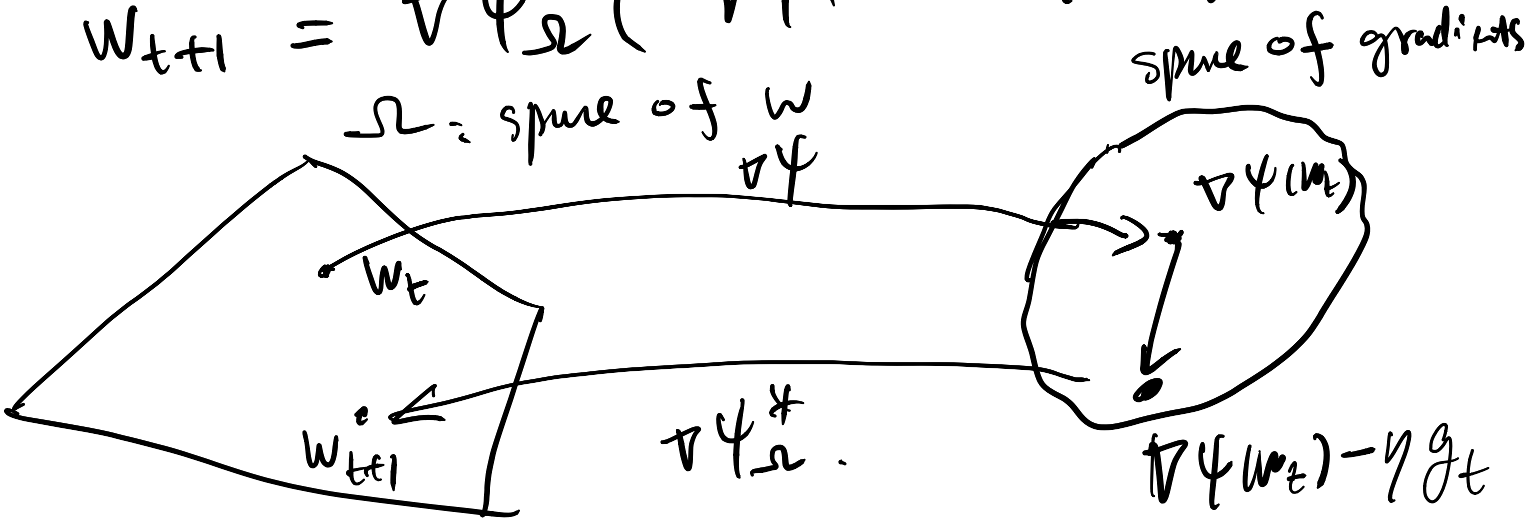

Guarantees of OMD:

Thm: If ψ is 1-SC wrt $\|\cdot\|$, then OMD w/ ψ and learning rate η has regret guarantee:

$$= \psi(u) - \psi(w_1) - \langle \nabla\psi(w_1), (u-w_1)\rangle$$

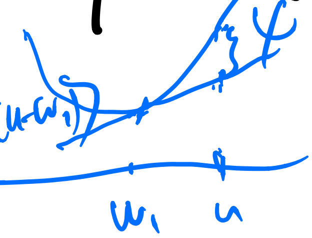

$$\forall u \in \Omega: \quad R_T(u) \le \frac{D_\psi(u, w_1)}{\eta} + \frac{\eta}{2}\sum_{t=1}^{T} \|g_t\|_*^2$$

Specifically, if $D_\psi(u, w_1) \le H^2$, and

$\forall t.\ \|g_t\|_* \le \rho.$ then.

$$\eta = \frac{H}{\rho}\sqrt{\frac{1}{T}} \Rightarrow R_T(u) \le H \cdot \rho \cdot \sqrt{T}.$$

Ex: ① p-norm algorithm

$$\psi(w) = \frac{1}{2(p-1)}\|w\|_p^2 \quad p \in (1,2].$$

$$\Omega = \mathbb{R}^d.$$

pick $\|\cdot\| = \|\cdot\|_p \Rightarrow \|\cdot\|_* = \|\cdot\|_q$

$$\frac{1}{p} + \frac{1}{q} = 1.$$

$$w_1 = \vec{0} \in \mathbb{R}^d.$$

OMD w/ ψ, step size $\eta \Rightarrow$

$$\forall u \in \mathbb{R}^d, \quad R_T(u) \le \frac{\overbrace{\|u\|_p^2}^{\le B_p^2}}{2(p-1)\eta} + \frac{\eta}{2}\underbrace{\sum_{t=1}^{T}\|g_t\|_q^2}_{\le R_q^2}$$

If f_t's are R_q-Lip ($\forall t\ \|g_t\|_q \le R_q$) wrt $\|\cdot\|_p$

$\forall u \in \{u \in \mathbb{R}^d : \|u\|_p \le B_p\}$,

$$R_T(u) \le \frac{B_p^2}{2(p-1)\eta} + \frac{\eta}{2} \cdot T R_q^2$$

$$\eta = \frac{B_p}{R_q} \sqrt{\frac{1}{T(p-1)}} \qquad B_p \cdot R_q \cdot \sqrt{\frac{T}{p-1}}.$$

($p = 2$ $\Rightarrow$ reduces to OGD guarantee)

(2). Exponential wt algorithm:

$$\psi(w) = \sum_i w_i \ln w_i \qquad \Omega = \Delta^{d-1}$$

$$\|\cdot\| = \|\cdot\|_1 \Rightarrow \|\cdot\|_* = \|\cdot\|_\infty$$

$$w_1 = \left(\frac{1}{d}, \ldots \frac{1}{d}\right).$$

$$D_\psi(u, w_1) = \sum_i u_i \ln \frac{u_i}{(1/d)}$$

$$= \sum_i u_i \ln u_i + \ln d. \le \ln d.$$

$u_i \in [0,1]$.

OMD w/ ψ, η.

$$\forall u \in \Delta^{d-1}, \quad R_T(u) \le \frac{\ln d}{\eta} + \frac{\eta}{2} \sum_{t=1}^{T} \|g_t\|_\infty^2$$

so if $\forall t$, $\|g_t\|_\infty \le R_\infty$, then

$$\ldots \quad \ldots \le \frac{\ln d}{\eta} + \frac{\eta}{2} \cdot T \cdot R_\infty^2$$

$$\left(\eta = \frac{1}{R_\infty} \cdot \sqrt{\frac{\ln d}{T}}\right) = R_\infty \cdot \sqrt{T \ln d}.$$

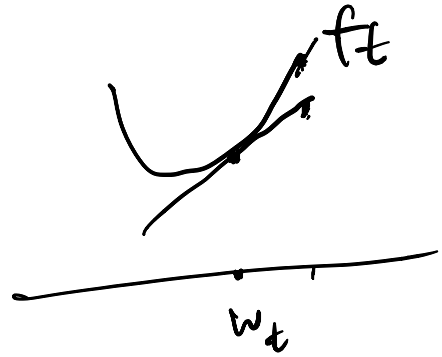

Pf of thm:

<u>step 1</u>: linearization:

$$R_T(u) = \sum_{t=1}^{T} f_t(w_t) - f_t(u)$$

$$\le \sum_{t=1}^{T} \langle \overbrace{g_t}^{\in \partial f_t(w_t)}, w_t - u \rangle$$

<u>step 2</u>: optimality condition on w_{t+1}.

$$w_{t+1} = \underset{w \in \Omega}{\arg\min} \underbrace{\langle \eta g_t, w \rangle + D_\psi(w, w_t)}_{f(w)}$$

$\psi(w) - \psi(w_t) - \langle \nabla\psi(w_t), w - w_t \rangle$

Ω

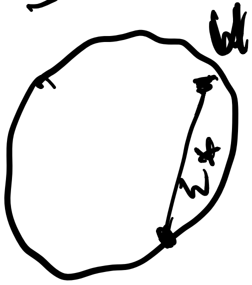

First order opt condition

$$\Rightarrow \langle \underbrace{\nabla f(w_{t+1})}_{\parallel}, u - w_{t+1} \rangle \geq 0$$

$$\Rightarrow \langle \nabla\psi(w_{t+1}) - \nabla\psi(w_t) + \eta g_t, u - w_{t+1} \rangle \geq 0.$$

$$\Rightarrow \langle g_t, w_{t+1} - u \rangle \leq \frac{1}{\eta} \left(\langle u - w_{t+1}, \nabla\psi(w_{t+1}) - \nabla\psi(w_t) \rangle \right)$$

Exercise

$$= \frac{1}{\eta} \left(D_\psi(u, w_t) - D_\psi(u, w_{t+1}) - \underbrace{D_\psi(w_{t+1}, w_t)}_{\geq \frac{1}{2}\|w_t - w_{t+1}\|^2} \right)$$

<u>step 3</u>: bounding $\langle g_t, w_t - u \rangle$:

$$\langle g_t, w_t - u\rangle = \langle g_t, w_{t+1} - u\rangle + \underline{\langle g_t, w_t - w_{t+1}\rangle}$$

$\le \|g_t\|_* \|w_t - w_{t+1}\|$

AM-GM. $\le \frac{\|w_t - w_{t+1}\|^2}{2\eta} + \frac{\eta}{2}\|g_t\|_*^2$

$$\le \frac{\eta}{2}\|g_t\|_*^2 + \frac{1}{\eta}\left(D_\psi(u, w_t) - D_\psi(u, w_{t+1})\right)$$

step 4 : summing over t :

$$R_T(u) \le \sum_{t=1}^{T} \langle g_t, w_t - u\rangle$$

w_{t+1}

$$\le \frac{\eta}{2}\sum_{t=1}^{T}\|g_t\|_*^2 + \frac{1}{\eta}\sum_{t=1}^{T}\left(D_\psi(u, w_t) - D_\psi(u, w_{t+1})\right)$$

$$\le \frac{\eta}{2}\sum_{t=1}^{T}\|g_t\|_*^2 + \frac{1}{\eta} \cdot D_\psi(u, w_1) .$$

Ex weather forecast problem.

$\underline{d}$ $\quad \Omega = \Delta^{d-1} \quad f_t(w) = \langle w, l_t\rangle$

$\forall t, \forall i \in [d]$ $\ell_t(i) \in [0,1]$.

want to design OMD algorithm that minimizes $R_T(\Omega)$.

How to choose ψ ?

① Exp wt alg:

$$R_T(\Omega) \le \frac{\ln d}{\eta} + \eta \underbrace{\sum_{t=1}^{T} \|\ell_t\|_\infty^2}_{\le T}$$

$$\le O\left(\sqrt{T \cdot \ln d}\right).$$

$\ell_t = (1 \cdots 1)$

② OGD:

$(1,0 \cdots 0)$ $(0,1 \cdots 0)$

$$R_T(\Omega) \le \frac{\overbrace{\max_{u,v \in \Omega} \|u-v\|_2^2}^{\le \text{const.}}}{2\eta} + \eta \underbrace{\sum_{t=1}^{T} \|\ell_t\|_2^2}_{\le T \cdot d}$$

$\forall u, v \in \Omega$, $\|u\|_1 = \|v\|_1 = 1$

$\Rightarrow \|u-v\|_1 \le 2$

$$\Rightarrow \|u-v\|_2 \le 2$$

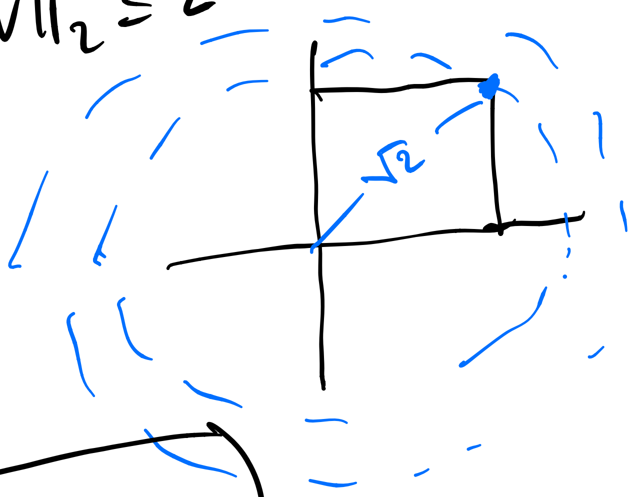

$$\le O\left(\sqrt{T\cdot d}\right)$$

③ p norm alg w/ $p = \dfrac{\ln d}{\ln d - 1} > 1 \Rightarrow q = \ln d$.

$$R_T(\Omega) \le \frac{\max_{u\in\Omega} \|u\|_p^2}{2\eta(p-1)} + \frac{\eta}{2}\sum_{t=1}^{T} \|l_t\|_q^2$$

$\|u\|_p^2 \le \|u\|_1^2 = 1$

$\dfrac{1}{p-1} = (\ln d - 1)$

$$\|l_t\|_q = \left(\sum_{i=1}^{d} l_t(i)^q\right)^{\frac{1}{q}}$$

$$\le d^{\frac{1}{q}} \le e.$$

Fact: $\forall\, p_1 < p_2 \quad \|u\|_{p_1} \ge \|u\|_{p_2}$.

$$\le \frac{1}{2\eta}(\ln d - 1) + \frac{\eta}{2} e^2 \cdot T.$$

$$\le O\left(\sqrt{(\ln d)\, T}\right).$$

Things we didn't cover for general OCO results:

① How to design algs when time horizon is unknown?

Idea 1: doubling trick:

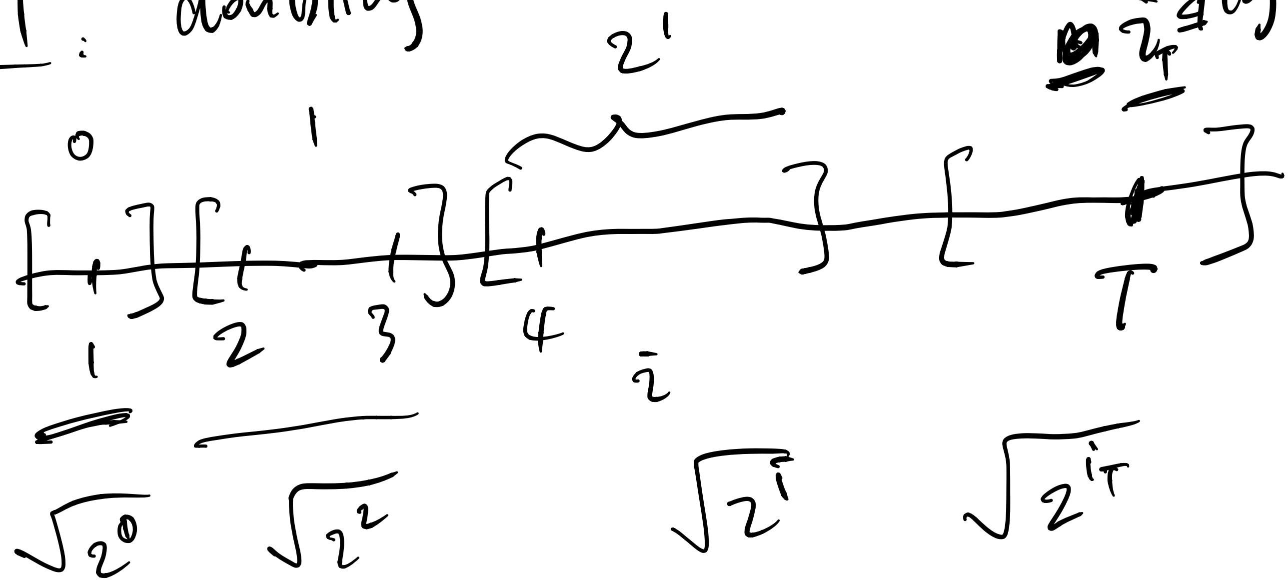

If A (that need to know T) has regret $\alpha\sqrt{T}$. then doubling trick gives an alg w/ regret $\le \frac{\sqrt{2}}{\sqrt{2}-1} \cdot \alpha \cdot \sqrt{T}$.

Idea 2: use time varying step size in

OMD: $\eta_t = \frac{H}{\rho}\sqrt{\frac{1}{t}}$

$O(H \cdot \rho \cdot \sqrt{T})$ regret.

(2) optimality of regret guarantees?

Δ^{d-1}

Thm: Let $\Omega \subseteq \mathbb{R}^d$ cvx. $D = \sup_{u,v \in \Omega} \|u-v\|_2$

T, d. large enough.

Let A be any alg. Let T be any integer. then there exists a sequence of linear fns $f_t(w) = \langle g_t, w \rangle$, $\|g_t\|_2 \le L$, $\|g_t\|_\infty \le L$.

such that $A \geq$ const $\cdot L \sqrt{T \ln d}$.

$$R_T(\Omega) \geq \frac{L \cdot D \sqrt{T}}{4}$$

$\Omega = \Delta^{d-1}$

Exp. wt

(OGD is optimal for this Ω, and assuming L-lip of the loss fns wrt $\|\cdot\|_2$ $\|\cdot\|_1$)

Caveat: does not rule out algs that can exploit "easy data" / "weak adversary".

③ Follow the regularized leader

For $t = 1, 2, \ldots T$

choose $w_t = \underset{w \in \Omega}{\arg\min} \sum_{s=1}^{t-1} f_s(w) + \frac{\psi(w)}{\eta}$

receive f_t, suffer loss $f_t(w_t)$.

Analysis: ~~intuition~~: w/o regularization (FTL), w_t's will oscillate a lot. ⇒ large regret.

w/ regularization, ⇒ alg stable, ⇒ amenable for analysis.

For some settings of ψ, Ω, FTRL may coincide w/ OMD.

when $f_t(w) = \langle w, g_t \rangle$

$$w_t = \underset{w \in \Omega}{\operatorname{argmin}} \frac{\psi(w)}{\eta} + \sum_{s=1}^{t-1} \langle g_s, w \rangle$$

$$= \nabla \psi_\Omega^*\left(-\eta \sum_{s=1}^{t-1} g_s\right)$$

<u>Next</u>: exploiting strong convexity in Online Convex Optimization (OCO).

$\underline{\underline{f_t(w)}}$ λ-SC.