- Online Learning

  Multiplicative Weights (MW)

- Using MW for Query Release

  "Learn" a synthetic distribution.

  "Synthetic Data"

Query Release

Adding Noise
(data independent)

- Laplace
- Gaussian
- Factorization
- Projection

Synthetic Data

Syn Data

$$\boxed{x} \xrightarrow{A_{\varepsilon,\delta}} \boxed{\hat{x}}$$

$$\hat{a} = F h_{\hat{x}}$$

$$\hat{a} \approx a = F h_x$$

Zero-Sum Game

Equilibrium Computation

Online Learning

Multiplicative Weights

MWEM

Online-PMW

Online Learning

"Sequential decision making"

Setting:
- a set of actions $\{1, \ldots, k\} = [k]$.
- "Game" between decison-maker D & Adversary A

" may know D's alg
not D's randomness"

For $t = 1, \ldots T$:

D chooses distribution $p^t \in \Delta(A)$

A chooses cost vector $c^t \in [0,1]^k$

$a^t \sim p^t$ sampled action

D pays cost $c^t_{a^t}$ and $\underbrace{\text{observes } c^t}_{\text{"Full information"}}$.

Focus on:

$$\mathbb{E}_{a \sim p^t}[c^t_a] = \langle p^t, c^t \rangle$$

Total cost: $\sum_{t=1}^{T} c^t_{a^t}$

$$\text{Regret} = \underbrace{\frac{1}{T} \sum_{t=1}^{T} c^t_{a^t}}_{\text{Cost of } D} - \underbrace{\min_{a \in [k]} \frac{1}{T} \sum_{t=1}^{T} c^t_a}_{\text{Cost of } a^*}$$

$\mathbb{E}[\text{Regret}]$

Online Learning

Give an algorithm Multiplicative Weights.

$$\mathbb{E}[\text{Regret}] \leq O\left(\sqrt{\frac{\ln(K)}{T}}\right).$$

logarithmic dependence on # actions.

"Bad" Algorithm: Follow-the-leader.

Select $a^t = \arg\min_{a \in [K]} C_a^{<t}$

$$C_a^{<t} = \sum_{\tau=1}^{t-1} C_a^{\tau}$$

Actions

t	1	2	FTL
C^1	1	0	?
C^2	0	1	2
C^3	1	0	1
	⋮		⋮

$\text{Reg} \approx \frac{1}{2}$.

Lack of Stability

Idea: At time t

$$p_a^t \propto (1-\eta)^{C_a^{<t}}$$

proportional to

Focus today

Also works

$$p_a^t \propto \exp(-\eta\, C_a^{<t})$$

Multiplicative Weights (MW)

$w_a^1 = 1$ for all $a \in [k]$

For $t=1$ to T:

$$Z_t = \sum_{a=1}^{k} w_a^t$$

$$\vec{p}^{\,t} = \frac{\vec{w}^{\,t}}{Z_t} \qquad \text{"probability vector"}$$

Observes $C^t \in [0,1]^k$

Update: for each a

$$w_a^{t+1} = w_a^t \cdot (1-\eta)^{C_a^t}$$

$$= \prod_{i=1}^{t} (1-\eta)^{C_a^i} = (1-\eta)^{C_a^{<t+1}}$$

Theorem. $\forall$ adversaries. $MW(\eta)$ has expected regret

$$\mathbb{E}[\text{Regret}] \le 2\sqrt{\frac{\ln(k)}{T}}, \quad \text{when } \eta = \sqrt{\frac{\ln(k)}{T}}.$$

Proof. Intuition. Two pieces

- Decision maker has high cost $\Rightarrow Z_t \downarrow$ (Z_t: Total weight)
- $Z_T \ge$ weight on a^* ($w_{a^*}^T$) (best in hindsight)

① Claim: $Z_{t+1} \le Z_t e^{-\eta \ell_t}$, $\ell_t = \langle c_t, p_t \rangle$.

$$Z_{t+1} = \sum_a w_a^{t+1} = \sum_a w_a^t (1-\eta)^{c_a^t}$$

$$\le \sum_a w_a^t (1-\eta c_a^t) \quad \leftarrow \boxed{\begin{array}{l}\forall \eta \in [0,\tfrac{1}{2}] \\ x \in [0,1] \\ (1-\eta)^x \le 1-\eta x\end{array}}$$

$$= \sum_a w_a^t - \sum_a w_a^t \cdot \eta \cdot c_a^t$$

$$= Z_t \Big(1 - \eta \underbrace{\sum_a p_a^t c_a^t}_{\langle p^t, c^t \rangle = \ell_t}\Big)$$

$$= Z_t(1-\eta \ell_t) \le Z_t e^{-\eta \ell_t} \quad \leftarrow \boxed{1-x \le e^{-x}}$$

$$Z_{T+1} \le \underbrace{Z_1}_{k} e^{-\eta \sum_t^T \ell_t}$$

② Claim: $Z_{t+1} \ge (1-\eta)^{OPT}$, $OPT = \min_a \sum_t c_a^t$

$$\ge e^{(-\eta-\eta^2)\,OPT}$$

$$Z_{t+1} \ge w_{a^*}^{T+1} = (1-\eta)^{c_{a^*}^{<T+1}}$$

$$= (1-\eta)^{OPT}$$

$$\ge e^{(-\eta-\eta^2)\,OPT} \quad \leftarrow \boxed{\begin{array}{l}1-\eta \ge e^{-\eta-\eta^2} \\ \text{for } 0 \le \eta \le \tfrac{1}{2}\end{array}}$$

$$e^{(-\eta-\eta^2)\,OPT} \leq Z_{t+1} \leq e^{-\eta \sum_t \ell^t} \cdot k.$$

Re-arranging:
$$\sum_t \ell^t - OPT \leq \eta \cdot \underbrace{OPT}_{\leq T} + \frac{\ln(k)}{\eta}.$$

Choose η to balance the two terms.

$\square$

Theorem. $\forall$ adversary $\forall\, p^* \in \Delta([m])$

$$\frac{1}{T}\sum_{t=1}^{T} \langle c^t, p^t \rangle - \frac{1}{T}\sum_{t=1}^{T} \langle c^t, p^* \rangle \leq 2\sqrt{\frac{\ln(m)}{T}}$$

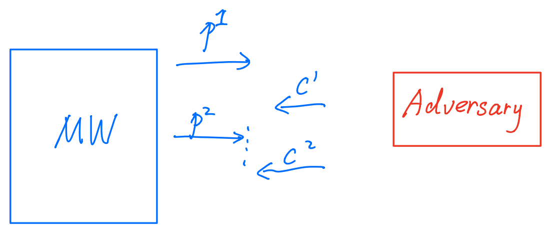

Query Release via Synthetic Data Distributions

→ Given $F = \{ f_1, \dots, f_k \}$, $\quad f_i(x) = \frac{1}{n} \sum_{j=1}^{n} \varphi_i(x_j)$

$$\varphi_i : \mathcal{X} \mapsto [0,1]$$

→ Histogram $(h_x)_u = \dfrac{\# \{ j \mid x_j = u \}}{n}$

→ Release answers $\hat{a} \approx F h_x$ ← Add noise previously

Idea: "Learn" a distribution $\hat{p}$ over $\mathcal{X} = \{1, \dots, m\}$

s.t. $\text{error}(\hat{p}) = \| F \hat{p} - F h_x \|_\infty \leq \alpha$

s.t. $\forall i = 1, \dots k$

$$\left| \mathbb{E}_{x \sim \hat{p}} [\varphi_i(x)] - \frac{1}{n} \sum_{j=1}^{n} \varphi_i(x_j) \right| \leq \alpha$$

$\varphi_i : \mathcal{X} \to [0,1]$

$$\left| \langle \varphi_i , \hat{p} \rangle - \langle \varphi_i , h_x \rangle \right| \leq \alpha.$$

$$\vec{\varphi}_i = \begin{pmatrix} \varphi_i(u_1) \\ \varphi_i(u_2) \\ \vdots \\ \varphi_i(u_m) \end{pmatrix}$$

Annoying: Absolute Value

Trick: consider F that is closed under complement.

For each $\varphi_i \in F$, $1 - \varphi_i \in F$.

Example (threshold):

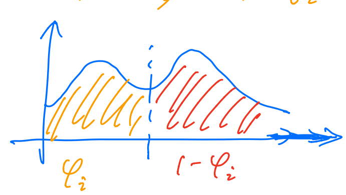

$$\left| \langle \varphi_i , \hat{p} \rangle - \langle \varphi_i , h_x \rangle \right|$$

$$= \max \Big(\big(\langle \varphi_i , \hat{p} \rangle - \langle \varphi_i , h_x \rangle \big) , \big(\langle 1-\varphi_i , \hat{p} \rangle - \langle 1-\varphi_i , h_x \rangle \big) \Big)$$

If F is closed under complement

$$
\begin{aligned}
\text{error}(\hat{p}) &= \max_i \left| \langle \varphi_i, \hat{p} \rangle - \langle \varphi_i, h_x \rangle \right| \\
&= \max_i \left(\langle \varphi_i, \hat{p} \rangle - \langle \varphi_i, h_x \rangle \right)
\end{aligned}
$$

From Online Learning to Query Release

Goal: Design M, $x \to \boxed{M} \to \hat{p}$, $\max_{i \in F} \langle \varphi_i, \hat{p} - h_x \rangle \leq \alpha$

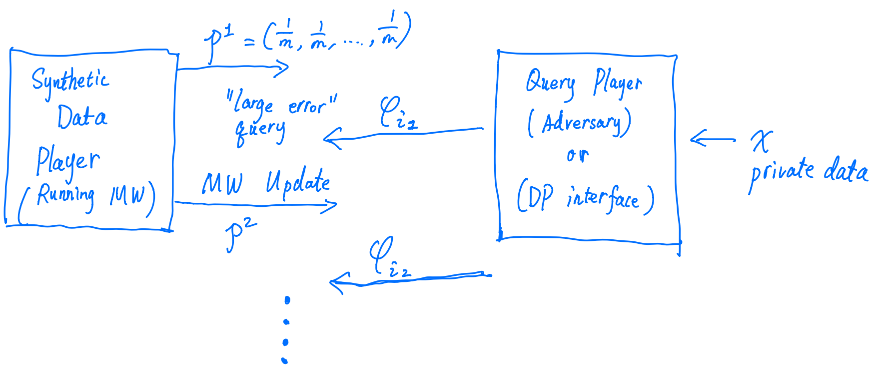

Multiplicative Weights w/ Exponential Mechanism (MWEM)

$p^1 \leftarrow (\frac{1}{m}, \dots, \frac{1}{m})$

for $t = 1, \dots T$.

$\quad i_t \leftarrow M_0(x, \varepsilon_0, p^t)$

$\quad c^t \leftarrow \varphi_{i_t}$

$\quad p^{t+1} \leftarrow \text{MW-Update}(p^t, c^t, \eta)$

Return $\hat{p} = \frac{1}{T} \sum_t p^t$

Privacy?

Selection Problem.

Exponential Mechanism (Report Noisy Max).

Score function. "error"

$q^t(i, x) = \langle \varphi_i, p^t - h_x \rangle$

Privacy Proof: A composition of T exponential mechanisms.