# Smoothed Analysis in Learning

By: Anish Jayant

# Standard Models of Learning

**Offline (PAC) learning**

- Adversary picks a worst-case distribution $P$ and learner receives *i.i.d.* samples
- Constrained to a function class $F$, the learner outputs predictor $f \in F$
- The learner's performance graded against best $f^* \in F$ on a fresh sample from $P$

**Online Learning**

- Adversary picks a worst case $(x_t, y_t)$ while player picks $h_t: X \to Y$ simultaneously
- Player incurs loss $\ell\big(y_t, h_t(x_t)\big)$
- Game continues for rounds $t \in [T]$ and performance graded on *regret,* the best fixed action $f^* \in F$ in hindsight

# Standard Models of Learning

**Offline (PAC) learning**

- Adversary picks a worst-case distribution $\mathrm{P} \in \mathcal{P}$ and learner receives *i.i.d.* samples
- Constrained to a function class $\mathcal{F}$, the learner outputs predictor $f \in \mathcal{F}$
- The learner's performance graded against best $f^* \in \mathcal{F}$ on a fresh sample from $P$

Characterized by **VC Dimension** of $\mathcal{F}$, originates from *empirical process theory*

**Online Learning**

- Adversary picks a worst case $(x_t, y_t)$ while player picks $h_t: X \to Y$ simultaneously
- Player incurs loss $\ell(y_t, h_t(x_t))$
- Game continues for rounds $t \in [T]$ and performance graded on *regret,* the best fixed action $f^* \in F$ in hindsight

Characterized by **Littlestone Dimension** of $F$, generalizes VC analysis using martingales

# Smoothed Online Learning

Question [RST'11]: If the adversary's suggests distributions instead of points (*i.e.*, noise perturbation, semi-random behavior), can we recover learnability? (in the VC theory sense)

# Smoothed Online Learning

Question [RST'11]: If the adversary's suggests distributions instead of points (*i.e.*, noise perturbation, semi-random behavior), can we recover learnability?

Smoothed Online Learning: At each $t \in [T]$ adversary may suggest $p_t \in \Delta(\mathcal{X})$ as long as $\frac{dp_t(x)}{d\mu(x)} \leq \sigma^{-1}$, for $\sigma \in [0, 1]$.

# Smoothed Online Learning

<u>Question</u> [RST'11]: If the adversary's suggests distributions instead of points (*i.e.*, noise perturbation, semi-random behavior), can we recover learnability?

<u>Smoothed Online Learning:</u> At each $t \in [T]$ adversary may suggest $p_t \in \Delta(\mathcal{X})$ as long as $\frac{dp_t(x)}{d\mu(x)} \leq \sigma^{-1}$, for $\sigma \in [0, 1]$.

<u>"Answer"</u> [*e.g.* Haghtalab'18]: If the adversary suggests all distributions *in advance* then VC dimension characterizes learnability. ~ (Learning without *i.i.d* assumption)

# Smoothed Online Learning

Question [RST'11]: If the adversary's suggests distributions instead of points (*i.e.*, noise perturbation, semi-random behavior), can we recover learnability?

Smoothed Online Learning: At each $t \in [T]$ adversary may suggest $p_t \in \Delta(\mathcal{X})$ as long as $\frac{dp_t(x)}{d\mu(x)} \leq \sigma^{-1}$, for $\sigma \in [0, 1]$.

"Answer" [*e.g.* Haghtalab'18]: If the adversary suggests all distributions *in advance* then VC dimension characterizes learnability. ~ (Learning without *i.i.d* assumption)

"Answer" [*e.g.* HRS'24]: Even if the adversary suggests distributions adaptively, then VC dimension again characterizes learnability ~ (Reduces to oblivious setting by coupling arguments)

# Smoothed Online Learning

Question [RST'11]: If the adversary's suggests distributions instead of points (*i.e.*, noise perturbation, semi-random behavior), can we recover learnability?

Smoothed Online Learning: At each $t \in [T]$ adversary may suggest $p_t \in \Delta(\mathcal{X})$ as long as $\frac{dp_t(x)}{d\mu(x)} \leq \sigma^{-1}$, for $\sigma \in [0, 1]$.

"Answer" [*e.g.* Haghtalab'18]: If the adversary suggests all distributions *in advance* then VC dimension characterizes learnability. ~ (Learning without *i.i.d* assumption)

"Answer" [*e.g.* HRS'24]: Even if the adversary suggests distributions adaptively, then VC dimension again characterizes learnability ~ (Reduces to oblivious setting by coupling arguments)

"Answer" [*e.g.* BRS'24]: Learning is possible in smoothed online setting *without* knowledge of base measure through ERM – which fails in traditional online! ~ (natural exploration/exploitation, sharpened by [Bla'24])

# Smoothed Online Learning

Question [RST'11]: If the adversary's suggests distributions instead of points (*i.e.*, noise perturbation, semi-random behavior), can we recover learnability?

Smoothed Online Learning: At each $t \in [T]$ adversary may suggest $p_t \in \Delta(\mathcal{X})$ as long as $\frac{dp_t(x)}{d\mu(x)} \leq \sigma^{-1}$, for $\sigma \in [0, 1]$.

"Answer" [*e.g.* Haghtalab'18]: If the adversary suggests all distributions *in advance* then VC dimension characterizes learnability. ~ (Learning without *i.i.d* assumption)

"Answer" [*e.g.* HRS'24]: Even if the adversary suggests distributions adaptively, then VC dimension again characterizes learnability ~ (Reduces to oblivious setting by coupling arguments)

"Answer" [*e.g.* BRS'24]: Learning is possible in smoothed online setting *without* knowledge of base measure through ERM – which fails in traditional online! ~ (natural exploration/exploitation, sharpened by [Bla'24])

...is the story over?

- How can we expand the toolkit from offline and online algorithms to smoothened setting?
- Is the "bounded derivative" setting the appropriate interpretation of smoothening for learning problems?

# Theoretical Thinking

- Relaxing bounded likelihood ratio restriction to more general adversary distributions $\mathcal{U}$.
  - Initially needed for a coupling step in the proof that uses *rejection sampling*
  - A first step made by [BP'23] relaxes the ratio to general closeness in $f$-divergence, and shows rejection sampling technique continues to apply
  - Is rejection sampling *really* the correct way to think about this?

# Theoretical Thinking

- Relaxing bounded likelihood ratio restriction to more general adversary distributions $\mathcal{U}$.
  - Initially needed for a coupling step in the proof that uses *rejection sampling*
  - A first step made by [BP'23] relaxes the ratio to general closeness in $f$-divergence, and shows rejection sampling technique continues to apply
  - Is rejection sampling *really* the correct way to think about this?

- Efficient Algorithms
  - For binary classification, ERM is the staple yet notably fails for online learning [HK16]. How should we learn in smoothened online learning, given that ERM works again [BRS24]
  - A clever Hedge-based technique of [Bla24] further improves rates but is utterly inefficient (constructing coverings)

# Theoretical Thinking

- Relaxing bounded likelihood ratio restriction to more general adversary distributions $\mathcal{U}$.
  - Initially needed for a coupling step in the proof that uses *rejection sampling*
  - A first step made by [BP'23] relaxes the ratio to general closeness in $f$-divergence, and shows rejection sampling technique continues to apply
  - Is rejection sampling *really* the correct way to think about this?

- Efficient Algorithms
  - For binary classification, ERM is the staple yet notably fails for online learning [HK16]. How should we learn in smoothened online learning, given that ERM works again [BRS24]
  - A clever Hedge-based technique of [Bla24] further improves rates but is utterly inefficient (constructing coverings)

- A *universal* learning characterization?
  - Does there exist a characterization that generalizes both VC and Littlestone dimension?
  - A first step made by [BK25] by studying the pair $(\mathcal{F}, \mathcal{U})$ on an "interaction tree," which recovers near-optimal rates in *both* offline and prior smoothed online learning.