Do you pay for Privacy in Online learning?

Conference On Learning Theory 2022

Giorgia Ramponi, Amartya Sanyal

ETH AI Center, Zurich, Switzerland

Online Learning in the Mistake Bound Model

- Let X be the instance space

Online Learning in the Mistake Bound Model

- Let X be the instance space
- At each time step t, the learner

Online Learning in the Mistake Bound Model

- Let X be the instance space
- At each time step t, the learner
    - receives an unlabelled example $x_t \in X$,

Online Learning in the Mistake Bound Model

- Let X be the instance space
- At each time step t, the learner
  - receives an unlabelled example $x_t \in X$,
  - predicts a label $\hat{y}_t$ for x_t, and then

Online Learning in the Mistake Bound Model

- Let X be the instance space
- At each time step t, the learner
  - receives an unlabelled example $x_t \in X$,
  - predicts a label $\hat{y}_t$ for x_t, and then
  - receives the true label y_t for x_t.

Online Learning in the Mistake Bound Model

- Let X be the instance space
- At each time step t, the learner
  - receives an unlabelled example $x_t \in X$,
  - predicts a label $\hat{y}_t$ for x_t, and then
  - receives the true label y_t for x_t.

The performance is measured by the number of mistakes:

$$\text{Mistakes}(T) := \sum_{t=1}^{T} (\hat{y}_t \neq y_t).$$

Online Learning in the Mistake Bound Model

- Let X be the instance space
- At each time step t, the learner
  - receives an unlabelled example $x_t \in X$,
  - predicts a label $\hat{y}_t$ for x_t, and then
  - receives the true label y_t for x_t.

The performance is measured by the number of mistakes:

$$\text{Mistakes}(T) := \sum_{t=1}^{T} (\hat{y}_t \neq y_t).$$

Realisable case: $\exists h^* \in \mathcal{H}$ such that $y_t = h^*(x_t)$ for all $t \leq T$.

Online Learnability

A hypothesis class $\mathcal{C}$ on X is **online learnable** if there exists a learner L that makes at most $M < \infty$ mistakes on any sequence of length T labelled by any $h^* \in \mathcal{H}$, for all T.

[1] Nick Littlestone. "Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm". In: *Machine learning* 2.4 (1988), pp. 285–318.

Online Learnability

A hypothesis class $\mathcal{C}$ on X is **online learnable** if there exists a learner L that makes at most $M < \infty$ mistakes on any sequence of length T labelled by any $h^* \in \mathcal{H}$, for all T.

The combinatorial measure **Littlestone dimension** or $\mathrm{Ldim}(\mathcal{H})$ exactly characterises online learnability of $\mathcal{H}$[1].

[1] Nick Littlestone. "Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm". In: *Machine learning* 2.4 (1988), pp. 285–318.

Online Learnability

A hypothesis class $\mathcal{C}$ on X is **online learnable** if there exists a learner L that makes at most $M < \infty$ mistakes on any sequence of length T labelled by any $h^* \in \mathcal{H}$, for all T.

The combinatorial measure **Littlestone dimension** or $\mathrm{Ldim}(\mathcal{H})$ exactly characterises online learnability of $\mathcal{H}$[1].

$$\mathrm{Ldim}(\mathcal{H}) \geq \mathrm{VC}(\mathcal{H}) \text{ (possibly arbitrarily larger)}$$

[1] Nick Littlestone. "Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm". In: *Machine learning* 2.4 (1988), pp. 285–318.

Online Learnability

A hypothesis class $\mathcal{C}$ on X is **online learnable** if there exists a learner L that makes at most $M < \infty$ mistakes on any sequence of length T labelled by any $h^* \in \mathcal{H}$, for all T.

The combinatorial measure **Littlestone dimension** or $\mathrm{Ldim}(\mathcal{H})$ exactly characterises online learnability of $\mathcal{H}$[1].

$$\mathrm{Ldim}(\mathcal{H}) \geq \mathrm{VC}(\mathcal{H}) \text{ (possibly arbitrarily larger)}$$

Thus, **online learnability** is harder than offline learnability.

[1] Nick Littlestone. "Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm". In: *Machine learning* 2.4 (1988), pp. 285–318.

Differential Privacy

An **offline learning** algorithm $\mathcal{A} : \mathcal{X} \to \mathcal{Y}$ is said to be (ϵ, δ)-differentially private if for any two datasets S_1, S_2 that differ in just one element:

$$\mathbb{P}\left[\mathcal{A}\left(S_1\right) \in Q\right] \le e^{\epsilon}\mathbb{P}\left[\mathcal{A}\left(S_2\right) \in Q\right] + \delta$$

[2] Prateek Jain, Pravesh Kothari, and Abhradeep Thakurta. "Differentially private online learning". In: *Conference on Learning Theory*. JMLR Workshop and Conference Proceedings. 2012, pp. 24–1.

[3] Naman Agarwal and Karan Singh. "The price of differential privacy for online learning". In: *International Conference on Machine Learning*. PMLR. 2017, pp. 32–40.

Differential Privacy

An **offline learning** algorithm $\mathcal{A} : \mathcal{X} \to \mathcal{Y}$ is said to be (ϵ, δ)-differentially private if for any two datasets S_1, S_2 that differ in just one element:

$$\mathbb{P}\left[\mathcal{A}\left(S_1\right) \in Q\right] \leq e^{\epsilon}\mathbb{P}\left[\mathcal{A}\left(S_2\right) \in Q\right] + \delta$$

An **online learning** algorithm $\mathcal{A}$ is $\{\epsilon, \delta\}$-online differentially private[2][3] if for all $T \in \mathbb{N}$, for any two sequences of T points $S_T, S'_T \in (X \times \mathcal{Y})^T$ that differ in one entry the following holds:

$$\mathbb{P}\left[\mathcal{A}(S_T) \in Q\right] \leq e^{\epsilon}\mathbb{P}\left[\mathcal{A}(S'_T) \in Q)\right] + \delta$$

[2] Prateek Jain, Pravesh Kothari, and Abhradeep Thakurta. "Differentially private online learning". In: *Conference on Learning Theory*. JMLR Workshop and Conference Proceedings. 2012, pp. 24–1.
[3] Naman Agarwal and Karan Singh. "The price of differential privacy for online learning". In: *International Conference on Machine Learning*. PMLR. 2017, pp. 32–40.

Abbreviations

- Non-Private Offline Learning (PAC) → NP-Off
- Non-Private Online Learning (Online learnability) → NP-MB
- Private Offline Learning (PAC + Offline DP) → P-Off
- Private Online Learning → P-MB

Learning hierarchy

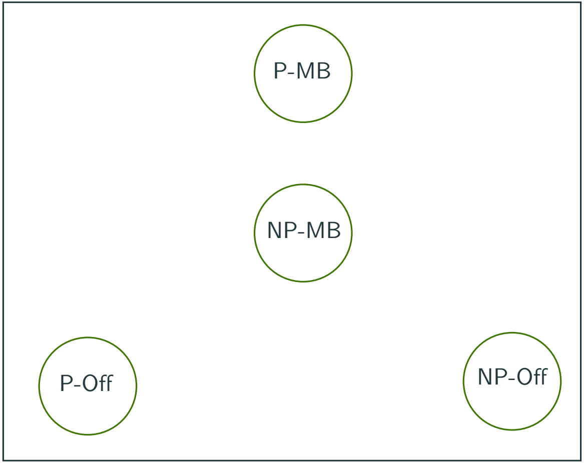

Learning hierarchy

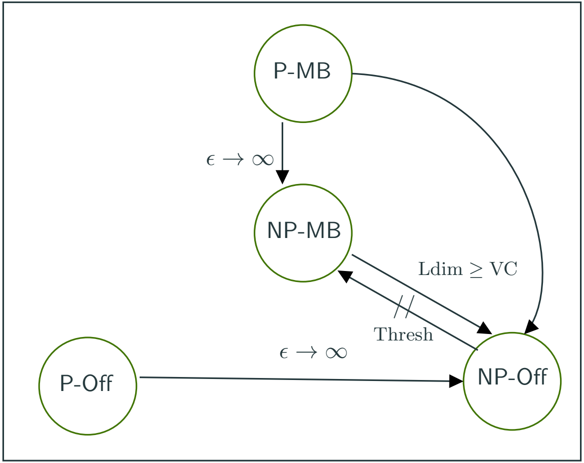

Learning hierarchy

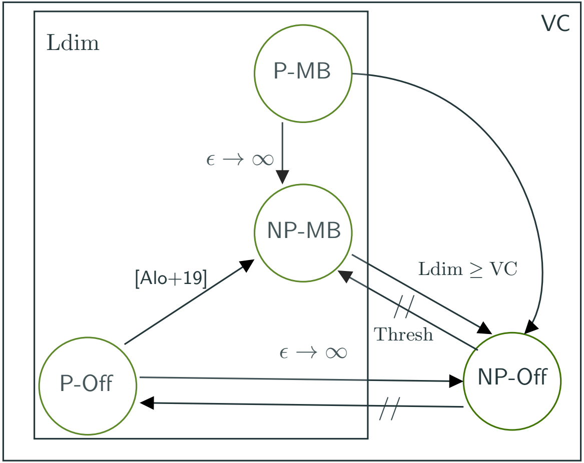

Learning hierarchy

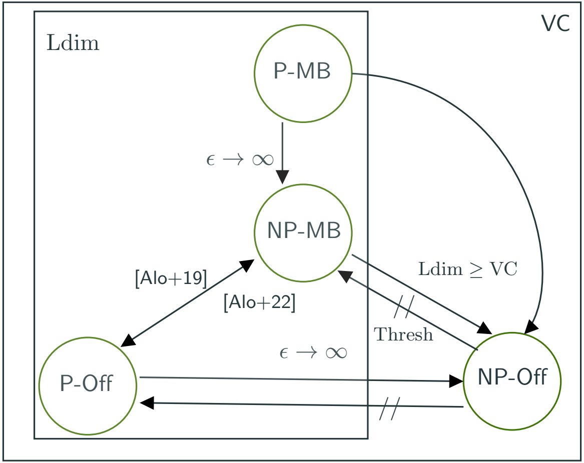

Learning hierarchy

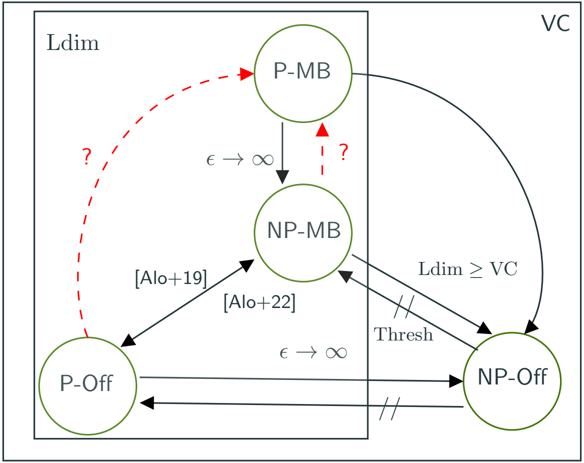

Learning hierarchy

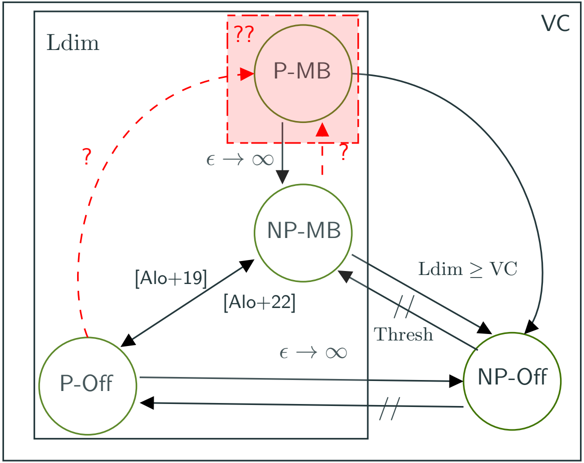

Open problem P-MB $\not\Rightarrow$ NP-MB

Theorem

There exists a hypothesis class such that for any sequence of points of length T labelled by $h^* \in \mathcal{H}$,

1. *(Online learnable) there exists an online algorithm $\mathcal{A}$ that does not make more than M mistakes on the sequence S_T for some $M < \infty$.*
2. *(Not privately online learnable) any (ϵ, δ)-differentially private online algorithm makes at least $M' \geq M + \alpha(\epsilon, \delta, T)$ mistakes,*

where $\alpha : \mathbb{R} \times [0, 1] \to \mathbb{N}$ is such that $\alpha(\epsilon, \delta, T) \gtrsim_\delta \frac{\sqrt{T}}{\epsilon}$.

Open Problem: P-MB $\implies$ NP-MB

Theorem
For any online-learnable hypothesis class $\mathcal{H}$, there exists a positive, monotonically decreasing function $\gamma = o\left(\sqrt{\cdot}\right)$ such that for all $T \in \mathbb{N}$, and $\epsilon : \gamma\left(T\right) \gtrsim \epsilon \gtrsim \sqrt{T}$ there exists an (ϵ, δ)-differentially private online algorithm that makes at most $M < \infty$ mistakes for any sequence of points $S_T = \{\left(x_1, h^*\left(x_1\right)\right), \ldots, \left(x_T, h^*\left(x_T\right)\right)\}$ of length T labelled by $h \in \mathcal{H}$.

Thanks for the attention!

- P-MB $\not\Longrightarrow$ NP-MB 
- P-MB $\Longrightarrow$ NP-MB