Understanding the Role of Adaptivity in Machine Teaching
The Case of Version Space Learners

Caltech

Yuxin Chen — Adish Singla — Oisin Mac Aodha — Pietro Perona — Yisong Yue





MAX PLANCK INSTITUTE FOR SOFTWARE SYSTEMS

1 Introduction

Motivating Applications

Citizen science, crowdsourcing services, medical diagnosis

iNaturalist — Explore Learn Record

Learning vs. Teaching Setting

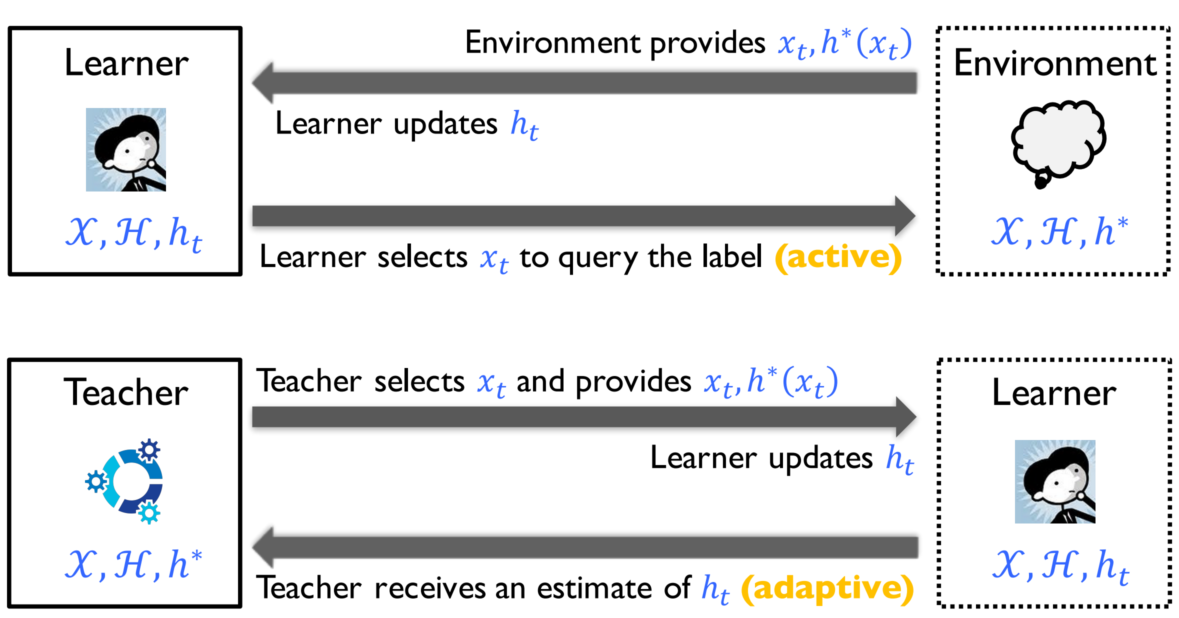

Canonical Example: 1-D Threshold Classifier

Discrete points on a line $\mathcal{X}=\{1, 2, \dots, n\}$

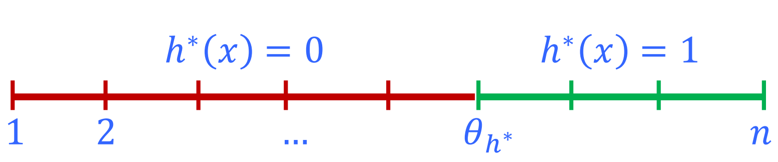

Threshold classifier $h(x)=1$ iff $x \geq \theta_h$ where $\theta_h \in \{1, 2, \dots, n\}$

	complexity
Passive learning	Θ(n)
Active learning	Θ(log(n))
Non-adaptive teaching	2
Adaptive teaching	1*

How much speed up a teacher can achieve from adaptivity?

1*: under additional restriction on learner's update rule

2 Teaching Model

Teaching a "Version Space" Learner

$\mathcal{X}, \mathcal{H}$: Discrete, finite sets

Learner starts from initial hypothesis $h_0 \in \mathcal{H}$

At time t:

- teacher provides $x_t, h^*(x_t)$
- learner updates the hypothesis space $\mathcal{H}_{t+1} = \mathcal{H}_t \cap \mathcal{H}(x_t)$
- learner selects a new hypothesis $h_{t+1} \in \mathcal{H}_{t+1}$ randomly

Teaching stops when $h_t = h^*$

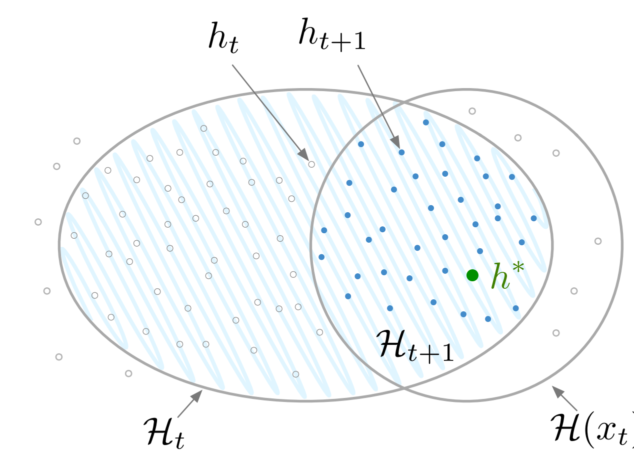

State-dependent Preference

Learner's preference of next hypothesis depends on the **version space**, as well as the **current hypothesis**

Learner's preference function $\sigma: \mathcal{H} \times \mathcal{H} \to \mathbb{R}$

Given current hypotheses h_t and two hypotheses

- $\sigma(h_i; h_t) < \sigma(h_j; h_t)$: Learner prefers to pick h_i instead of h_j
- $\sigma(h_i; h_t) = \sigma(h_j; h_t)$: Learner could pick either one of these two randomly

At time t, the learner selects a new hypothesis h_{t+1} randomly from $\{h \in \mathcal{H}_{t+1}: \sigma(h; h_t) = \min_{h' \in \mathcal{H}_{t+1}} \sigma(h'; h_t)\}$

Special Cases: State-independent Preference

Classical model (TD) [Goldman, Kearns '92]:
Learner picks next hypothesis at random
$\forall h, h' \in \mathcal{H}: \sigma(h'; h) = c$

TS(h*): *Optimal teaching sequence for* h^*
Equivalent to set cover of $\mathcal{H} \setminus \{h^*\}$ by $\mathcal{X}$

Teaching Dimension $TD(\mathcal{H}, \mathcal{X}) := \max_{h^* \in \mathcal{H}} |TS(h^*)|$

Global preference-based model ($PBTD$) [Gao et al. '16]
Learner picks next hyp. based on a global preference
$\forall h' \in \mathcal{H}: \sigma(h'; h) = c_{h'}$ ⟵ *a constant only depends on* h'

TS(h*): *Optimal teaching sequence for* h^*
Given by the following notion of set cover:
$\min_X |X|$, s.t. $\forall h \in \mathcal{H}(X) \setminus \{h^*\}: \sigma(h; \cdot) > \sigma(h^*; \cdot)$

Proposition [*a necessary condition to gain from adaptivity*]
Learner must have **state-dependent** preferences: Choice of next hypothesis $h_{t+1} \in \mathcal{H}_{t+1}$ depends on h_t

Theorem There exist hypothesis classes with state-dependent preferences, where the optimal non-adaptive teacher, in the worst case, requires exponentially more teaching examples than the optimal adaptive teacher.

3 Examples: State-dependent Pref.

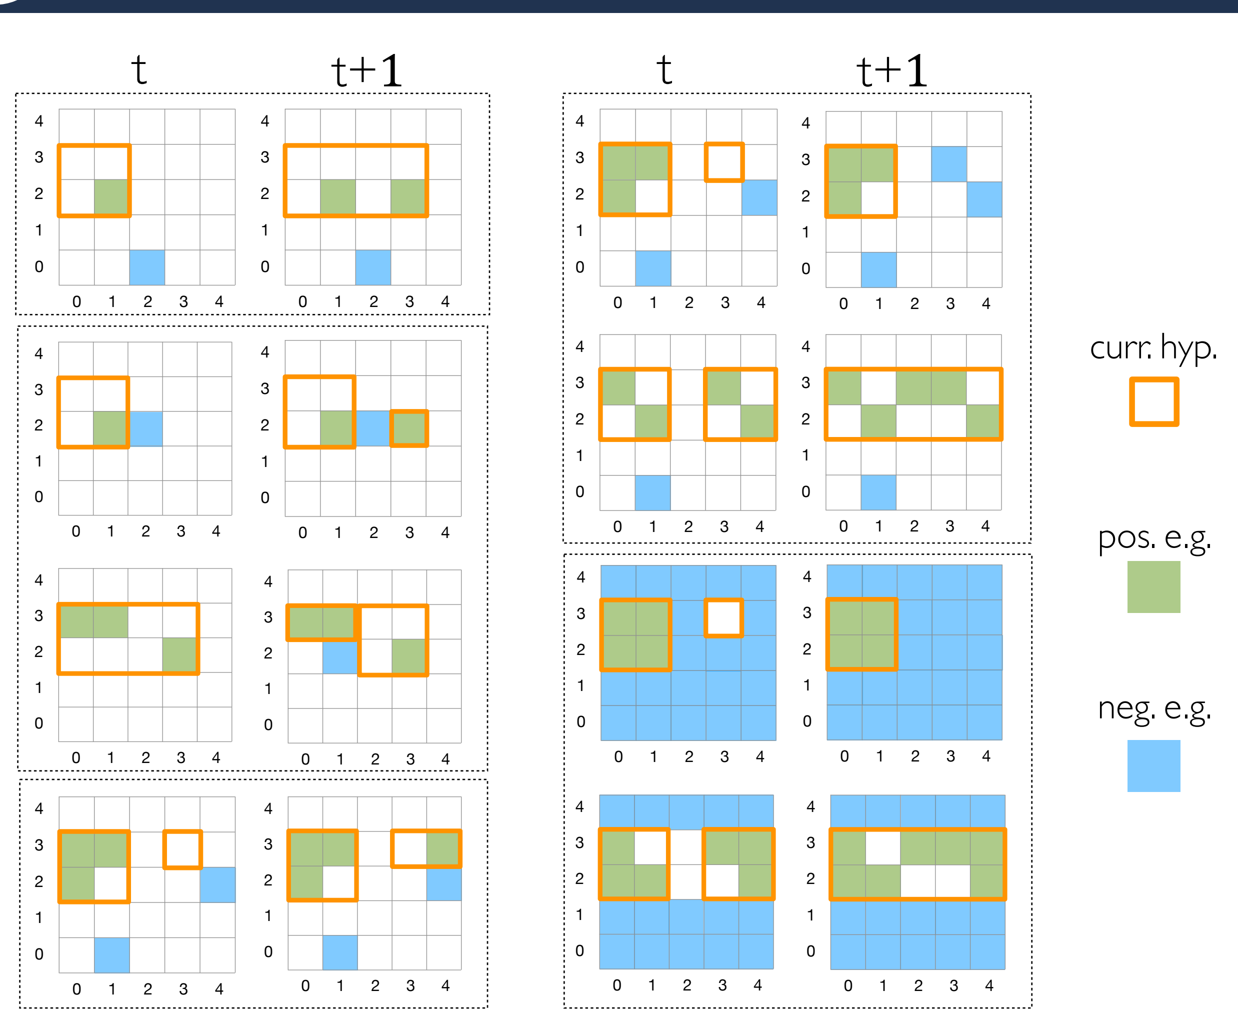

2-Rec: Disjoint union of geometric objects [Gao et al. '17]

Difficulty of teaching: $TD(\mathcal{H}, \mathcal{X}) = O(n^2)$ for $n \times n$ grid size

Our model of preferences $\sigma_{2\text{-}Rec}$

$\sigma_{2\text{-}Rec}$: Prefer hypotheses in the same complexity subclass

$\sigma_{2\text{-}Rec}$: Within same subclass, prefer hypothesis with min. edge edits

4 Experimental Results

Teaching Algorithms

Random
Randomly chosen examples; stops when $h_t = h^*$

Classical
Set cover for $\mathcal{H} \setminus \{h^*\}$; stops when $h_t = h^*$

2R-NonAdaT
Observes h_0
Uses $\sigma_{2\text{-}Rec}$ and h_0 to optimally select examples
Teaching stops when $h_t = h^*$

2R-AdaT
Observes h_t $\forall$ t
Uses $\sigma_{2\text{-}Rec}$ and h_t to optimally select an example at t
Teaching stops when $h_t = h^*$

2-Rec Class: Simulated Learners

- $n \times n$ grid size; $h_0 \in \mathcal{H}^2$, $h^* \in \mathcal{H}^1$
- 50 simulated learners with $\sigma_{2\text{-}Rec}$ preferences

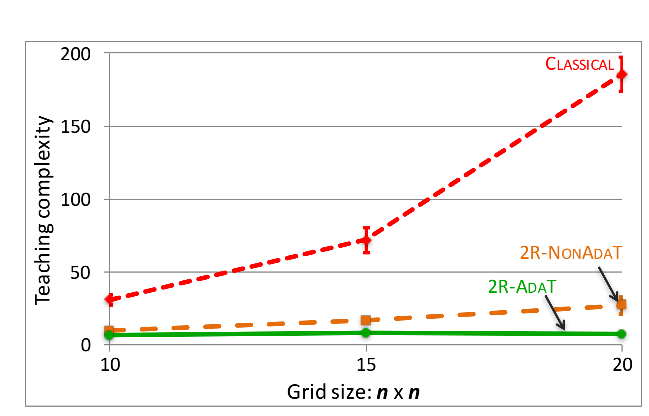

Classical: $O(n^2)$

2R-NonAdaT: $O(|h_0|) = O(n^2)$

2R-AdaT: $O(\log(|h_0|)) = O(\log(n^2))$

2-Rec Class: Human Learners

Preference elicitation:
Users were asked to update the position of the orange rectangle so that green cells are inside and blue cells are outside

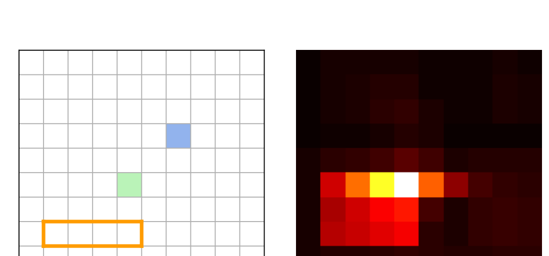

Preference elicitation:
Participants favor staying at their current hypothesis if it remains valid, along with preferring smaller updates.

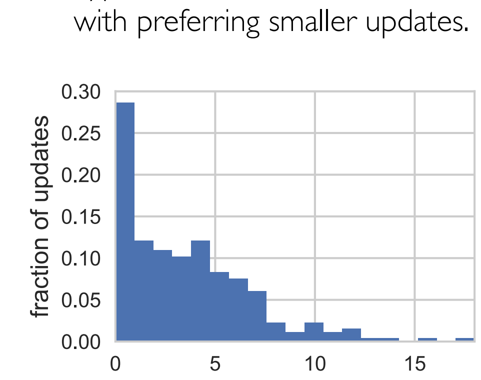

- 8×8 grid size; $h_0 \in \mathcal{H}^2$, $h^* \in \mathcal{H}^1$
- 200 participants from a crowdsourcing platform
- **2R-AdaT** and **2R-NonAdaT** teachers use $\sigma_{2\text{-}Rec}$