Testably learning polynomial threshold functions

Thirty-Eighth Annual Conference on Neural Information Processing Systems

Lucas Slot

ETH Zurich

Stefan Tiegel

ETH Zurich

Manuel Wiedmer

ETH Zurich

Introduction: Agnostic learning of a concept class $\mathcal{F}$

Given: samples $(x_i, y_i) \in \mathbb{R}^n \times \{\pm 1\}$.

Goal: Find a classifier $\hat{f}$ such that

$$\mathbb{P}[\hat{f}(x) \neq y] \leq \mathrm{opt} + \varepsilon,$$

where $\mathrm{opt} = \min_{f \in \mathcal{F}} \mathbb{P}[f(x) \neq y]$, using as few samples and time as possible.

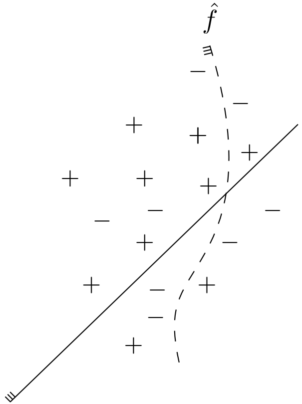

$$f_{\mathrm{opt}} = \arg\min_{f \in \mathcal{F}} \mathbb{P}[f(x) \neq y]$$

$\mathrm{opt} = \frac{6}{18}$

$\mathbb{P}[\hat{f}(x) \neq y] = \frac{7}{18}$

Introduction: Agnostic learning of a concept class $\mathcal{F}$

Given: samples $(x_i, y_i) \in \mathbb{R}^n \times \{\pm 1\}$.

Goal: Find a classifier $\hat{f}$ such that

$$\mathbb{P}[\hat{f}(x) \neq y] \leq \mathrm{opt} + \varepsilon,$$

where $\mathrm{opt} = \min_{f \in \mathcal{F}} \mathbb{P}[f(x) \neq y]$, using as few samples and time as possible.

Issue: Generally computationally hard
$\rightarrow$ Add distributional assumption that $x_i \sim \mathcal{D}$ for some (known) $\mathcal{D}$.

For this work: Assume samples come from the standard Gaussian.

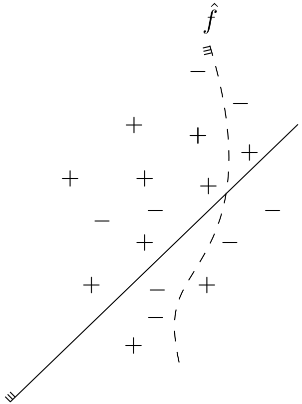

$$f_{\mathrm{opt}} = \arg\min_{f \in \mathcal{F}} \mathbb{P}[f(x) \neq y]$$

$\mathrm{opt} = \frac{6}{18}$
$\mathbb{P}[\hat{f}(x) \neq y] = \frac{7}{18}$

Testable learning — Definition

Question: How to verify this assumption?

Checking the distributional assumption directly is computationally hard or impossible.

Testable learning — Definition

Question: How to verify this assumption?

Checking the distributional assumption directly is computationally hard or impossible.

Recently introduced model of testable learning [RV23]:

1. (Tester) Check computationally tractable relaxation of distributional assumption. Accept or reject the samples.
2. (Learner) If we accept the sample, run learning algorithm.

[RV23]: Rubinfeld and Vasilyan. "Testing distributional assumptions of learning algorithms" (2023)

Testable learning — Definition

Question: How to verify this assumption?

Checking the distributional assumption directly is computationally hard or impossible.

Recently introduced model of testable learning [RV23]:

1. (Tester) Check computationally tractable relaxation of distributional assumption. Accept or reject the samples.
2. (Learner) If we accept the sample, run learning algorithm.

Conditions:

- (Completeness) Samples from the target distribution are accepted.
- (Soundness) Whenever tester accepts, learner needs to output a good hypothesis.

[RV23]: Rubinfeld and Vasilyan. "Testing distributional assumptions of learning algorithms" (2023)

Testable learning — Results

Concept class	Agnostic learning	Testable learning
Halfspaces	$n^{\tilde{O}(1/\varepsilon^2)}$ [KKMS08; DKN10]	$n^{\tilde{O}(1/\varepsilon^2)}$ [RV23; GKK23]
Degree-d PTFs		
Convex sets		

[KKMS08]: Kalai, Klivans, Mansour, and Servedio. "Agnostically learning halfspaces" (2008)
[DKN10]: Diakonikolas, Kane, and Nelson. "Bounded independence fools degree-2 threshold functions" (2010)
[RV23]: Rubinfeld and Vasilyan. "Testing distributional assumptions of learning algorithms" (2023)
[GKK23]: Gollakota, Klivans, and Kothari. "A moment-matching approach to testable learning and a new characterization of Rademacher complexity" (2023)

Testable learning — Results

Concept class	Agnostic learning	Testable learning
Halfspaces	$n^{\tilde{O}(1/\varepsilon^2)}$ [KKMS08; DKN10]	$n^{\tilde{O}(1/\varepsilon^2)}$ [RV23; GKK23]
Degree-d PTFs		
Convex sets	$2^{\tilde{O}(\sqrt{n}/\varepsilon^4)}$ [KOS08]	$2^{\Omega(n)}$ [RV23]

[KKMS08]: Kalai, Klivans, Mansour, and Servedio. "Agnostically learning halfspaces" (2008)
[DKN10]: Diakonikolas, Kane, and Nelson. "Bounded independence fools degree-2 threshold functions" (2010)
[RV23]: Rubinfeld and Vasilyan. "Testing distributional assumptions of learning algorithms" (2023)
[GKK23]: Gollakota, Klivans, and Kothari. "A moment-matching approach to testable learning and a new characterization of Rademacher complexity" (2023)
[KOS08]: Klivans, O'Donnell, and Servedio. "Learning Geometric Concepts via Gaussian Surface Area" (2008)

Testable learning — Results

Concept class	Agnostic learning	Testable learning
Halfspaces	$n^{\tilde{O}(1/\varepsilon^2)}$ [KKMS08; DKN10]	$n^{\tilde{O}(1/\varepsilon^2)}$ [RV23; GKK23]
Degree-d PTFs	$n^{O(d^2/\varepsilon^4)}$ [Kan11b]	?
Convex sets	$2^{\tilde{O}(\sqrt{n}/\varepsilon^4)}$ [KOS08]	$2^{\Omega(n)}$ [RV23]

[KKMS08]: Kalai, Klivans, Mansour, and Servedio. "Agnostically learning halfspaces" (2008)
[DKN10]: Diakonikolas, Kane, and Nelson. "Bounded independence fools degree-2 threshold functions" (2010)
[RV23]: Rubinfeld and Vasilyan. "Testing distributional assumptions of learning algorithms" (2023)
[GKK23]: Gollakota, Klivans, and Kothari. "A moment-matching approach to testable learning and a new characterization of Rademacher complexity" (2023)
[KOS08]: Klivans, O'Donnell, and Servedio. "Learning Geometric Concepts via Gaussian Surface Area" (2008)

Testable learning — Results

Concept class	Agnostic learning	Testable learning
Halfspaces	$n^{\tilde{O}(1/\varepsilon^2)}$ [KKMS08; DKN10]	$n^{\tilde{O}(1/\varepsilon^2)}$ [RV23; GKK23]
Degree-d PTFs	$n^{O(d^2/\varepsilon^4)}$ [Kan11b]	$n^{\mathrm{poly}(1/\varepsilon)}$ for constant d **Our result**
Convex sets	$2^{\tilde{O}(\sqrt{n}/\varepsilon^4)}$ [KOS08]	$2^{\Omega(n)}$ [RV23]

[KKMS08]: Kalai, Klivans, Mansour, and Servedio. "Agnostically learning halfspaces" (2008)
[DKN10]: Diakonikolas, Kane, and Nelson. "Bounded independence fools degree-2 threshold functions" (2010)
[RV23]: Rubinfeld and Vasilyan. "Testing distributional assumptions of learning algorithms" (2023)
[GKK23]: Gollakota, Klivans, and Kothari. "A moment-matching approach to testable learning and a new characterization of Rademacher complexity" (2023)
[KOS08]: Klivans, O'Donnell, and Servedio. "Learning Geometric Concepts via Gaussian Surface Area" (2008)

Testable learning — Our result

Theorem (Main result)

Degree-d PTFs can be testably learned with respect to the standard Gaussian in time and sample complexity $n^{\tilde{O}_d\left(\varepsilon^{-4d \cdot 7^d}\right)}$.

Technique: Use "fooling" technique from [GKK23]; proof of this condition is based on [Kan11a].

Open question: Can the dependence on d in the above result be improved or can lower bounds be shown?

[GKK23]: Gollakota, Klivans, and Kothari. "A moment-matching approach to testable learning and a new characterization of Rademacher complexity" (2023)
[Kan11a]: Kane. "k-independent Gaussians fool polynomial threshold functions" (2011)