Attribute-Efficient PAC Learning of Low-Degree Polynomial Threshold Functions with Nasty Noise

Shiwei Zeng
Stevens Institute of Technology

Jie Shen
Stevens Institute of Technology

Problem Setup

- Concept classes

- Polynomial threshold functions (PTFs)
  - $f\colon \mathbb{R}^n \to \{-1,1\}$
  - $f = \mathrm{sign}(\mathrm{g})$, g is polynomial

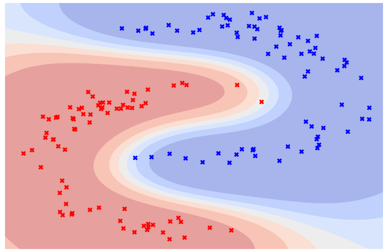

- Attribute-efficiency:
  - g is K-sparse (K attributes)
  - #Samples = $\mathrm{poly}(K \log n)$
- Nasty noise: [BEK02]
  - replace arbitrary η-fraction

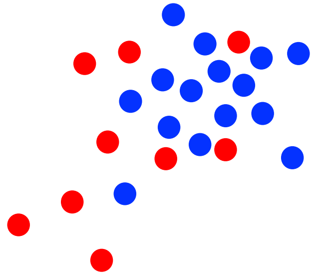

  - Strong contamination:
    - not only labels, but also their features!
    - allow the adversary to inspect

Comparison to Prior Works

	#Sample	Noise Tolerance	Efficient?
[KL88]	$\mathrm{poly}(K, \log n)$	$O\left(\frac{\epsilon}{K^d}\right)$	Y
[BEK02]	$\mathrm{poly}(K, \log n)$	$\Omega(\epsilon)$	N
[DKS18]	$\mathrm{poly}(n)$	$\epsilon^{O(1)}$	Y
This work	$\mathrm{poly}(K, \log n)$	$\epsilon^{O(1)}$	Y

[KL88] Kearns, M. J. and Li, M. Learning in the presence of malicious errors - STOC 1988
[BEK02] Bshouty, N. H., Eiron, N., and Kushilevitz, E. PAC learning with nasty noise – TCS 2022
[DKS18] Diakonikolas, I., Kane, D. M., and Stewart, A. Learning geometric concepts with nasty noise - STOC 2018

Our Results

- Assume that D is Gaussian
- PAC guarantees: $\mathrm{Pr}_D[\hat{f}(x) \neq f(x)] \leq \epsilon$, w.p. $1 - \delta$

- Sample complexity $\tilde{O}(\frac{d^{5d} K^{4d}}{\epsilon^{2d+2}} \log^{5d} n)$
- Noise tolerance $\eta \leq O(\frac{\epsilon^{d+1}}{d^{2d}})$
- Computational efficiency

Main Techniques

- Sparse Chow parameter recovery
  - Sparse f ➡ sparse Chow parameters $\chi_f = \mathbb{E}_D[f \cdot m(x)]$
  - Gaussian D, $m(x) =$ Hermite
  - Estimating sparse χ_f ➡ $\mathrm{poly}(K \log n)$ samples.

- Attribute-efficient outlier removal
  - Degree-d covariance matrix
  - Frobenius-norm filter that recovers D, and it alone!