

Improved learning of k -parities

Arnab Bhattacharyya

Indian Institute of Science
India

Ameet Gadekar

Aalto University
Finland

Ninad Rajgopal

University of Oxford
UK

COCOON, Qingdao, China, July 2018


The Learning Parity Problem

Push




Hidden Vector f


The Learning Parity Problem



The Learning Parity Problem



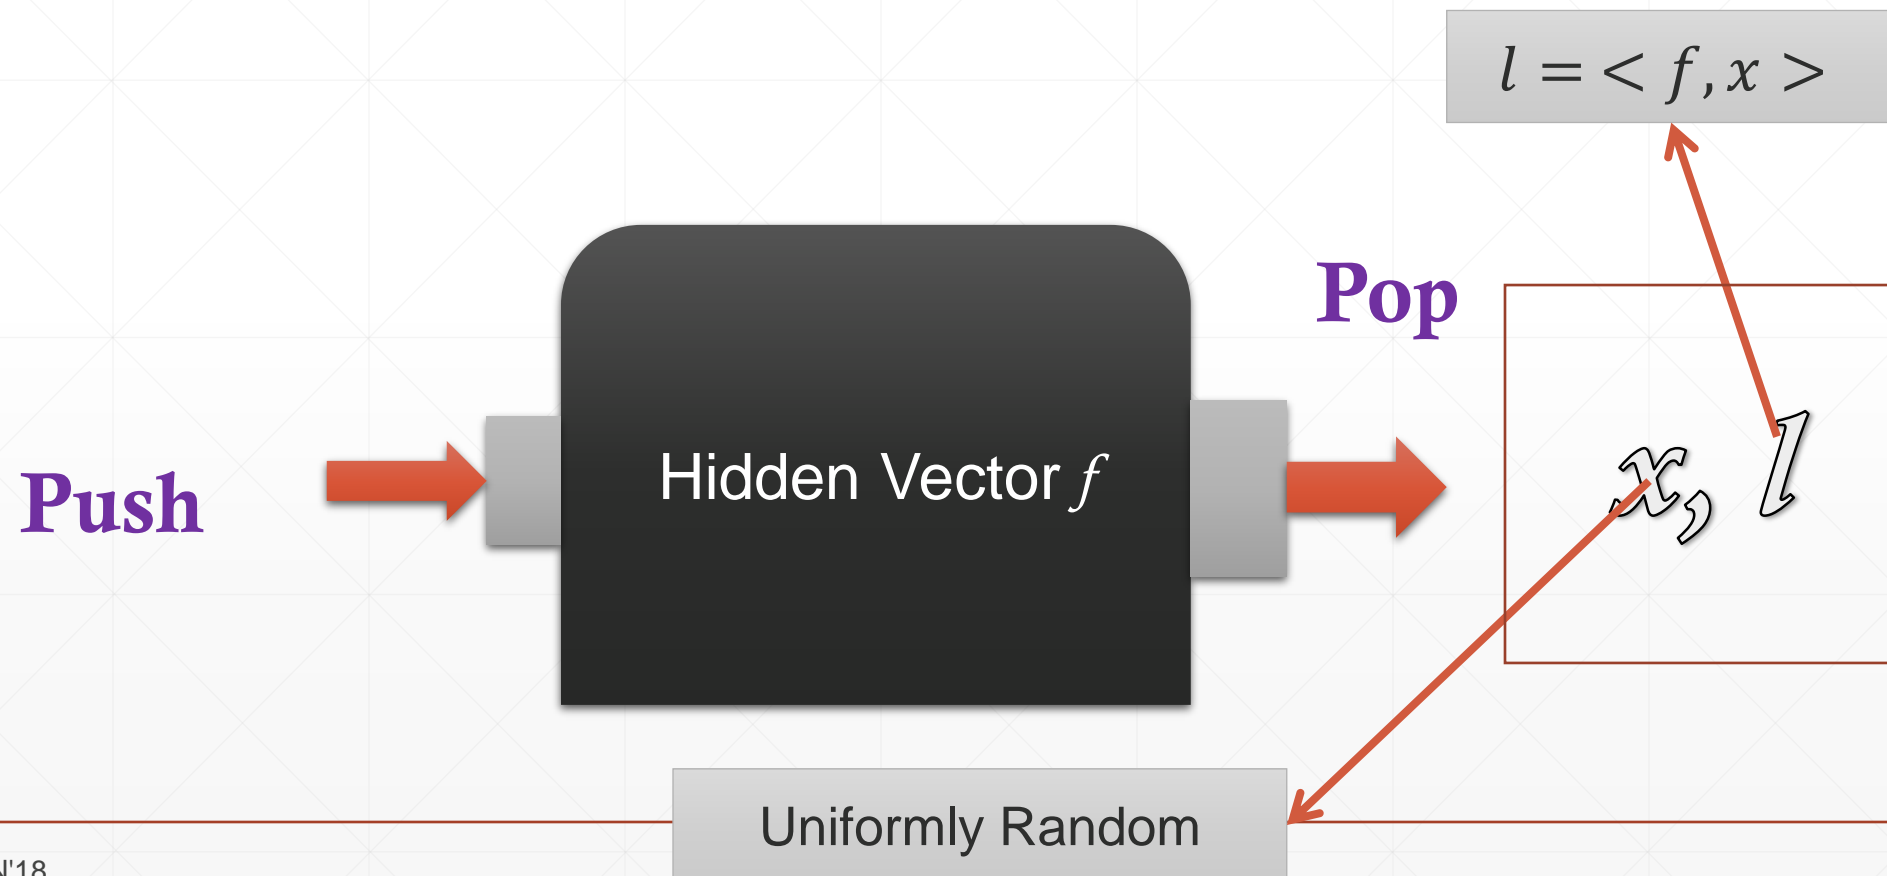
The Learning Parity Problem



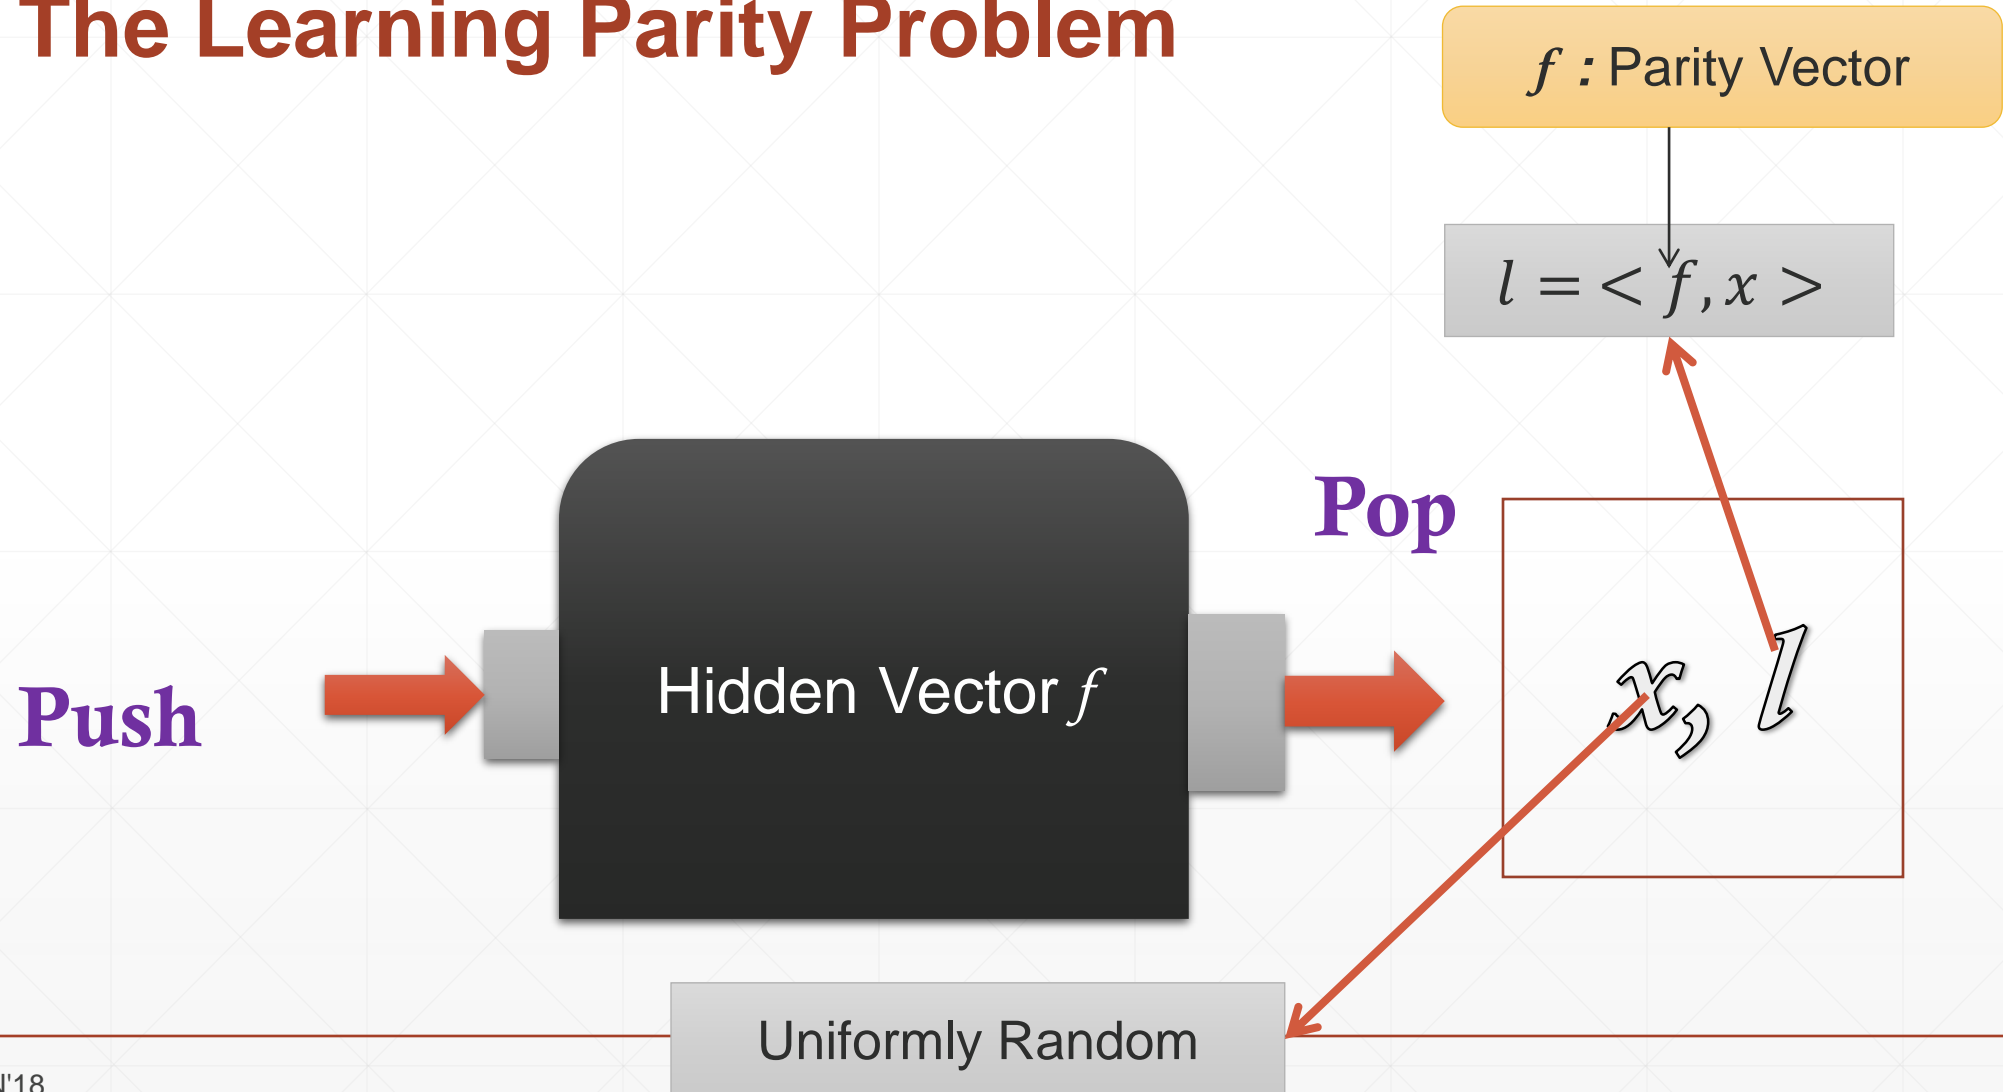
The Learning Parity Problem



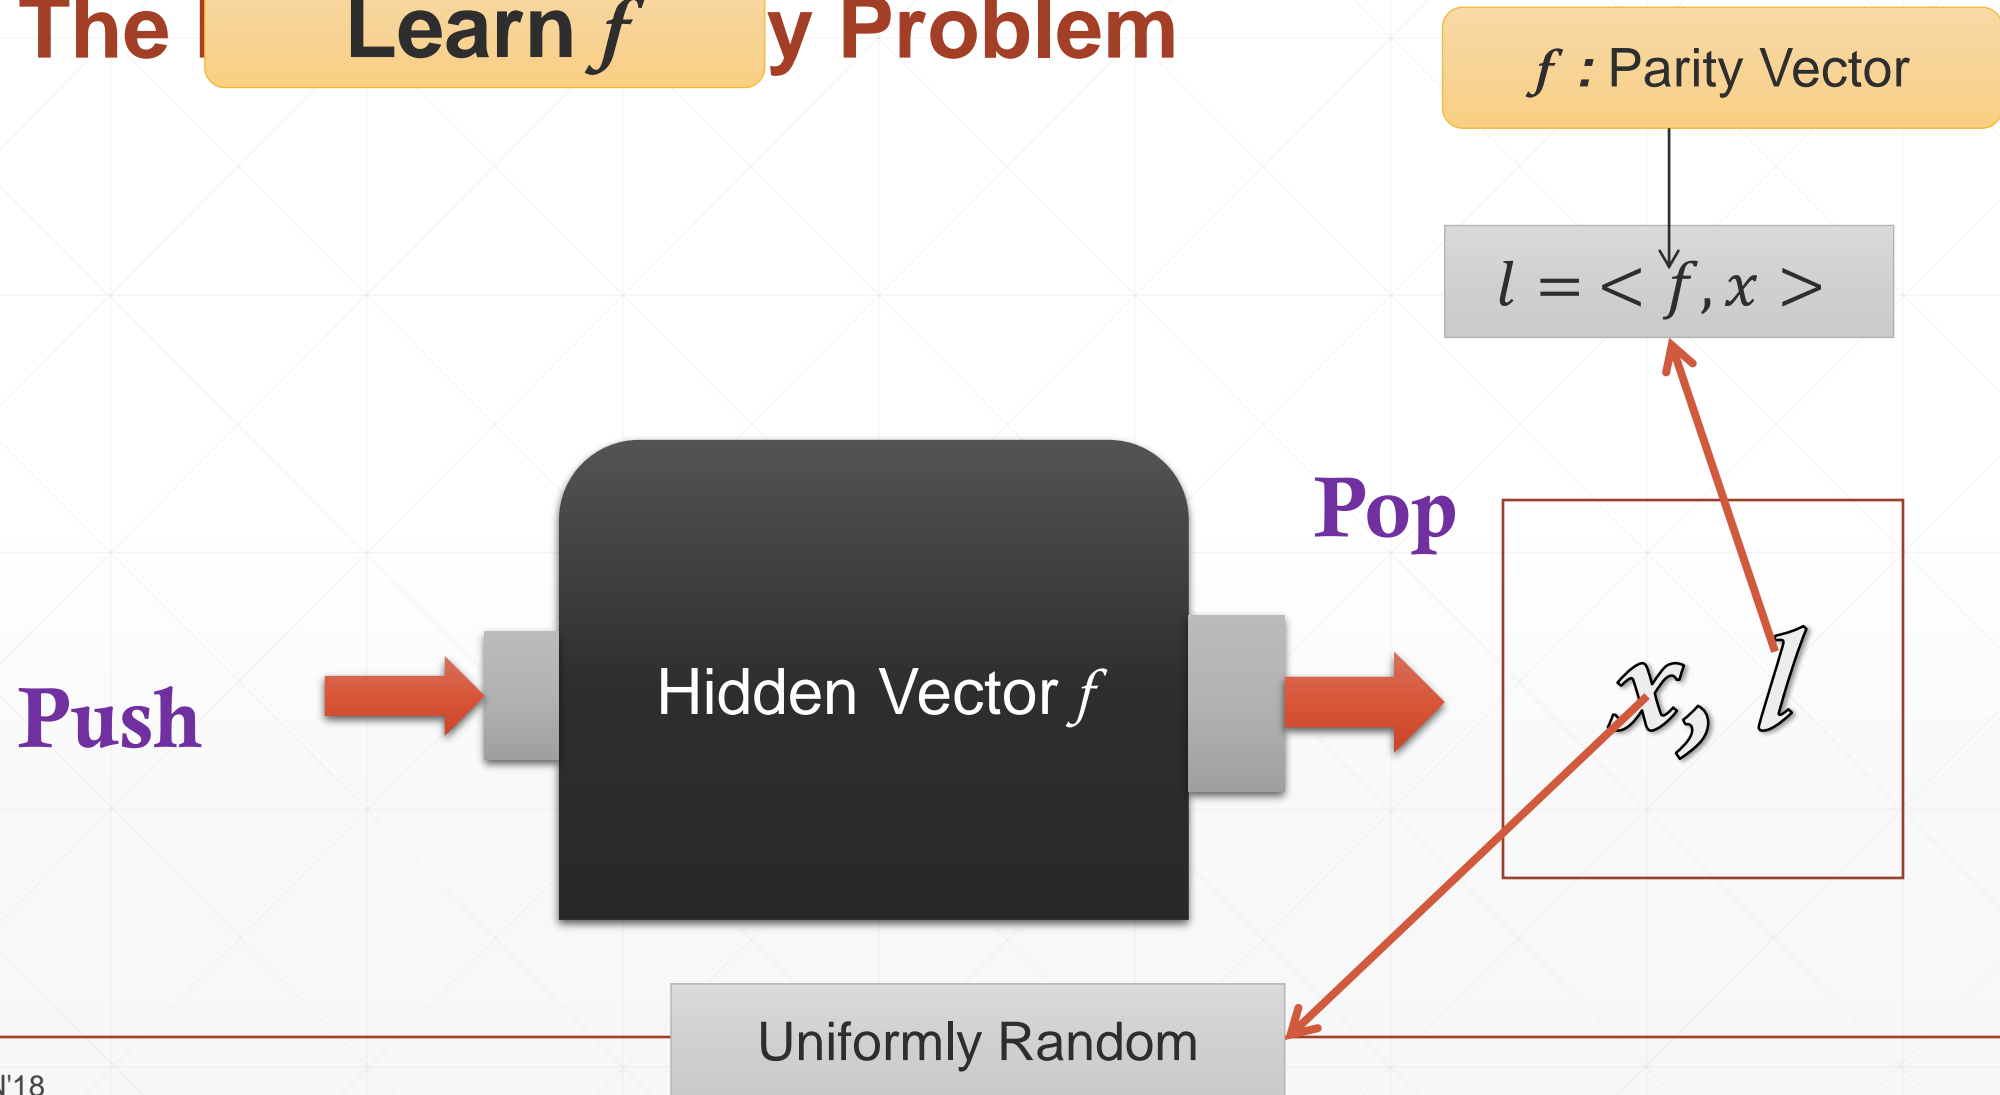
The Learning Parity Problem



The Learning Parity Problem



The Learn f Problem



Learn f minimizing

-
- Number of samples
 - Running time

The Learning Parity Problem

Gaussian elimination

- Uses $O(n)$ samples
- Runs in time $O(n^3)$

Learning Parity with Noise

- Same setup
- But the environment is noisy with **noise rate η**
- The labels are **flipped independently** with probability η
- Learn f **minimizing** the number of **samples** and **running time**

Learning Parity with Noise

- Can be solved using brute force algorithm that runs in time $\mathbf{O}(2^n)$
- Best known algorithm running time $\mathbf{O}(2^{\frac{n}{\log n}})$ by [*Blum, Kalai, Wasserman '03*]

Learning Parities

- Central Problem in Learning theory [*Feldman et al '09*]
- Coding theory
- Cryptography
- Lower bounds : **Open**

Learning k -Parity

- In this paper, study the variant problem in which f is k -sparse i.e. $|f| = k$ and $k \ll n$
- *First result* - Learning k -Parity **without noise**
- *Second result* - Learning k -Parity **with noise**

Learning k -Parity without noise

Learning k -Parity without noise

- Two approaches to learn f

Learning k -Parity without noise

- Two approaches to learn f

	Gaussian Elimination	Halving Algorithm
Sample Complexity	$O(n)$	$O\left(\log\binom{n}{k}\right)$
Time Complexity	$O(n^3)$	$O\left(n^{\frac{k}{2}}\right)$

Learning k -Parity without noise

- Current best *trade-offs* between sample complexity and running time given by (BGM) *Buhrman, Garcia-Soriano and Matsliah (2010)* in the stronger **Mistake bound model**.
- This paper - we **improve** the current *best trade offs*.

Online Mistake Bound model

- Oracle provides an **unlabeled** example x
- Learner **predicts** the label $\tilde{f}(x)$
- Oracle gives the **correct** label $f(x)$
- Learner can **update** its solution space depending upon the answer revealed

Each Round

The process repeats

Online Mistake Bound model

- **Mistake:** $f(x) \neq \tilde{f}(x)$
- *Learn f minimizing*
 - **Mistake bound**
 - Per round running **time**
- **Adversarial model, more difficult** than PAC model [*Blum'94*].

Our results for noiseless case

	This paper
Running time:	$e^{-k/4.01} \cdot \binom{t}{k} \cdot \tilde{O}\left(\left(\frac{kn}{t}\right)^2\right)$
Mistake bound:	$(1 + o(1)) \frac{kn}{t} + \log \binom{t}{k}$

Our results for noiseless case


	This paper	BGM'10
Running time:	$e^{-k/4.01} \cdot \binom{t}{k} \cdot \tilde{O}\left(\left(\frac{kn}{t}\right)^2\right)$	$O\left(\binom{t}{k} \left(\frac{kn}{t}\right)^2\right)$
Mistake bound:	$(1 + o(1)) \frac{kn}{t} + \log \binom{t}{k}$	$k \lceil \frac{n}{t} \rceil + \lceil \log \binom{t}{k} \rceil$

Idea behind the BGM algorithm

- Two approaches to learn f

	Gaussian Elimination	Halving Algorithm
Sample Complexity	$O(n)$	$O\left(\log\binom{n}{k}\right)$
Time Complexity	$O(n^3)$	$O\left(n^{\frac{k}{2}}\right)$


Gaussian Elimination



*Gaussian Elimination:
Geometrically*

Consider the vector space \mathbb{F}_2^n

Gaussian Elimination




Gaussian Elimination: Geometrically

Consider the vector space \mathbb{F}_2^n

An unlabeled example x specifies
a hyper plane

Gaussian Elimination



COCOON'18

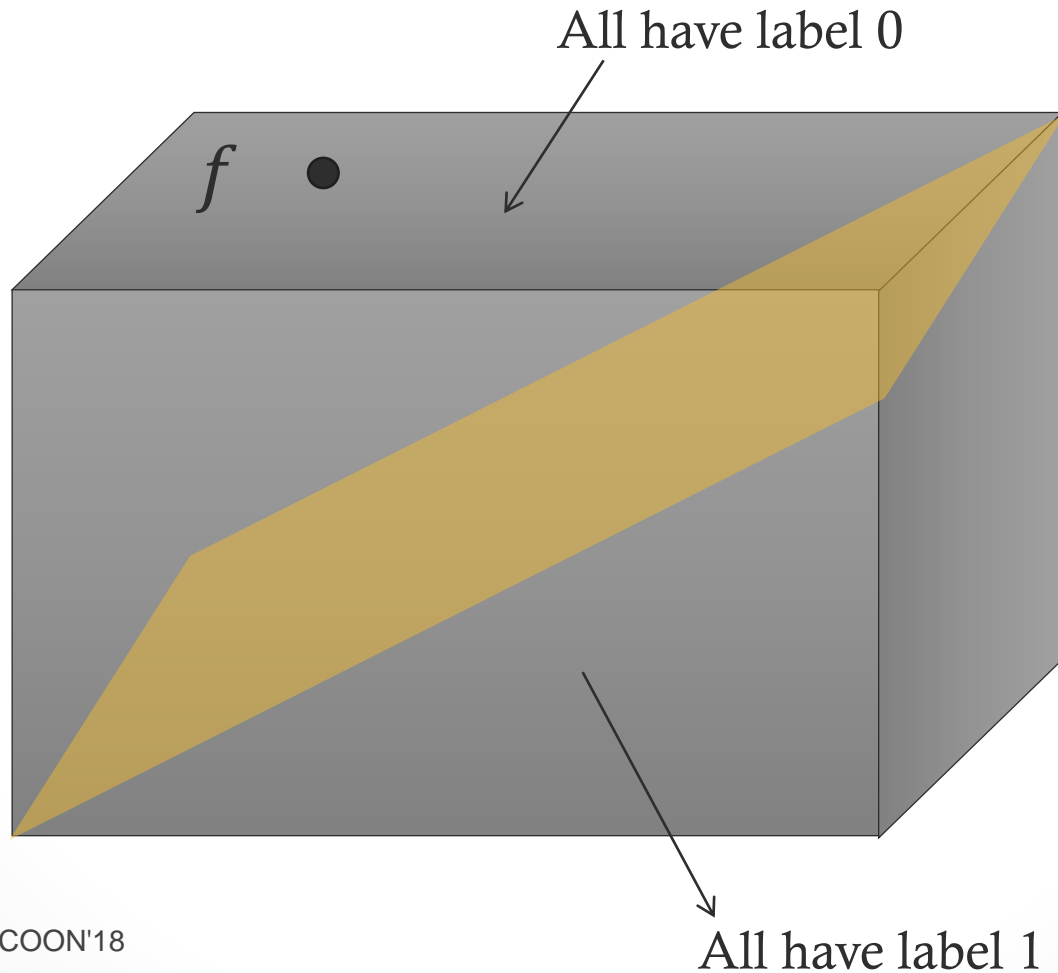
Gaussian Elimination: Geometrically

Consider the vector space \mathbb{F}_2^n

An unlabeled example x specifies
a hyper plane

The hyperplane divides the space
into two halves

Gaussian Elimination



COCOON'18

Gaussian Elimination: Geometrically


Consider the vector space \mathbb{F}_2^n

An unlabeled example x specifies
a hyper plane

The hyperplane divides the space
into two halves

Predict the majority – say,
predicted 1

Gaussian Elimination



COCOON'18

Gaussian Elimination: Geometrically

Consider the vector space \mathbb{F}_2^n


An unlabeled example x specifies
a hyper plane

The hyperplane divides the space
into two halves

Predict the majority – say,
predicted 1

If the true label is 0, throw half
space corresponding to 1

Gaussian Elimination



COCOON'18

Gaussian Elimination: Geometrically

Consider the vector space \mathbb{F}_2^n

An unlabeled example x specifies a hyper plane

The hyperplane divides the space into two halves

Predict the majority – say, predicted 1

If the true label is 0, throw half space corresponding to 1

Repeat with new example


Gaussian Elimination

- *Gaussian Elimination - Analysis*
 - Start with **one set** containing 2^n vectors as possible solutions
 - **Predict** the majority of the labels of the remaining solutions by performing the **intersection of the halfspace** with the remaining subset
 - At each mistake, throw **at least half** of the vectors

Gaussian Elimination

- *Gaussian Elimination - Analysis*
 - After, at most $\log_2 2^n = n$ mistakes, only 1 vector remains = hidden vector f
 - Computing intersection in time $O(n^3)$ by Gaussian elimination

Halving Algorithm




COCOON'18

$\binom{n}{k}$ many points

Halving Algorithm:
Geometrically

Consider all k -sparse vectors in
vector space \mathbb{F}_n^2

Halving Algorithm



COCOON'18

All have label 1


$\binom{n}{k}$ many points

Halving Algorithm:
Geometrically

Consider all k -sparse vectors in
vector space \mathbb{F}_n^2

An unlabeled example x specifies
a hyper plane

Halving Algorithm



COCOON'18

All have label 1

$\binom{n}{k}$ many points

Halving Algorithm: Geometrically

Consider all k -sparse vectors in vector space \mathbb{F}_n^2

An unlabeled example x specifies a hyper plane

Predict the majority – say, predicted 1

If the true label is 0, throw half space corresponding to 1

Repeat with new example

Halving Algorithm

- *Halving Algorithm* - Analysis

- Start with $\binom{n}{k}$ sets as possible solutions such that each k -sparse vector is in one subset.
- **Predict** the majority of the labels of the remaining solutions by performing the **intersection of the halfspace** with the remaining subset
- At each mistake, throw **at least half** of the vectors


Halving Algorithm

- *Halving Algorithm - Analysis*
 - After, at most $\log_2 \binom{n}{k} = k \log n$ mistakes, only 1 vector remains which is the hidden vector f
 - Computing the intersection with all the sets in time $\binom{n}{k}$

The BGM algorithm

- Tries to **balance** both the extremes
- Consider a set of **fewer subsets** such that each k -sparse vector is in at least one subset.
- Predict the label which has more **weighted majority** of subsets where weights are proportional to their sizes

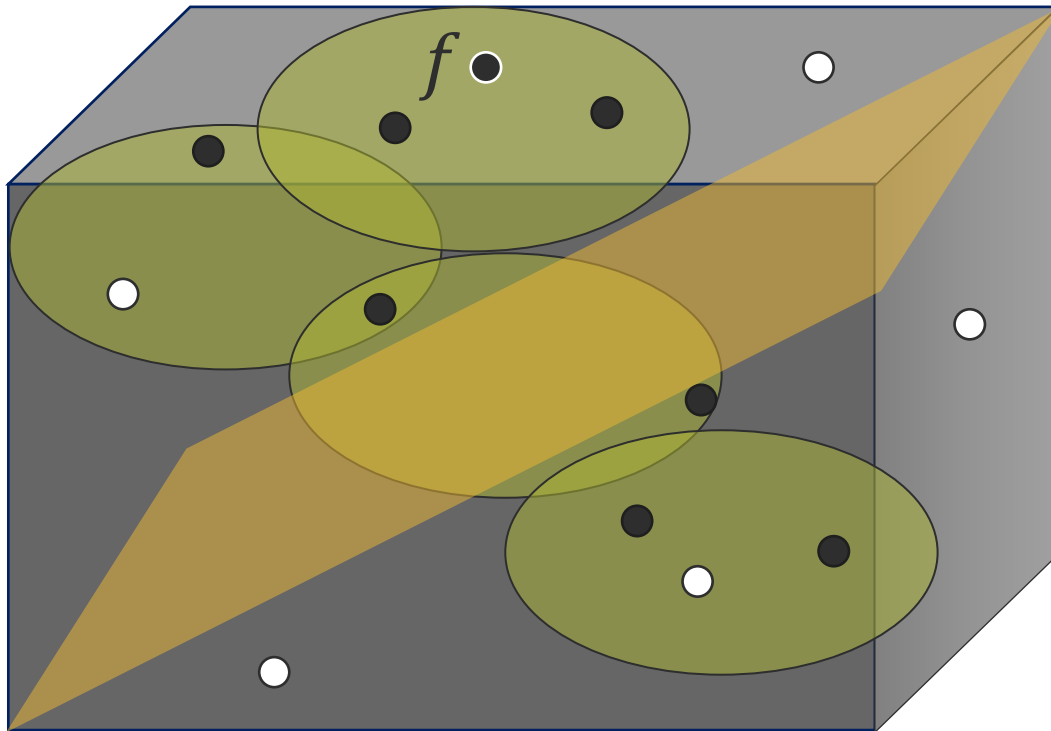
The BGM algorithm



The BGM algorithm:
Geometrically

Consider larger subsets of points
such that each k -sparse point is
present in some subset

The BGM algorithm




The BGM algorithm: Geometrically

Consider larger subsets of points such that each k -sparse point is present in some subset

An unlabeled example x specifies a hyper plane

The BGM algorithm




The BGM algorithm: Geometrically

Consider larger subsets of points such that each k -sparse point is present in some subset

An unlabeled example x specifies a hyper plane

Once true label is revealed, throw the irrelevant halfspace

The BGM algorithm



The BGM algorithm: Geometrically

Consider larger subsets of points such that each k -sparse point is present in some subset

An unlabeled example x specifies a hyper plane

Once true label is revealed, throw the irrelevant halfspace

Repeat with next example (or hyper plane)

The BGM algorithm formally

- *Initialization:*
 - Let $f \in \{0,1\}^n$ be the k -sparse parity vector
 - Let $e = \{e_1, e_2, \dots, e_n\}$ be the set of standard basis vectors of $\{0,1\}^n$
 - Arbitrarily partition e into $t \leq n$ parts $C = C_1, C_2, \dots, C_t$
 - Let $S := k$ -subsets of C
 - For each $s \in S$, let M_s be the span of $e_i \in s$. Thus, $|M_s| \leq 2^{k\lceil n/t \rceil}$

The BGM algorithm formally

- *On receiving an example $x \in \{0,1\}^n$:*
 - For each M_S , let $M_S^1 :=$ affine space of $\{M_S\} \cup \{x = 1\}$
 - Similarly, let $M_S^0 :=$ affine space of $\{M_S\} \cup \{x = 0\}$
 - Note that $|M_S^1| = 0, |M_S|$ or $\frac{|M_S|}{2}$
 - Predict $y \in \{0,1\}$ such that $\sum_{S \in \mathcal{S}} |M_S^y| \geq \sum_{S \in \mathcal{S}} |M_S^{1-y}|$

The BGM algorithm formally

- *On receiving answer $l \in \{0,1\}$:*
 - Update each $M_s = M_s^Z$

The BGM algorithm - analysis

- Mistakes:

- Total number of vectors in the beginning = $\binom{t}{k} 2^{k\lceil n/t \rceil}$
- At each mistake, throw away at least half of the vectors
- Number of mistake $\leq \log \left(\binom{t}{k} 2^{k\lceil n/t \rceil} \right) = k \left\lceil \frac{n}{t} \right\rceil + \log \binom{t}{k}$

- Running time:

- Per Round $O \left(\binom{t}{k} \left(\frac{kn}{t} \right)^2 \right)$

Our Algorithm

- **Idea** - Have **slightly bigger** subsets and pick **slightly fewer** of them
 - The setup is same as *BGM*, *but.....*
 - Partition e into $T = 1000t$ parts $C = C_1, C_2, \dots, C_T$
 - **Randomly** pick m , **1000k-sized** subsets of $[T]$
 - $m = O\left(\frac{\binom{1000t}{1000k}}{\binom{1000t-k}{1000k-k}}\right)$ ensures that with **non zero probability** each **k-sized** subset of $[T]$ is present in some S_i

Our Algorithm

- *Crucial claim:*

$$\frac{\binom{T}{1000k}}{\binom{T-k}{1000k-k}} \leq e^{-k/4.01} \binom{t}{k}$$

- The analysis is same as BGM
- Mistake bound = Mistake bound in BGM *up to constant terms*
- Running time = $e^{-k/4.01} \times$ BGM

Relating the results to PAC model

- Standard conversion techniques [*Angluin'88, Littlestone'89, Haussler'88*]
- Allow our result to get an improvement in the PAC model

Learning k -Parity with noise

Learning k -Parity with noise (k -LPN)

- Best known algorithm - *Grigorescu et al. (2011)*

Time: $\binom{n}{k/2}^{1+4\eta^2+o(1)}$

Samples: $\frac{k \log n}{(1-2\eta)^2} \cdot \omega(1)$

- When $\eta \rightarrow 1/2$, *G. Valiant (2012)* in time $n^{0.8k} \cdot \text{poly}\left(\frac{1}{1-2\eta}\right)$
- Barrier of $\binom{n}{k/2}$ in running time!

Breaking the Barrier...

- We show an algorithm that for **polynomially small** but non trivial range of noise rates, it is possible to break this barrier
- For example, when $\eta = \Theta\left(\frac{1}{n^{2/5}}\right)$ and $k = \sqrt{n}$, then our algorithm
 - Runs in time $\mathcal{O}\left(\binom{n}{k}^{\frac{1}{4}}\right)$
 - With $\mathcal{O}(k \cdot n^{3/8})$ samples

Breaking the Barrier... Algorithm

- Draw **sufficiently** many examples
- Guess a **set** of locations of a particular size (say $\frac{3\eta}{2}$) of the **mis-labelings** and correct them
- Use the **previous learner** from the noiseless setting to get a candidate parity vector
- **Repeat** this for every guess set of that size
- Draw few more examples and pick the **candidate parity vector** which **agrees** with the most number of newly drawn samples

Open Questions

- Noiseless case:
 - poly(n) algorithm with $O\left(\log\binom{n}{k}\right)$ samples – *attribute efficient learning of parities*
 - Improving our trade-offs
- Noisy case:
 - Lower bounds!
 - Better algorithms [*E.g., Karppa et.al. (2016)*]

Thank you

Attribute efficient learning k -Parity without noise

- Learn k -parity in **polynomial** time with only $\mathit{poly}(\log\binom{n}{k})$ samples
- Best known algorithm using $O\left(\log\binom{n}{k}\right)$ samples, in time $O\left(n^{\frac{k}{2}}\right)$ [Spielman]

Our Result for noisy case

THEOREM

Suppose $k(n) = n/f(n)$ for some function $f : \mathbb{N} \rightarrow \mathbb{N}$ for which $f(n) \ll n/\log \log n$, and suppose $\eta(n) = o\left(\frac{1}{((f(n))^\alpha \log n)}\right)$ for some $\alpha \in [1/2, 1)$. Then, for constant confidence parameter, there exists an algorithm for k -lpn with noise rate η with running time $e^{-k/4.01+o(k)} \cdot \binom{n}{k}^{1-\alpha} \cdot \text{poly}(n)$ and sample complexity $O(k(f(n))^\alpha)$.

For example, consider

$$k = \sqrt{n} \text{ and } \eta = \frac{1}{n^{2/5}} < \frac{1}{n^{3/8}},$$

$$\text{then } \alpha = \frac{3}{4}.$$

In this case, the running time would be $O\left(\binom{n}{k}^{\frac{1}{4}}\right)$ and the sample complexity would $O\left(k \left(\frac{n}{k}\right)^{\frac{3}{4}}\right)$.



Learning k -Parity without noise

- Information theoretically, $O\left(\log\binom{n}{k}\right)$ samples
- Running time is $O(n^k)$, improved to $O(n^{\frac{k}{2}})$
- **Open question** to get a **polynomial** algorithm with $O\left(\log\binom{n}{k}\right)$ samples

Online Mistake Bound model

- Different than the “black box” model (PAC)
- Learning proceeds in *rounds*
- Each *round*: “**Oracle**” *teaches* the “**Learner**”

Our results for noiseless case

- Let $k, t: \mathbb{N} \rightarrow \mathbb{N}$ be two functions such that $\log \log n \ll k(n) \ll t(n) \ll n$. Then for every $n \in \mathbb{N}$, there is an algorithm that learns k -parity in the mistake-bound model, with mistake bound at most $(1 + o(1)) \frac{kn}{t} + \log \binom{t}{k}$ and running time per round $e^{-k/4.01} \cdot \binom{t}{k} \cdot \tilde{O}\left(\left(\frac{kn}{t}\right)^2\right)$.

BGM'10

- Let $k, t: \mathbb{N} \rightarrow \mathbb{N}$ be two functions such that $k(n) \leq t(n) \leq n$. For every $n \in \mathbb{N}$, there is a deterministic algorithm that learns k -parity in the mistake-bound model, with mistake bound $k \lceil \frac{n}{t} \rceil + \lceil \log \binom{t}{k} \rceil$ and running time per round $O\left(\binom{t}{k} \left(\frac{kn}{t}\right)^2\right)$.

Our results for noiseless case

- Let $k, t: \mathbb{N} \rightarrow \mathbb{N}$ be two functions such that $\log \log n \ll k(n) \ll t(n) \ll n$. Then for every $n \in \mathbb{N}$, there is an algorithm that learns k -parity in the mistake-bound model, with mistake bound at most $(1 + o(1)) \frac{kn}{t} + \log \binom{t}{k}$ and running time per round $e^{-k/4.01} \cdot \binom{t}{k} \cdot \tilde{O}\left(\left(\frac{kn}{t}\right)^2\right)$.

BGM'10

- Let $k, t: \mathbb{N} \rightarrow \mathbb{N}$ be two functions such that $k(n) \leq t(n) \leq n$. For every $n \in \mathbb{N}$, there is a deterministic algorithm that learns k -parity in the mistake-bound model, with mistake bound $k \lceil \frac{n}{t} \rceil + \lceil \log \binom{t}{k} \rceil$ and running time per round $O\left(\binom{t}{k} \left(\frac{kn}{t}\right)^2\right)$.

Add a Slide Title - 2