A Polynomial Lower Bound for Monotonicity Testing of Boolean Functions

Joint work with Xi Chen and Rocco Servedio

 

Property Testing

Simplest question about a Boolean function: Does it have some property ***P***?

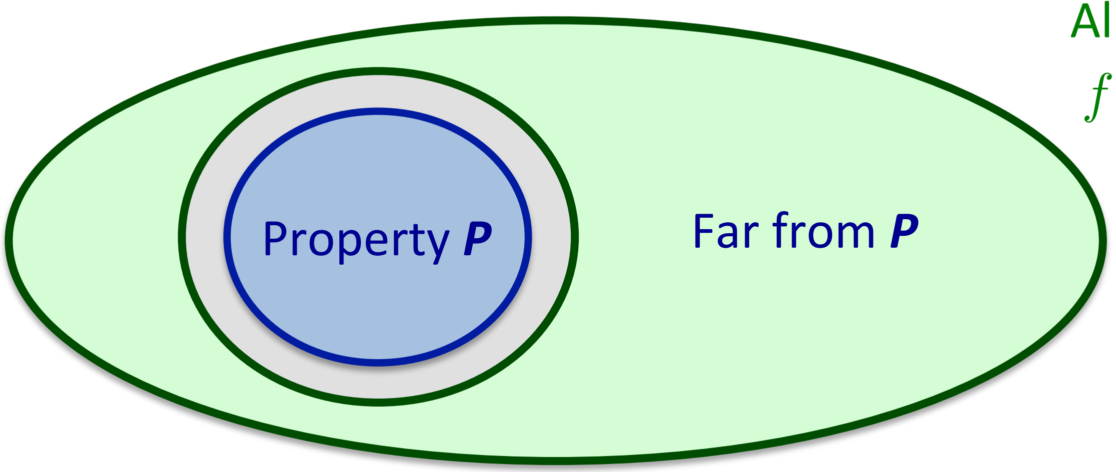

All Boolean functions
$f : \{0,1\}^n \to \{0,1\}$

- Query access to unknown *f* on any input *x*
- With as few queries as possible, decide if

Unreasonable:
Easy lower bound of $\Omega(2^n)$

f has Property ***P*** *vs.* ~~*f* does not have Property ***P***~~

f is **far** from having Property ***P***

Rules of the game

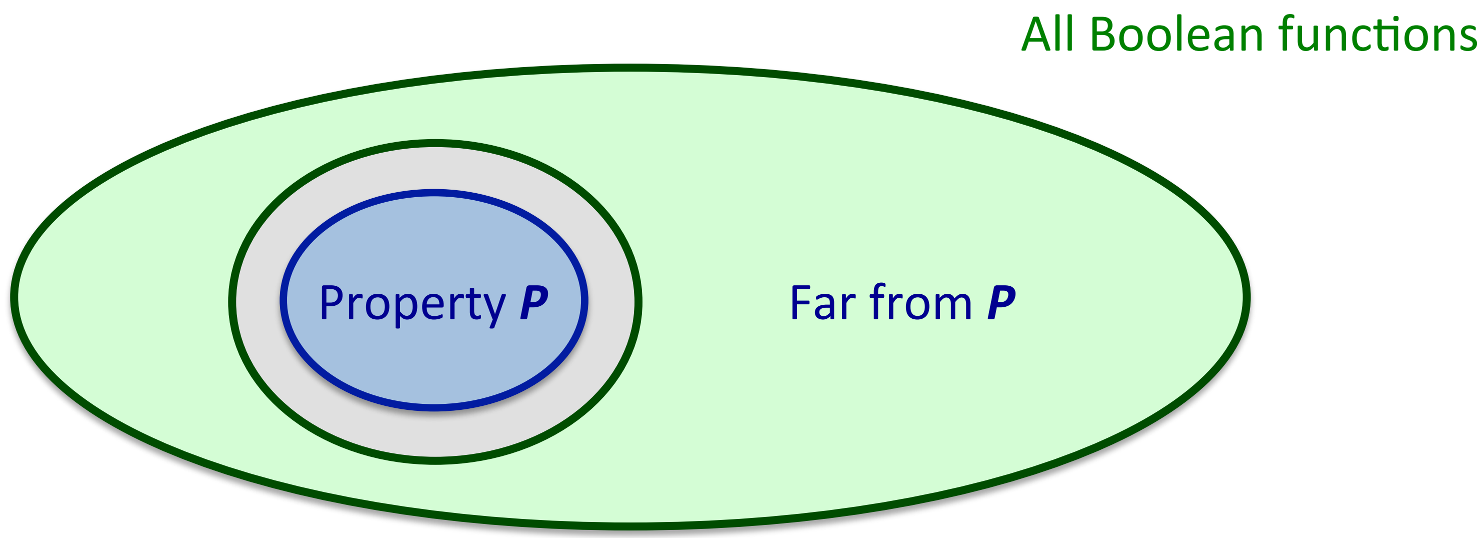

Query access to unknown f on any input x

1. If f has Property ***P***, accept w.p. > 2/3.
2. If f is ***ε*-far** from having Property ***P***, reject w.p. > 2/3.
3. Otherwise: doesn't matter what we do.

Super-efficient algorithms

Sublinear Space	**Sublinear Time**	Sublinear Measurement
Streaming, Sketching	**Property Testing**	Sparse Recovery, Compressed Sensing

Two recurring messages:

1. Many properties ***P*** testable with surprisingly few queries.
2. Rich connections with many other areas:
   - Learning theory
   - Hardness of approximation
   - Communication complexity
   - ...

This work: ***P*** **= monotonicity**

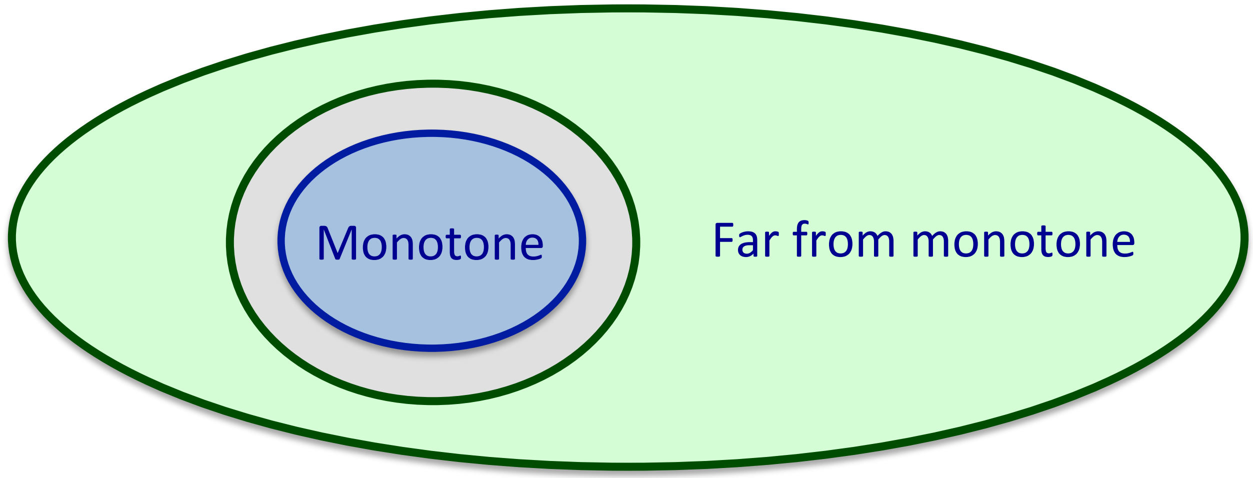

Well-studied problem:

[GGR98, GGL+98, DGL+99, FLN+02, HK08, BCGM12, RRS+12, BBM12, BRY13, CS13, …]

but still significant gaps in our understanding.

This talk

- The natural tester and its analysis [Goldreich *et al.* 1998]
- **Our main result:**

  A **polynomial lower bound** on query complexity
- **Our main technical ingredient:**

  Multidimensional Central Limit Theorems
- Generalizing our main result: **Testing monotonicity on *hypergrids***

A quick reminder

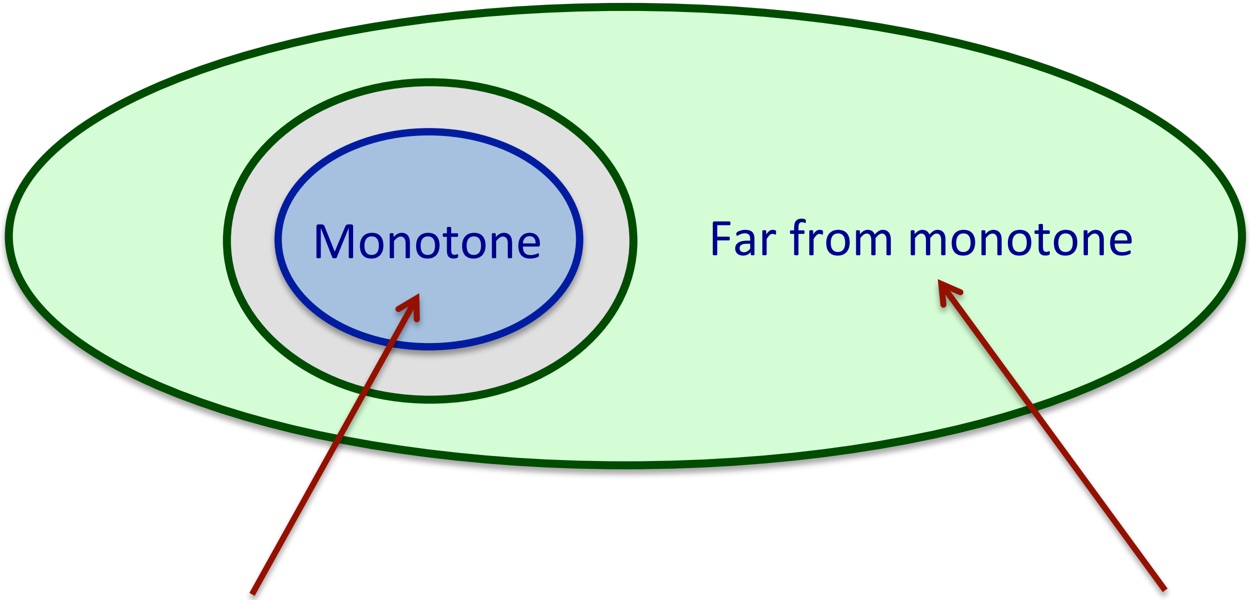

A monotone function is one that satisfies:

$$\forall \underbrace{x \preceq y}_{x_i \leq y_i \ \forall i \in [n]}, \ f(x) \leq f(y)$$

For all monotone functions *g:*

$$\Pr_{\boldsymbol{x} \in \{0,1\}^n} [f(\boldsymbol{x}) \neq g(\boldsymbol{x})] \geq \varepsilon$$

"Flipping an input bit from **0** to **1** cannot make *f* go from **1** to **0**"

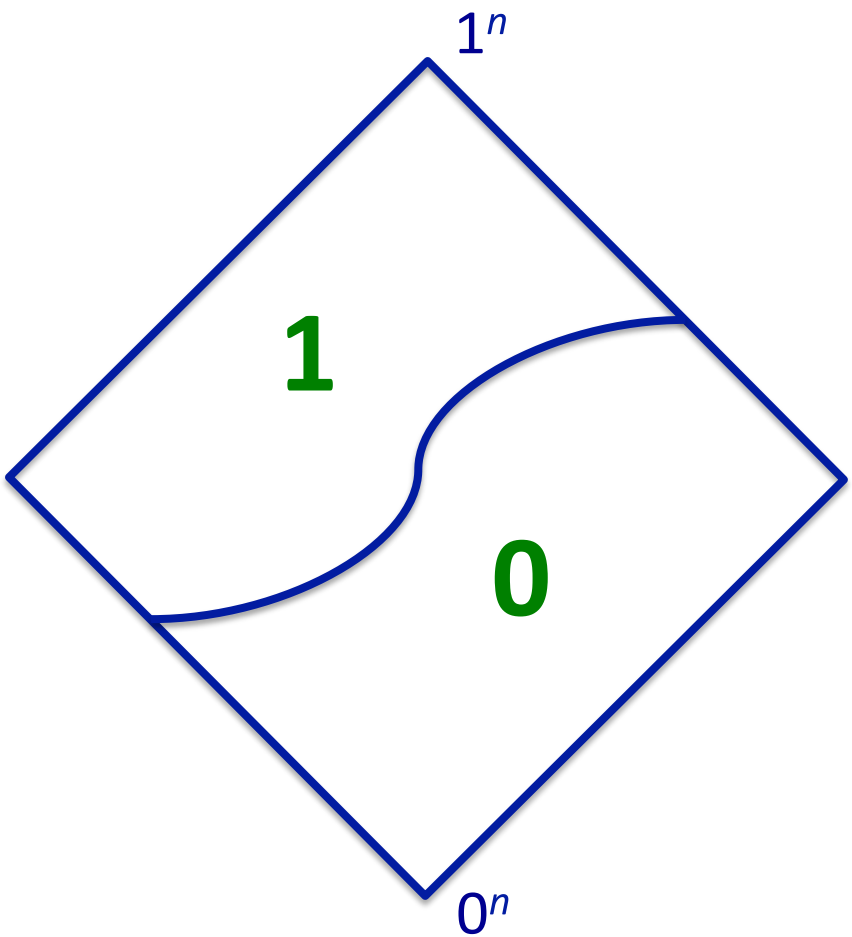

Monotone

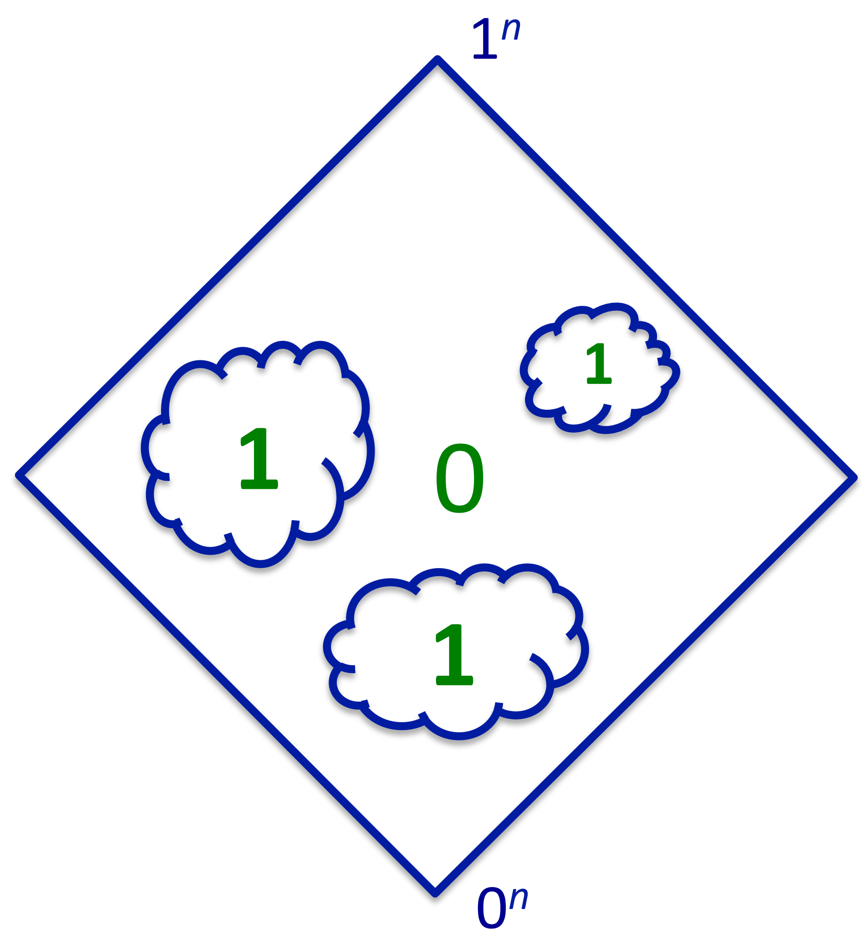

Far from monotone

First test that comes to mind

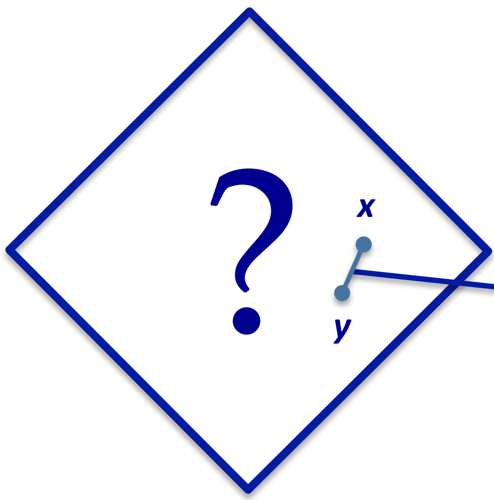

x = 1 1 1 1 0 0 1 **1** 1 0 1 0 0 1

Pick random edge

y = 1 1 1 1 0 0 1 **0** 1 0 1 0 0 1

$f(x) = 0$, $f(y) = 1$

Reject! 🙂

Call such an edge a "**violating edge**"

$f(x) = 1$, $f(y) = 1$ | $f(x) = 0$, $f(y) = 0$ | $f(x) = 1$, $f(y) = 0$

Repeat

An observation and a theorem

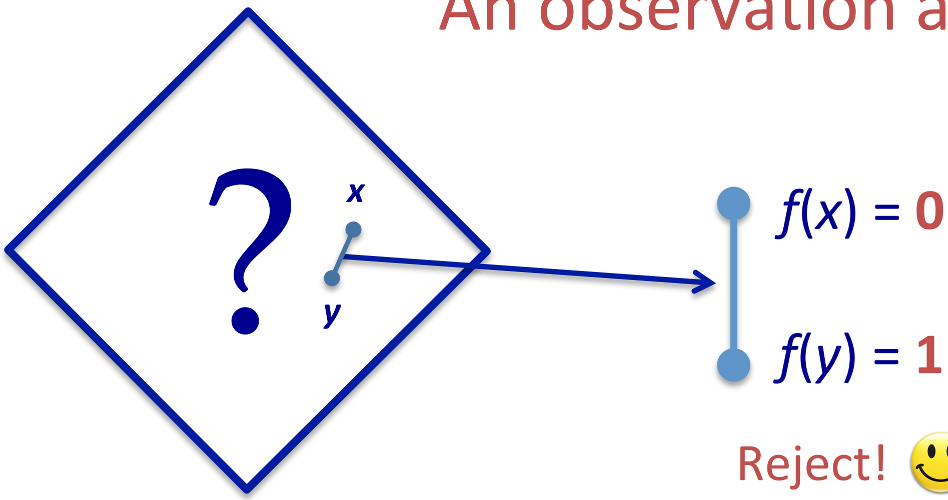

Call such an edge a “**violating edge**”

Reject!

Tester = Sample random edges, check for violations.

Simple observation: If f is monotone, tester never rejects.

Question: If f is ε-far from monotone, how likely to catch violating edge?

Theorem [Goldreich *et al*. 1998, 2000]

If f is ε-far from monotone, $\Omega(\varepsilon/n)$ fraction of edges are violations.
Therefore tester will reject within $O(n/\varepsilon)$ queries.

An exponential gap

- **Goldreich *et al.*** [FOCS 1998, SICOMP 2000]
  - Introduced problem, gave tester with $O(n)$ query complexity.
- **Fischer *et al.*** [STOC 2002]
  - Any non-adaptive tester must make $\Omega(\log n)$ queries.
  - Therefore, any adaptive tester must make $\Omega(\log \log n)$ queries.

[GGR98, DGL+99, HK08, BCGM12, RRS+12, BRY13, CS13, …]

- **Chakrabarty-Seshadhri** [STOC 2013]
  - $O(n^{7/8})$-query non-adaptive tester!

Theorem [Chen-Servedio-**T** 2014]

Any non-adaptive tester must make $\widetilde{\Omega}(n^{1/5})$ queries.

Therefore, any adaptive tester must make $\Omega(\log n)$ queries.

Exponential improvements over $\Omega(\log n)$ and $\Omega(\log \log n)$ lower bounds of Fischer *et al.* (2002)

Theorem [Chen-Servedio-**T** 2014]

There is a non-adaptive tester that makes $\widetilde{O}(n^{5/6})$ queries.

Polynomial improvement over $O(n^{7/8})$ upper bound of Chakrabarty and Sheshadhri (2013)

Yao's minimax principle

Lower bound against **randomized** algorithms ← implies — Tricky distribution over inputs to **deterministic** algorithms

Yao's principle in our setting

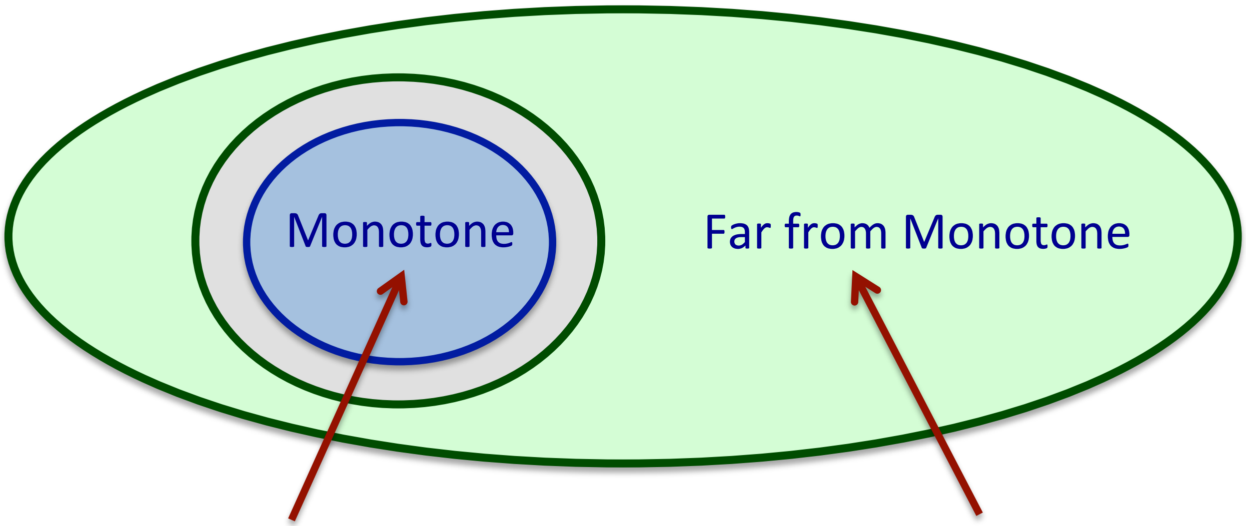

Distribution $\boldsymbol{D}_{\text{yes}}$ supported on monotone functions

Distribution $\boldsymbol{D}_{\text{no}}$ supported on far from monotone functions

Indistinguishability. For all T = deterministic tester that makes $o(n^{1/5})$ queries,

$$\left| \Pr_{\boldsymbol{f}_{yes} \sim \mathcal{D}_{yes}} \left[\mathcal{T} \text{ accepts } \boldsymbol{f}_{yes} \right] - \Pr_{\boldsymbol{f}_{no} \sim \mathcal{D}_{no}} \left[\mathcal{T} \text{ accepts } \boldsymbol{f}_{no} \right] \right| = o_n(1)$$

Our D_{yes} and D_{no} distributions

Both supported on *Linear Threshold Functions* (LTFs) over $\{-1,1\}^n$:

$$f(x) = \text{sign}(w_1 x_1 + \ldots + w_n x_n) \quad \vec{w} \in \mathbb{R}^n$$

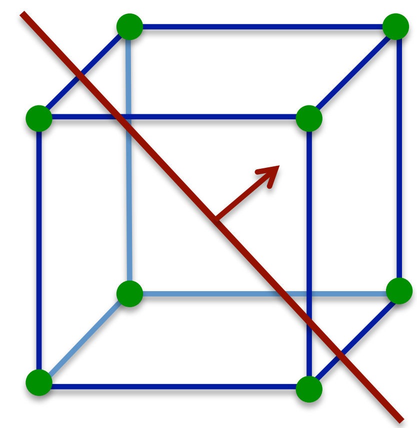

D_{yes} : $\boldsymbol{\sigma_i}$ = uniform from {1,3}

D_{no} : $\boldsymbol{\nu_i}$ = −1 with prob 0.1, 7/3 with prob 0.9

Verify: D_{yes} LTFs are monotone, D_{no} LTFs far from monotone w.h.p.

Main Structural Result: **Indistinguishability**

Any deterministic tester that makes few queries cannot tell D_{yes} from D_{no}

Key property: $\mathbb{E}[\boldsymbol{\sigma_i}] = \mathbb{E}[\boldsymbol{\nu_i}]$, $\mathbf{Var}[\boldsymbol{\sigma}_i] = \mathbf{Var}[\boldsymbol{\nu}_i]$.

Indistinguishability

q = 1 query

Claim. For all T = deterministic tester that makes ~~$q = o(n^{1/5})$ queries,~~

$$\left| \Pr_{\boldsymbol{f}_{yes} \sim \mathcal{D}_{yes}} \left[\mathcal{T} \textit{ accepts } \boldsymbol{f}_{yes} \right] - \Pr_{\boldsymbol{f}_{no} \sim \mathcal{D}_{no}} \left[\mathcal{T} \textit{ accepts } \boldsymbol{f}_{no} \right] \right| = o_n(1)$$

Non-trivial proof of a triviality

Claim. Let T = deterministic tester that makes 1 query **z**. Then:

$$\underbrace{\left| \Pr_{\boldsymbol{f}_{yes} \sim \mathcal{D}_{yes}} \left[\mathcal{T} \textit{ accepts } \boldsymbol{f}_{yes} \right] - \Pr_{\boldsymbol{f}_{no} \sim \mathcal{D}_{no}} \left[\mathcal{T} \textit{ accepts } \boldsymbol{f}_{no} \right] \right|}_{(*)\ \le\ d_{\mathrm{TV}}(\mathbf{R}_{yes}, \mathbf{R}_{no})} = o_n(1)$$

Tester sees:

$$\mathbf{R}_{yes} = \mathrm{sign}(\boldsymbol{\sigma_1} z_1 + \cdots + \boldsymbol{\sigma_n} z_n) \quad \textit{vs.} \quad \mathbf{R}_{no} = \mathrm{sign}(\boldsymbol{\nu_1} z_1 + \cdots + \boldsymbol{\nu_n} z_n)$$

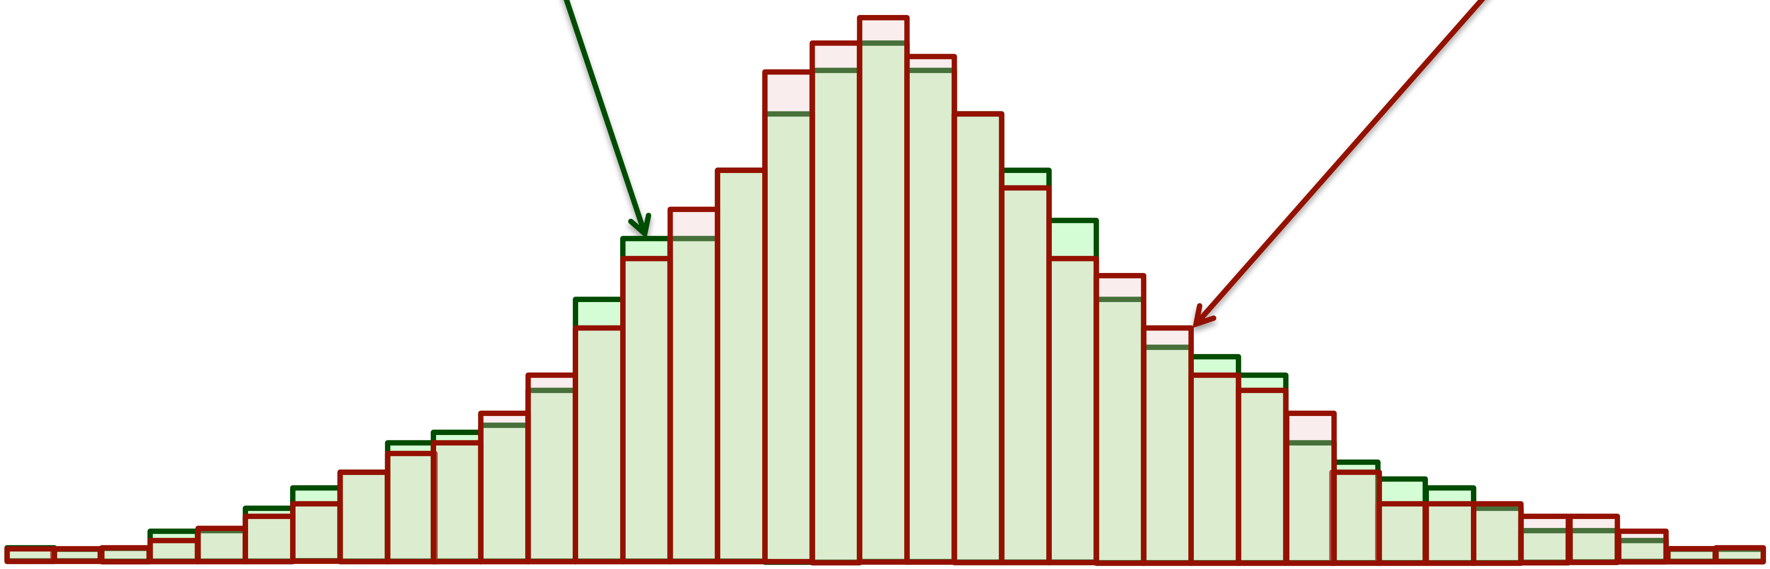

Central Limit Theorems. Sum of many independent "reasonable" random variables converges to Gaussian of same mean and variance.

Main analytic tool (Baby version):

Berry–Esséen CLT. *Let* $\mathbf{S} = \mathbf{X}_1 + \cdots + \mathbf{X}_n$ *where* $\mathbf{X}_1, \ldots, \mathbf{X}_n$ *are independent real-valued random variables satisfying* $|\mathbf{X}_j - \mathbf{E}[\mathbf{X}_j]| \leq \tau$ *with probability* 1 *for all* $j \in [n]$. *Let* $\mathcal{G}$ *be a Gaussian with mean* $\mathbf{E}[\mathbf{S}]$ *and variance* $\mathbf{Var}[\mathbf{S}]$. *Then for all* $\theta \in \mathbb{R}$,

$$\underbrace{\left|\mathbf{Pr}[\mathbf{S} \leq \theta] - \mathbf{Pr}[\mathcal{G} \leq \theta]\right|} \leq \frac{O(\tau)}{\mathbf{Var}[\mathbf{S}]^{1/2}}.$$

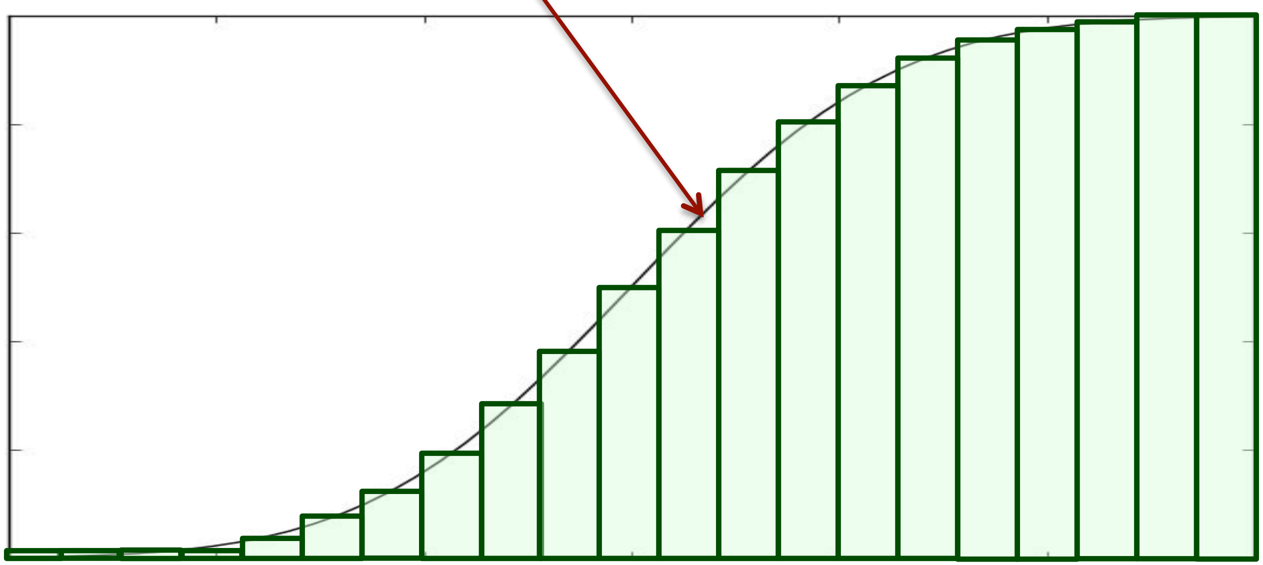

Goal: Upper bound $d_{\mathrm{TV}}(\mathrm{sign}(\mathbf{S}_{yes}), \mathrm{sign}(\mathbf{S}_{no}))$

$$\mathbf{S}_{yes} = \boldsymbol{\sigma_1} z_1 + \cdots + \boldsymbol{\sigma_n} z_n \qquad \mathbf{S}_{no} = \boldsymbol{\nu_1} z_1 + \cdots + \boldsymbol{\nu_n} z_n$$

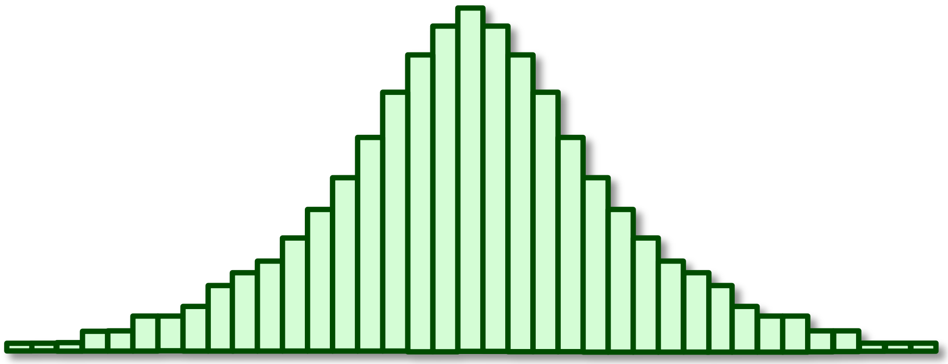

$\mathbf{S}_{yes} \to \mathcal{G}_1$

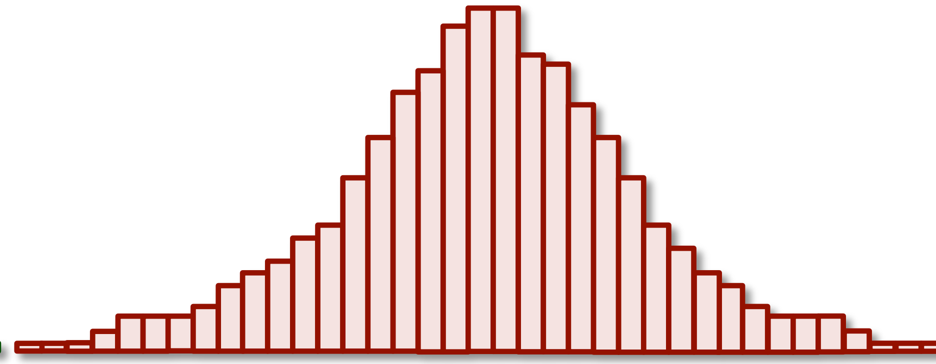

$\mathbf{S}_{no} \to \mathcal{G}_2$

Recall key property:

$$\left.\begin{aligned} \mathbb{E}[\boldsymbol{\sigma_i}] &= \mathbb{E}[\boldsymbol{\nu_i}] \\ \mathbf{Var}[\boldsymbol{\sigma}_i] &= \mathbf{Var}[\boldsymbol{\nu}_i] \end{aligned}\right\} \longrightarrow \begin{aligned} \mathbb{E}[\mathbf{S}_{yes}] &= \mathbb{E}[\mathbf{S}_{no}] \\ \mathbf{Var}[\mathbf{S}_{yes}] &= \mathbf{Var}[\mathbf{S}_{no}] \\ \mathcal{G}_1 &= \mathcal{G}_2 \end{aligned}$$

We just proved:

Claim. Let T = deterministic tester that makes **1 query**. Then:

$$\left| \Pr_{\boldsymbol{f}_{yes} \sim \mathcal{D}_{yes}} \left[\mathcal{T} \ accepts \ \boldsymbol{f}_{yes} \right] - \Pr_{\boldsymbol{f}_{no} \sim \mathcal{D}_{no}} \left[\mathcal{T} \ accepts \ \boldsymbol{f}_{no} \right] \right| = O(n^{-1/2})$$

Plenty of room to spare!

Would be happy with < 0.1

q queries instead of 1

Main analytic tool (Grown-up version):

Multidimensional CLTs. Sum of many independent "reasonable" ***q*-dimensional** random variables converge to ***q*-dimensional** Gaussian of same mean and variance.

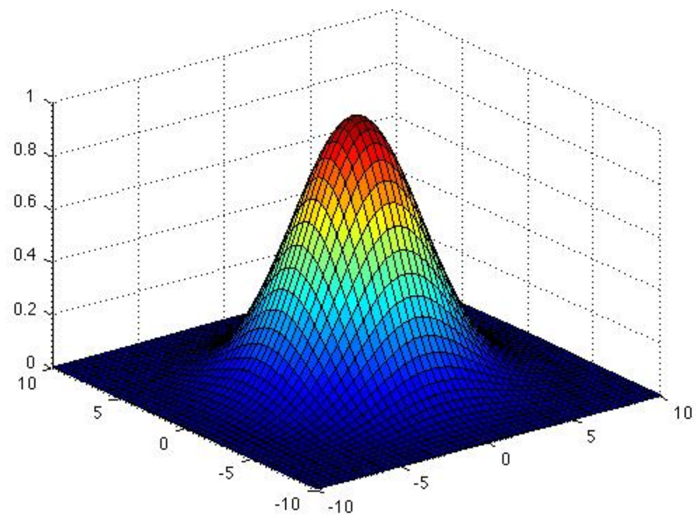

Multidimensional CLTs

[Mossel 08, Gopalan-O'Donnell-Wu-Zuckerman 10]

$\longrightarrow$ $\Omega(n^{1/12})$

Lower bound

[Valiant-Valiant 11]

Main technical work:

Adapting multidimensional CLT for Earth Mover Distance to get

$$\widetilde{\Omega}(n^{1/5})$$

Let's prove the real thing:

Claim. For all T = deterministic tester that makes $q = o(n^{1/5})$ queries,

$$\left| \Pr_{\boldsymbol{f}_{yes} \sim \mathcal{D}_{yes}} \left[\mathcal{T} \ accepts \ \boldsymbol{f}_{yes} \right] - \Pr_{\boldsymbol{f}_{no} \sim \mathcal{D}_{no}} \left[\mathcal{T} \ accepts \ \boldsymbol{f}_{no} \right] \right| = o_n(1)$$

Setting things up

Arrange the *q* queries of tester *T* in a *q* x *n* matrix $\boldsymbol{Q} \in \{-1,1\}^{q \times n}$

$$
\begin{array}{c|c|}
\cline{2-2}
Q_1 & \\
Q_2 & \\
Q_3 & Q \in \{-1,1\}^{q \times n} \\
\vdots & Q_i = i\text{-th query string} \\
Q_q & \\
\cline{2-2}
\end{array}
$$

Recall: Tester's goal is to distinguish

$\boldsymbol{f}_{yes} = \mathrm{sign}(\boldsymbol{\sigma} x)$ versus $\boldsymbol{f}_{no} = \mathrm{sign}(\boldsymbol{\nu} x)$

What does the tester see?

Tester sees:

$\boldsymbol{f}_{yes} \leftarrow \mathcal{D}_{yes}$

$\boldsymbol{f}_{yes} = \mathrm{sign}(\boldsymbol{\sigma} x)$

$Q \in \{-1,1\}^{q\times n}$

$\in \mathbb{R}^q$ $\in \{-1,1\}^q$

$Q\boldsymbol{\sigma}$ $\mathbf{R}_{yes} = \mathrm{sign}(Q\boldsymbol{\sigma})$

(coordinate-wise sign)

$\boldsymbol{f}_{no} \leftarrow \mathcal{D}_{no}$

$\boldsymbol{f}_{no} = \mathrm{sign}(\boldsymbol{\nu} x)$

$\in \mathbb{R}^q$ $\in \{-1,1\}^q$

$Q\boldsymbol{\nu}$ $\mathbf{R}_{no} = \mathrm{sign}(Q\boldsymbol{\nu})$

Goal: Upper bound $d_{\mathrm{TV}}(\mathbf{R}_{yes}, \mathbf{R}_{no}) = d_{\mathrm{TV}}(\mathrm{sign}(Q\boldsymbol{\sigma}), \mathrm{sign}(Q\boldsymbol{\nu}))$

Goal: Upper bound $d_{\mathrm{TV}}(\underbrace{\mathbf{R}_{yes}, \mathbf{R}_{no}}_{\{\pm 1\}^q}) = d_{\mathrm{TV}}(\mathrm{sign}(Q\boldsymbol{\sigma}), \mathrm{sign}(Q\boldsymbol{\nu}))$

$\mathbb{R}^q$

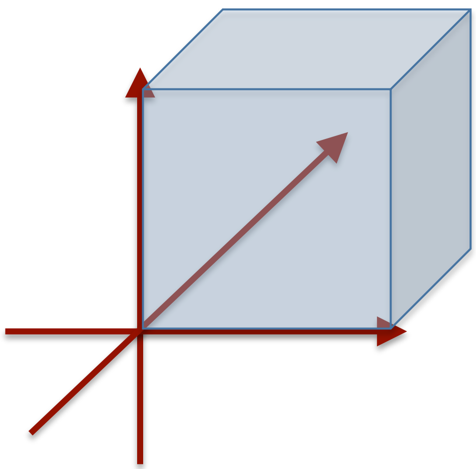

$\mathbf{R}_{yes} \equiv$ orthant of $\mathbb{R}^q$ that $Q\boldsymbol{\sigma}$ falls in

$\mathbf{R}_{no} \equiv$ orthant of $\mathbb{R}^q$ that $Q\boldsymbol{\nu}$ falls in

Random variables supported on 2^q orthants of R^q

$$d_{\mathrm{TV}}(\mathbf{R}_{yes}, \mathbf{R}_{no}) = \sum_{\substack{2^q \text{ orthants} \\ O_i \text{ of } \mathbb{R}^q}} \left| \Pr[Q\boldsymbol{\sigma} \in O_i] - \Pr[Q\boldsymbol{\nu} \in O_i] \right|$$

$$= \max_{\mathcal{O} \subseteq \mathbb{R}^q} \left| \Pr[Q\boldsymbol{\sigma} \in \mathcal{O}] - \Pr[Q\boldsymbol{\nu} \in \mathcal{O}] \right|$$

union of orthants

“roughly equal weight on any union of orthants”

$$d_{\mathrm{TV}}(\mathbf{R}_{yes}, \mathbf{R}_{no}) = \max_{\mathcal{O}\subseteq\mathbb{R}^q} \Big|\Pr[Q\boldsymbol{\sigma}\in\mathcal{O}] - \Pr[Q\boldsymbol{\nu}\in\mathcal{O}]\Big|$$

$$\begin{array}{c|cccc|}
 & Q_{\star 1} & Q_{\star 2} & \cdots & Q_{\star n} \\
\hline
Q_1 & & & & \\
Q_2 & & & & \\
\vdots & & Q \in \{-1,1\}^{q\times n} & & \\
Q_q & & & & \\
\hline
\end{array}$$

Fixed

$$\underbrace{Q\boldsymbol{\sigma} = \sum_{i=1}^{n} Q_{\star i}\boldsymbol{\sigma_i}}_{\text{sum of } n \text{ independent vectors in } \mathbb{R}^q}$$

from product distribution over R^n

Ditto: $Q\boldsymbol{\nu} = \sum_{i=1}^{n} Q_{\star i}\boldsymbol{\nu_i}$

$$\left.\begin{aligned}\mathbb{E}[\boldsymbol{\sigma_i}] &= \mathbb{E}[\boldsymbol{\nu_i}]\\ \mathbf{Var}[\boldsymbol{\sigma}_i] &= \mathbf{Var}[\boldsymbol{\nu}_i]\end{aligned}\right\} \longrightarrow \begin{aligned}\mathbb{E}[Q\boldsymbol{\sigma}] &= \mathbb{E}[Q\boldsymbol{\nu}]\\ \mathbf{Cov}[Q\boldsymbol{\sigma}] &= \mathbf{Cov}[Q\boldsymbol{\nu}]\end{aligned}$$

The final setup

Goal: Two sums of *n* independent vectors in $\mathbb{R}^q$ are "close"

$$Q\boldsymbol{\sigma} = \sum_{i=1}^{n} Q_{\star i}\boldsymbol{\sigma_i} \qquad Q\boldsymbol{\nu} = \sum_{i=1}^{n} Q_{\star i}\boldsymbol{\nu_i}$$

where closeness = roughly equal weight on any union of orthants.

Furthermore, since

$$\mathbb{E}[Q\boldsymbol{\sigma}] = \mathbb{E}[Q\boldsymbol{\nu}]$$

$$\mathbf{Cov}[Q\boldsymbol{\sigma}] = \mathbf{Cov}[Q\boldsymbol{\nu}]$$

suffices to show each are close to *q*-dimensional Gaussian with matching mean and covariance matrix.

Valiant-Valiant Multidimensional CLT

Sum of many independent “reasonable” ***q*-dimensional** random variables is close to ***q*-dimensional** Gaussian of same mean and variance.

with respect to **Earth Mover Distance**:

Minimum amount of work necessary to “get one PDF to look like the other”, where work := mass x distance

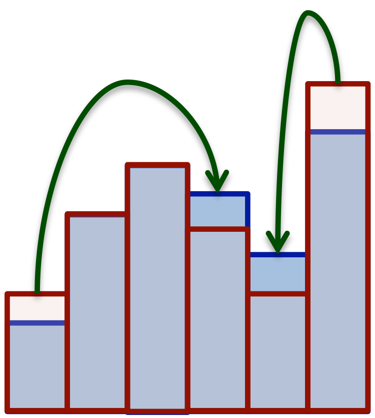

Key technical lemma:

Closeness in EMD $\longrightarrow$ roughly equal weight on any union of orthants

Valiant-Valiant Multidimensional CLT

Let $\mathbf{S} = \mathbf{X}_1 + \cdots + \mathbf{X}_n$, where the $\mathbf{X}_j$'s are independent $\mathbb{R}^q$-valued random variables satisfying $\|\mathbf{X}_j - \mathbf{E}[\mathbf{X}_j]\|_2 \le \tau$ with probability 1 for all $j \in [n]$. Let $\mathcal{G}$ be the q-dimensional Gaussian with the same mean and covariance matrix as $\mathbf{S}$. Then

$$d_{\mathrm{EMD}}(\mathbf{S}, \mathcal{G}) \le O(\tau q \log n).$$

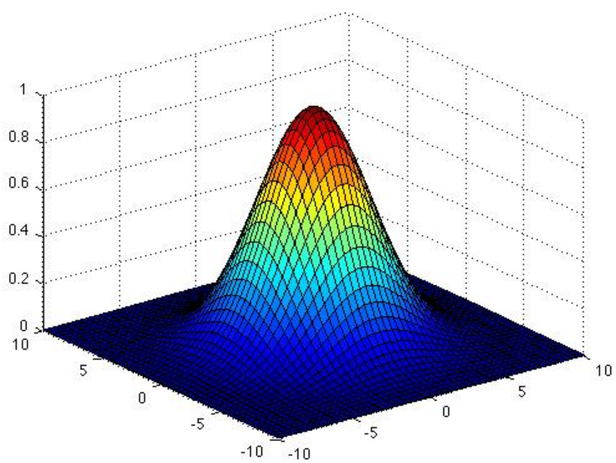

Key technical lemma

Closeness in EMD

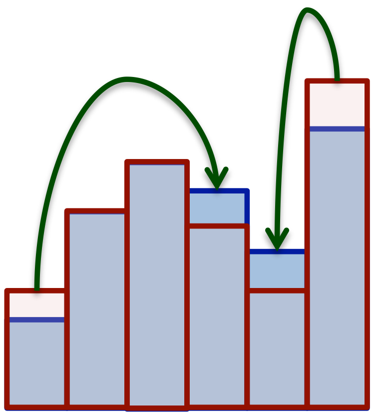

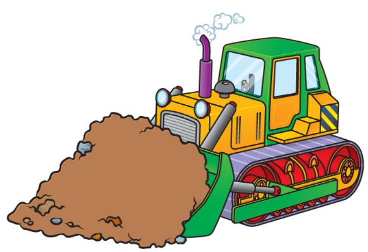

Roughly equal weight on any union of orthants

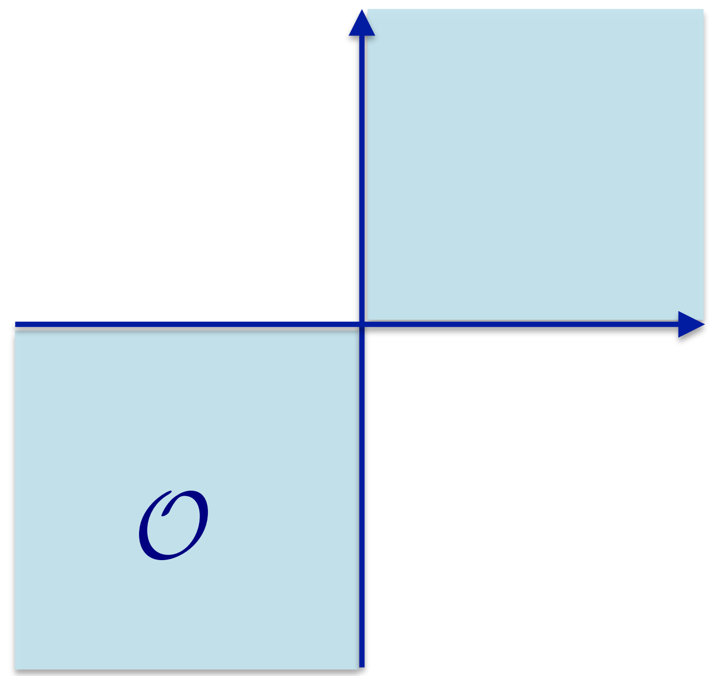

$$d_{\mathrm{EMD}}(\mathbf{S}, \mathcal{G}) \text{ small} \implies |\Pr[\mathbf{S} \in \mathcal{O}] - \Pr[\mathcal{G} \in \mathcal{O}]| \text{ small}$$

$$d_{\mathrm{EMD}}(\mathbf{S},\mathcal{G}) \text{ small} \implies |\Pr[\mathbf{S}\in\mathcal{O}] - \Pr[\mathcal{G}\in\mathcal{O}]| \text{ small}$$

for all unions of orthants $\mathcal{O}$

Let's consider the contrapositive:

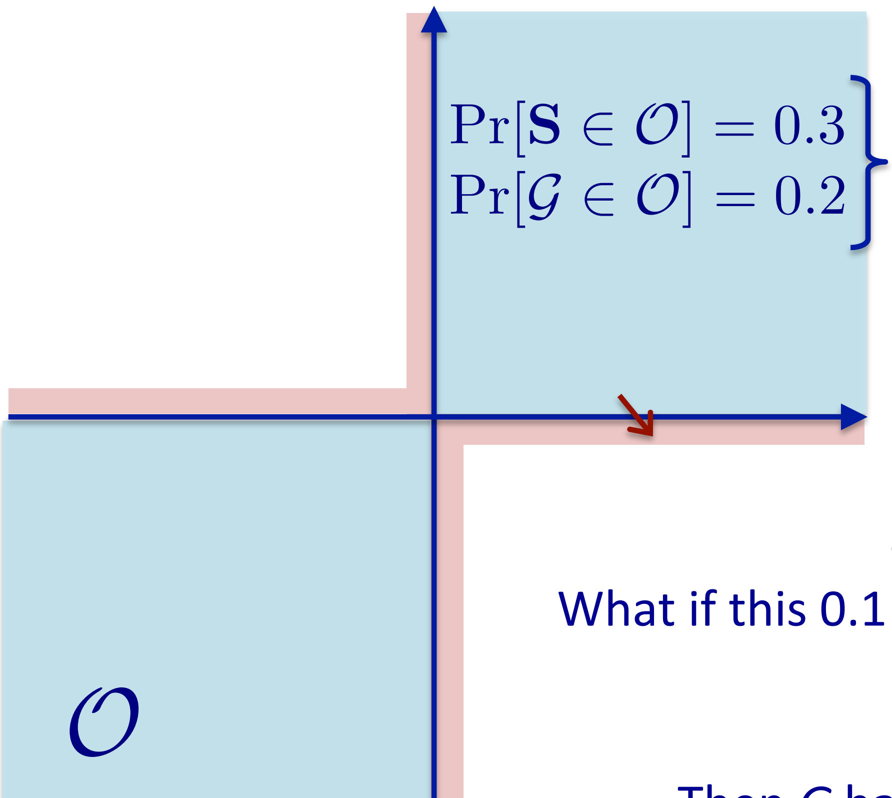

If we want to make **S** look like *G*, ≥ 0.1 unit of mass must be moved out of $\mathcal{O}$

Recall:
work = mass x distance

Slight technical obstacle:

What if this 0.1 unit of mass only moves a *tiny* distance?

Solution:

Then *G* has ≥ 0.1 weight on red boundary.

Not possible thanks to anti-concentration of *G*

In slightly more detail

$\Pr[\mathbf{S} \in \mathcal{O}] - \Pr[\mathcal{G} \in \mathcal{O}] = \Delta$, has to be moved out of $\mathcal{O}$

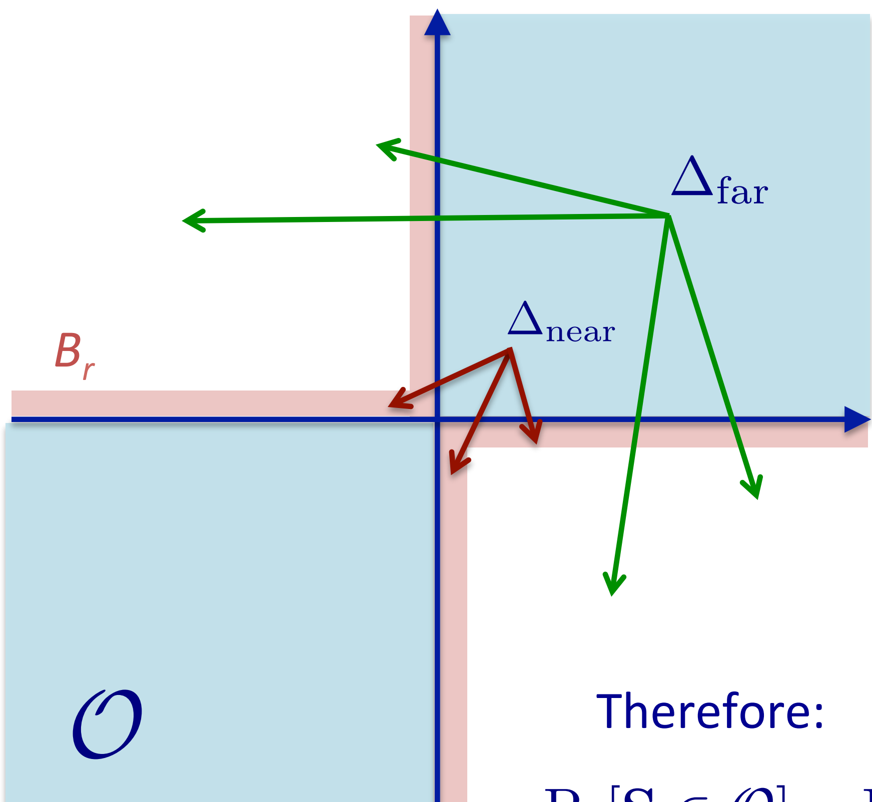

For all $r > 0$,
define B_r := radius r boundary around $\mathcal{O}$

$$\Delta = \Delta_{\text{near}} + \Delta_{\text{far}}$$

$$\Delta_{near} \leq \Pr[\mathcal{G} \in B_r]$$

$$r \cdot \Delta_{far} \leq d_{\text{EMD}}(\mathbf{S}, \mathcal{G})$$

Therefore:

$$\Pr[\mathbf{S} \in \mathcal{O}] - \Pr[\mathcal{G} \in \mathcal{O}] \leq \frac{d_{\text{EMD}}(\mathbf{S}, \mathcal{G})}{r} + \Pr[\mathcal{G} \in B_r]$$

We just proved

Let $\mathbf{S} = \mathbf{X}_1 + \cdots + \mathbf{X}_n$, where the $\mathbf{X}_j$'s are independent $\mathbb{R}^q$-valued random variables. Let $\mathcal{G}$ be the q-dimensional Gaussian with the same mean and covariance matrix as $\mathbf{S}$. Then for all unions of orthants $\mathcal{O} \subseteq \mathbb{R}^q$ and for all $r > 0$,

$$\left| \Pr[\mathbf{S} \in \mathcal{O}] - \Pr[\mathcal{G} \in \mathcal{O}] \right| \leq \underbrace{\frac{d_{\mathrm{EMD}}(\mathbf{S}, \mathcal{G})}{r}}_{\text{Valiant-Valiant CLT}} + \underbrace{\Pr[\mathcal{G} \in B_r]}_{\text{Gaussian anti-concentration}}$$

Recap

Indistinguishability. For all T = deterministic tester that makes $q = o(n^{1/5})$ queries,

$$\left| \Pr_{\boldsymbol{f}_{yes} \sim \mathcal{D}_{yes}} \left[\mathcal{T} \ accepts \ \boldsymbol{f}_{yes} \right] - \Pr_{\boldsymbol{f}_{no} \sim \mathcal{D}_{no}} \left[\mathcal{T} \ accepts \ \boldsymbol{f}_{no} \right] \right| = o_n(1)$$

$$\Downarrow$$

Theorem

Any non-adaptive tester must make $\widetilde{\Omega}(n^{1/5})$ queries.

Therefore, any adaptive tester must make $\Omega(\log n)$ queries.

Testing monotonicity of Boolean functions over general hypergrid domains

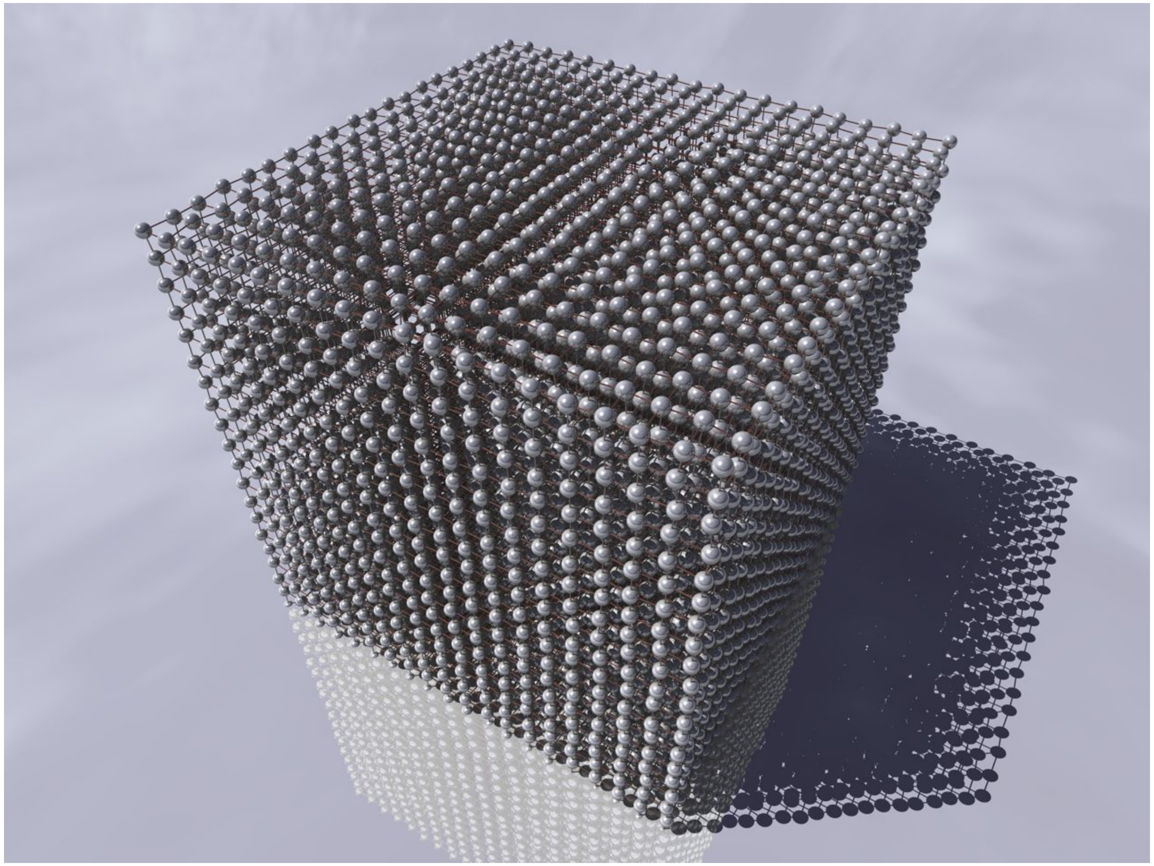

Boolean functions over hypergrids

$$F : \{1, 2, \ldots, m\}^n \to \{0, 1\}$$

Testing monotonicity of Boolean functions over hypergrids

All hypergrid functions

~~All Boolean functions~~

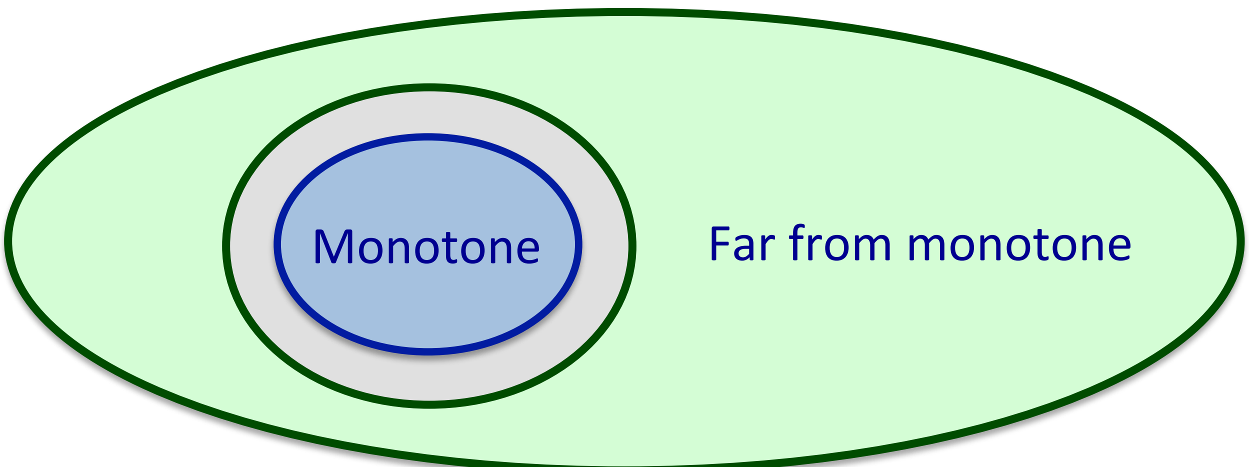

Theorem [Chen-Servedio-**T** 2014]

Any non-adaptive tester for testing monotonicity of $f : [m]^n \to \{0,1\}$ must make $\widetilde{\Omega}(n^{1/5})$ queries.

Proof by reduction to *m*=2 case (Boolean hypercube)

A useful characterization

Theorem. [Dodis *et al.* 1999]

$F : \{1, 2, \ldots, m\}^n \to \{0, 1\}$ is ε-far from monotone

$\updownarrow$

There exists $\varepsilon \cdot m^n$ vertex-disjoint pairs $(x_i, y_i) \in [m]^n$ such that $x_i \preceq y_i$ and $f(x_i) > f(y_i)$.

"violating pair"

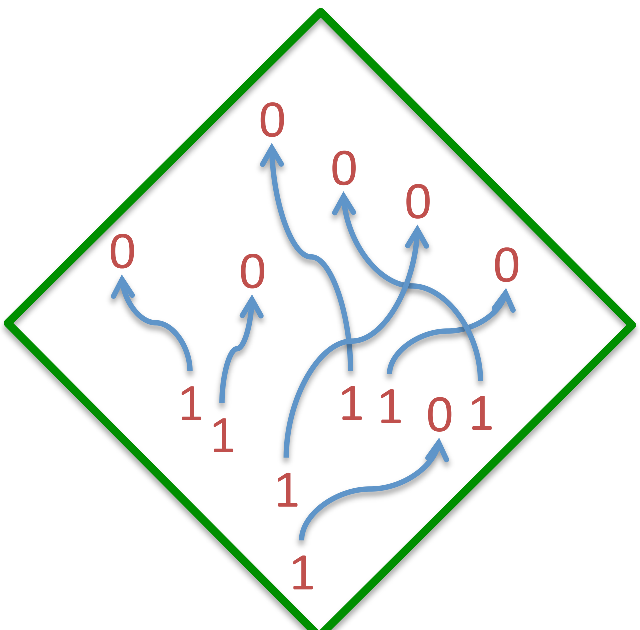

(Upward direction is easy)

Reducing to m = 2

Given any $f : \{0,1\}^n \to \{0,1\}$, define $F : [m]^n \to \{0,1\}$ as follows:

$$\underbrace{F(x_1, \ldots, x_n)}_{\text{numbers in } [m]} = \underbrace{f(\mathbf{1}[x_1 > \tfrac{m}{2}], \ldots, \mathbf{1}[x_n > \tfrac{m}{2}])}_{\text{bits in } \{0,1\}}$$

Easy: If *f* is monotone then so is *F.*

Remains to argue:

If *f* is *ε*-far from monotone then so is *F.*

Useful characterization

Useful characterization

suffices to show:

Exists $\varepsilon 2^n$ vertex-disjoint pairs in $\{0,1\}^n$ that are violations w.r.t. *f* $\longrightarrow$ Exists $\varepsilon \cdot m^n$ vertex-disjoint pairs in $[m]^n$ that are violations w.r.t. *F*

Each violating pair $\longrightarrow$ $(m/2)^n$ vertex-disjoint violating pairs

Each violating pair in $\{0,1\}^n \implies \left(\frac{m}{2}\right)^n$ violating pairs in $[m]^n$

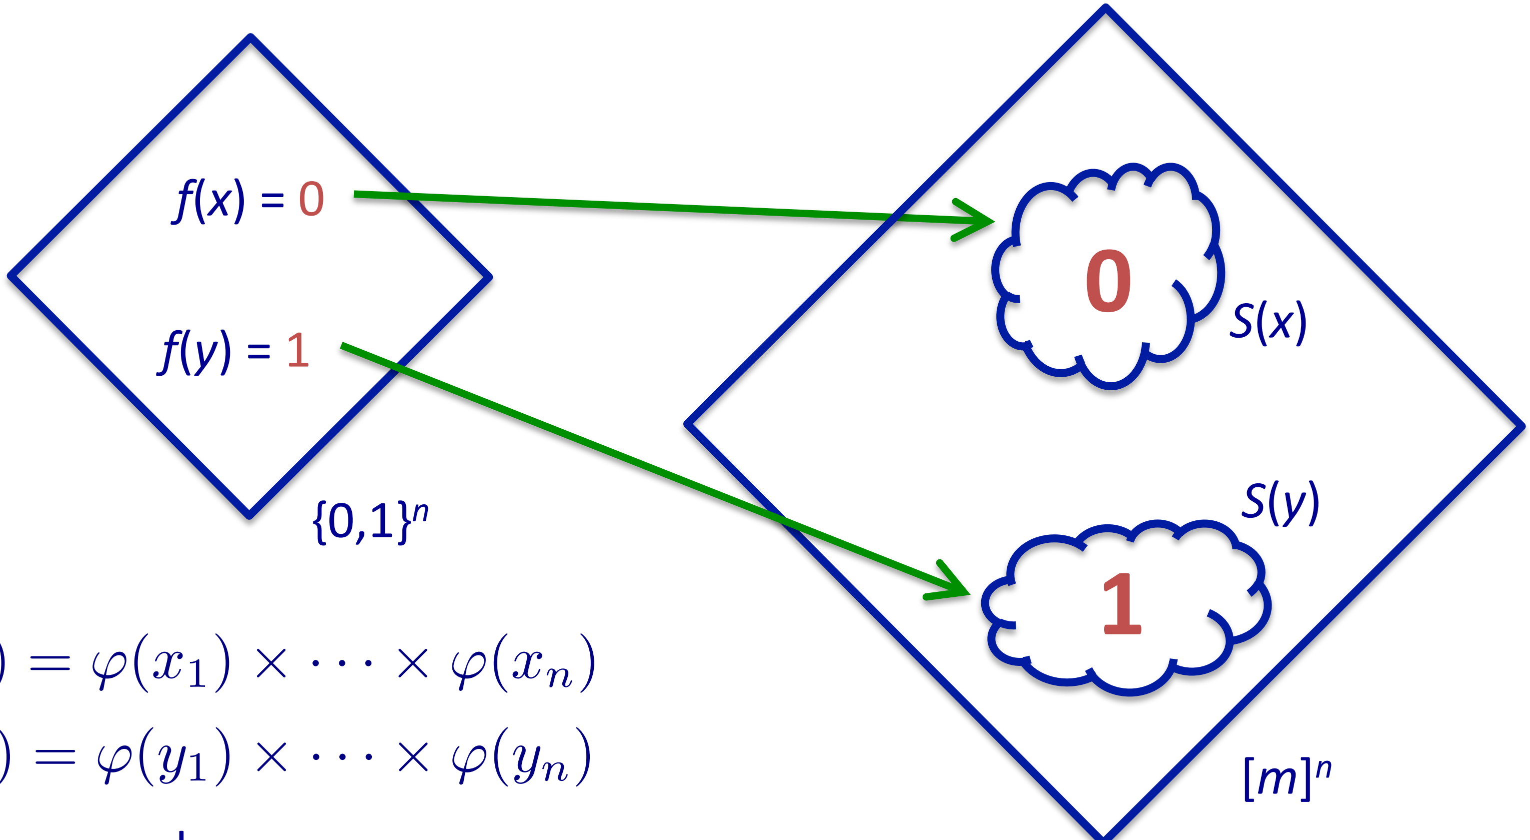

$$S(x) = \varphi(x_1) \times \cdots \times \varphi(x_n)$$
$$S(y) = \varphi(y_1) \times \cdots \times \varphi(y_n)$$

where

$$\varphi(0) = \{0, 1, \ldots, \tfrac{m}{2}\}$$
$$\varphi(1) = \{\tfrac{m}{2} + 1, \ldots, m\}$$

$|S(x)| = |S(y)| = (m/2)^n$

Easy end game: exhibit order-preserving bijection between $S(x)$ and $S(y)$

Conclusion

- A polynomial lower bound for testing monotonicity of Boolean functions

Theorem

Any non-adaptive tester must make $\widetilde{\Omega}(n^{1/5})$ queries.

Therefore, any adaptive tester must make $\Omega(\log n)$ queries.

- Main technical ingredient: multidimensional central limit theorems
- Proof extends to testing monotonicity over general hypergrid domains

Open Problem: Polynomial lower bounds against *adaptive* testers?

Thank you!