HARDNESS VIA DICTATOR VS "SPREAD-OUT" FUNCTION TESTS I

Dictator Testing → "Flipped view": If cut is large must be a dictator

Input: Truth table access to a $f: \{\pm 1\}^n \to \{-1,1\}$

Goal: 1) Choose q points $x^1, x^2, \ldots, x^q$ (appropriate prob distⁿ)

2) Apply some check to the q points & accept or reject accordingly.

HOPE: 1) Every dictator passes the check with some large probability $\geq c$

2) Every function that is far from a dictator passes the check w.p $\lesssim s$.

far from dictator $\equiv$ All coordinates have small "influence"

Def (C-degree Influence)

$$\mathrm{Inf}_i^c(f) = \sum_{S \ni i, |S| \leq c} \hat{f}(S)^2$$

Avoids the somewhat unintuitive phenomenon that there can be many coordinates of large influence (e.g. PARITY function)

<u>Lemma</u>: Let $f : \{-1,1\}^n \to [-1,1]$ ← this slight generality will be useful.

$|\{i \mid Inf_i^C(f) \geq \varepsilon\}| \leq \frac{C}{\varepsilon}$

<u>Proof</u>: $\sum_i Inf_i^C = \sum_i \sum_{\substack{S \ni i \\ |S| \leq C}} \hat{f}(S)^2 \leq C \cdot \sum_S \hat{f}(S)^2 \leq C$

By Markov's inequality, $\# i : Inf_i^C(f) \geq \varepsilon \leq \frac{C}{\varepsilon}$

How does this relate to hardness reductions?

<u>E.g.</u> A 2 query tester

1) Choose $x \in \{\pm 1\}^n$ uniformly at random

2) Choose $\mu_1, \ldots \mu_n$: each coord $= \begin{cases} -1 & \text{w.p. } \frac{1}{2} - \frac{1}{2}\rho \\ +1 & \text{w.p. } \frac{1}{2} + \frac{1}{2}\rho \end{cases}$

3) Accept if $f(x) \neq f(x \cdot \mu)$

Remind you of something?

$(x,y) \to$ edge distⁿ in the hypercube
$\overset{\shortparallel}{x \cdot \mu}$ rounding gap example for Max-Cut.

Check $\to$ "Is edge (x,y) cut".

What's the probability that dictators pass the check?

<u>Lemma</u> [Dictators on 2-query test]
Let $f: \{-1,1\}^n \to \{-1,1\}$ be a dictator.
Then f passes the test w.p. $\frac{1}{2} - \frac{1}{2}\rho$.

<u>Proof:</u> $\Pr_{(x,y)}[f(x) \neq f(y)] = \Pr_{(x,y)}[D_i(x) \neq D_i(y)]$

$$= \Pr_{(x,y)}[x_i \neq y_i] = \frac{1}{2} - \frac{1}{2}\rho.$$

So analyzing the test

~ proving that if f is far from a dictator, then it must pass the test with a much smaller probability.

Eg. $MAJ_3(x_1, x_2, x_3)$

passes the test with prob. $\frac{1}{2} - \frac{3}{8}\rho - \frac{1}{8}\rho^3$

Low Influence Functions?

What about $MAJ_n = \text{sign}(\sum_i x_i)$?

Lemma: $\forall \varepsilon > 0$, $\forall n$ large enough,

$$\Pr[\, MAJ_n \text{ passes}\,] \leq \frac{\text{arc-cos}(\rho)}{\pi} + \varepsilon$$

Sketch of a proof:

Test picks $\vec{x} \sim \{\pm 1\}^n$ uniformly

$\vec{\mu} \sim \{\pm 1\}^n \begin{cases} -1 & \text{w.p. } 1/2 - 1/2\rho \\ +1 & \text{w.p. } 1/2 + 1/2\rho \end{cases}$

& tests if

$$\underbrace{\langle \vec{x}, 1 \rangle}_{= \sum x_i} \;(\uparrow\, \sqrt{n}\,\vec{a}) \quad \& \quad \langle \vec{x} \cdot \vec{\mu}, 1 \rangle = \langle \vec{x}, \vec{b}\sqrt{n} \rangle$$

have different signs. $\qquad \vec{b} = \vec{\mu}/\sqrt{n}$.

Imagine picking $\vec{\mu}$ first and fixing it to a typical value. In particular, this fixes the inner product $\frac{1}{n}\sum_i x_i \cdot y_i = \frac{1}{n}\sum_i x_i^2 \cdot \mu_i$

$$= \frac{1}{n}\sum_i \mu_i$$

$$\sim \rho.$$

Conditioned on $\vec{\mu}$, test looks like

$\vec{x} \to \{\pm 1\}^n$ u.a.r.

Check if $\langle \vec{a}, \vec{x} \rangle$ & $\langle \vec{b}, \vec{x} \rangle$ have different signs.

Notice that now $\vec{a}, \vec{b}$ are fixed & only $\vec{x}$ is random.

Key Idea: Central Limit Theorem

$\langle \vec{a}, \vec{x} \rangle = \frac{1}{\sqrt{n}} \sum_i x_i \to$ gaussian with same mean & variance
$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad\; \shortparallel \qquad\qquad \shortparallel$
$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad\; 0 \qquad\qquad 1$

$\langle \vec{b}, \vec{x} \rangle = \frac{1}{\sqrt{n}} \sum_i \mu_i x_i \to N(0,1).$

Better way to state the same fact

$$\underbrace{\langle \vec{a}, \vec{x} \rangle}_{\text{uniform on } \{\pm 1\}^n} \overset{\text{dist}^n}{\approx} \underbrace{\langle \vec{a}, \vec{g} \rangle}_{\text{standard gaussian vector}}$$

$$\langle \vec{b}, \vec{x} \rangle \overset{\text{dist}^n}{\approx} \langle \vec{b}, \vec{g} \rangle$$

INVARIANCE PRINCIPLE

"2-D" CLT

$$\left(\langle \vec{a}, \vec{x} \rangle, \langle \vec{b}, \vec{x} \rangle\right) \overset{\text{dist}^n}{\approx} \left(\langle \vec{a}, \vec{g} \rangle, \langle \vec{b}, \vec{g} \rangle\right)$$

joint dist^n

What's $\Pr\left[\operatorname{Sign}(\langle \vec{a}, \vec{g} \rangle) \neq \operatorname{sign}(\langle \vec{b}, \vec{g} \rangle)\right]$

if $\langle \vec{a}, \vec{b} \rangle = \rho$?

$$\frac{\arccos(\rho)}{\pi} \quad !!$$

FACT: If $\vec{g}$ is std. gaussian vector, then $\dfrac{\vec{g}}{\|g\|_2}$ is uniformly distributed on the unit sphere.

It turns out that you can greatly generalize such "Invariance Principles" to say formally that

$\forall\ \rho, \varepsilon,\ \exists\ \tau, \delta > 0$

If $\mathrm{Inf}_i^{1/\delta}(f) \leq \tau \quad \forall i$, then

$$\Pr[f(\vec{x}) \neq f(\vec{y})]$$

$$\lesssim \Pr[f(\vec{g}) \neq f(\vec{g}')] + \varepsilon.$$

Max-Cut value on ρ-corr. sphere graph.

$\wr\wr$

Analysis of Feige-Schechtman Integrality gap.

So, to summarize,

the test

1) $\vec{x} \sim$ u.a.r. $\{\pm 1\}^n$

2) $\vec{\mu} \sim$ $\frac{1}{2}-\frac{1}{2}\rho$: -1
$\frac{1}{2}+\frac{1}{2}\rho$: 1

3) Acc $f(\vec{x}) \neq f(\vec{x}\cdot\vec{\mu})$.

Completeness

1) dictators pass w.p. $\frac{1}{2}-\frac{1}{2}\rho = C$

2) All f where all$_i$ have low $\frac{1}{\delta}$-inf pass w.p. $\sim$ soundness of integrality gap $\lesssim \dfrac{\text{arc-cos}(\rho)}{\pi} + \varepsilon$

Theorem [Reduction]

Suppose $\exists$ a dictator test with following properties ($f : \{\pm1\}^n \to \{\pm1\}$)

1) all n dictators pass w.p. $\geq c$.

2) if $h : \{\pm1\}^n \to \{-1,1\}$ s.t $\mathrm{Inf}_i^{1/\delta}(f) \leq \tau$, h passes w.p. $\leq s+\eta$.

3) checks of the test are of the form $P(x^1, x^2, \ldots, x^q) = 1$
"predicate" where $P \in \mathcal{F}$
family of predicates

then, Max-P problem is $(c-\eta, s+\eta)$ NP-Hard assuming the U.G.C.

We discussed such a test for

$P =$ "$\neq$" pred $=$ "max-cut" type

We'll prove Max-Cut hardness using it.

But the construction & analysis is general to any P.

Instance of U.G $=$ MAX 2LIN (p).

1) A bipartite, d-regular graph G

2) A $b_e = b_{i,j} \in \{0,1,\ldots,p-1\}$ for every edge $\{i,j\}$ of E.

We need to take this instance and produce a graph H $\rightarrow$ max-cut instance

Vertices $= n \cdot 2^p$.

For each $v \in V_G$, we will have

2^p vertices $\rightarrow$ "hypercube" for each $v \in V_G$

$\hookrightarrow$ p dim

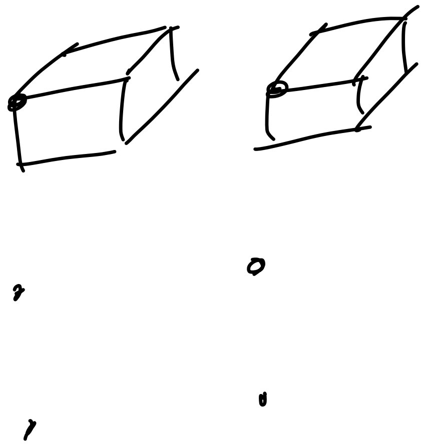

<u>Edges</u>? We'll use the tester to design an edge distribution.

Before that, note that any <u>cut</u> of H is described by

$$\{f_v : \{-1,1\}^p \rightarrow \{-1,1\}\}_{v \in V_G}.$$

$\downarrow$

which vertices of v-th hypercube are in

Idea: Use tester to "check" that the UG constraints are satisfied. that way passing the test will be related to value of the input UG instance.

In fact, we will like to ensure that we can learn a good assignment for UG instance by looking at a large cut $\{f_v\}_{v \in V_G}$

$\downarrow$

Intended assignment: $f_v = D_i$ for some $i \in \{0, 1, .., p-1\}$.

HOPE: this encodes that v should get assigned the label i.

Suppose all f_vs are dictators.
How should we check that they are "correct" dictators?

Need that if v is assigned i

w is assigned j

& $\{v, w\}$ is sat, then

$$j = i + b_{ij} \bmod p$$

<u>Def (g_v^w)</u> [Opinion of v-s nbrs]

For $w \sim v$ in G, let

$$g_v^w : \{-1,1\}^p \to \{-1,1\} \text{ by}$$

$$g_v^w = f_w \circ \sigma_{v \to w} \quad \text{where } \sigma_{v\to w} : [p] \to [p] \text{ perm.}$$

$$\sigma_{v \to w}(x) = y : y_i = x_{\sigma_{v \to w}(i)}$$

$\sigma_{v \to w}(i) = b_{v,w} - i$ (RHS in UG inst.)

Notice that if $f_v = D_i$

$f_w = D_j$

& (v,w) is sat by (i,j)

then $g_v^w = f_w \circ \sigma_{v \to w} = f_v$.

So if 1) f_v are dictators

2) that encode sat assignment to a nbr pair $\{v,w\}$

then $g_v^w = f_v$.

If not, then g_v^w may not be sensible...

TEST $\equiv$ Edge distribution with

1) pick $v \in V$ uniformly at random

2) pick 2 random nbrs
w_1, w_2 of v

3) Apply Tester to $g_v^{w_1}$ & $g_v^{w_2}$

$\|$

1) generate x at random from v-s hypercube

2) generate μ : $\frac{1}{2}-\frac{1}{2}\rho$ -1
$\frac{1}{2}+\frac{1}{2}\rho$ 1

3) add check $=$ $g_v^{w_1}(x) \neq g_v^{w_2}(x)$

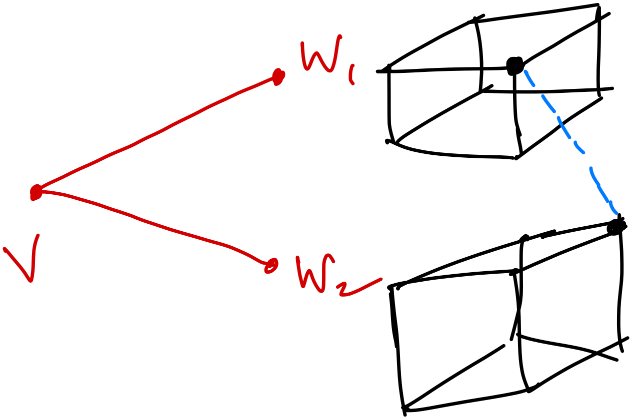

CAUTION: The edge distribution generates edges that are typically between vertices of two different hypercubes (corresponding to 2 diff nbrs of v)

NEXT TIME: Completeness, Soundness, discussions of generalization to other predicates P.

Analysis of the Reduction

Lemma (Completeness)

Suppose the U.G. instance G is $(1-\lambda)$ satisfiable. Then, there's an assignment $\{f_v\}_{v \in V_G}$ for vertices V_H of H that cuts $\geq (C - 2\lambda)$

$C \longrightarrow$ completeness of dictator test

<u>Proof:</u> Let A be assignment that satisfies $(1-\lambda)$-frac of constraints. We'll give an assignment $\{f_v\}$ to V_H and analyze its performance on the edge distribution of V_H.

$\rightarrow f_v = D_{A(v)} \leftarrow$ dictator on $A(v)$-th bit.

<u>Analysis</u>: distribution of $\{v, w_1\}$ & $\{v, w_2\}$ individually is uniform over all constraints of G.

<u>So</u> w.p. $\geq (1-2\lambda)$ over the choice of random v & two random nbrs w_1, w_2 of v, A satisfies (v, w_1) & (v, w_2).

Condition on this event.

In this case, note that $g_v^{w_1} = g_v^{w_2} = f_v$ & f_v is a dictator.

<u>So</u>, dictator test on v, w_1, w_2 passes w.p. $\geq c$.

Thus, in total, the test passes

w.p. $\geq \; c \cdot (1-2\lambda)$

$\Rightarrow$ frac of edges cut $\geq c(1-2\lambda)$

$\boxed{\text{Lemma (Soundness)}}$

Suppose $\{f_v\}$ passes the test w.p. $\geq s + \eta$. Then, there's an assignment A that satisfies $\geq$ constraints.

depends on η, τ but not p.

Proof:

$g_v^{w_1}, g_v^{w_2}$ pass the check.

w.p.

$$= \frac{1}{2} - \frac{1}{2} \mathop{\mathbb{E}}_{\vec{x},\vec{\mu}} \left[g_v^{w_1}(\vec{x}) \cdot g_v^{w_2}(\vec{x}\cdot\vec{\mu}) \right]$$

$\{f_v\}$s pass the checks w.p. $\geq S + \eta$

This means that

$$S+\eta \leq \mathop{\mathbb{E}}_{v,w_1,w_2} \frac{1}{2} - \frac{1}{2} \mathop{\mathbb{E}}_{\vec{x},\vec{\mu}} \left[g_v^{w_1}(\vec{x}) \cdot g_v^{w_2}(\vec{x}\cdot\vec{\mu}) \right]$$

$$= \mathop{\mathbb{E}}_{v,\vec{x},\vec{\mu}} \left(\frac{1}{2} - \frac{1}{2} \mathop{\mathbb{E}}_{w_1,w_2} \left[g_v^{w_1}(\vec{x}) \cdot g_v^{w_2}(\vec{x}\cdot\vec{\mu}) \right] \right) \quad \text{— (1)}$$

<u>Def</u> (Average Opinion about v)

$$h_v = \mathop{\mathbb{E}}_{w \sim v} \left[g_v^w \right].$$

<u>NOTE</u>: $h_v : \{-1,1\}^p \to [-1,1]$ interval instead of $\{-1,1\}$.

$$s+\eta \le \mathbb{E}_{v,\vec{x},\vec{\mu}} \left(\frac{1}{2} - \frac{1}{2} h_v(\vec{x}) \cdot h_v(\vec{x}\cdot\vec{\mu}) \right) \quad (*)$$

average opinion.

FACT (averaging)

Suppose Z is a r.v. such that $|Z| \le 1$ & $\mathbb{E} Z \geq s+\eta$

$$\Pr\left[Z \geq s + \frac{\eta}{2}\right] \geq \frac{\eta}{2}$$

Proof:

$$s+\eta \le \mathbb{E} Z \le \Pr\left[Z \le s+\frac{\eta}{2}\right]\left(s+\frac{\eta}{2}\right) + \Pr\left[Z \geq s+\frac{\eta}{2}\right]\cdot 1$$

or $\Pr\left[Z \geq s+\frac{\eta}{2}\right] \geq \frac{\eta}{2}$

Let's apply averaging to the r.v.

$$\mathbb{E}_{\vec{x},\vec{\mu}}\left[\frac{1}{2}-\frac{1}{2}h_v(\vec{x})\cdot h_v(\vec{x}\cdot\vec{\mu})\right]$$

(where randomness is over v).

<u>Then</u> averaging gives that for $\frac{\eta}{2}$ frac of v, it must hold that

$$s+\frac{\eta}{2}\le \mathbb{E}_{\vec{x},\vec{\mu}}\left[\frac{1}{2}-\frac{1}{2}h_v(\vec{x})\cdot h_v(\vec{x}\cdot\vec{\mu})\right]$$

Call such v GOOD.

From <u>Strong Soundness</u> of the test, we must have $\exists i$, there must be an influential variable

$$\mathrm{Inf}_i^{1/8} \lesssim \tau$$

$$\tau \le \mathrm{Inf}_i^{1/8}(\tau)$$

$$= \sum_{S \ni i,\, |S| \le 1/8} \hat{h}_v(S)^2$$

$$= \sum_{\substack{S \ni i \\ |S| \le 1/8}} \widehat{\mathbb{E}_{w \sim v}[g_v^w]}(S)^2$$

$$\le \sum_{\substack{S \ni i \\ |S| \le 1/8}} \mathbb{E}_{w \sim v} \widehat{g_v^w}(S)^2$$

Cauchy-Schwarz inequality

Now, $g_v^w = f_w \circ \sigma_{v\to w}$

so, $\widehat{g_v^w}(S) = \hat{f}_w(\sigma_{v\to w}^{-1}(S))$.

So,
$$\sum_{\substack{S\ni i \\ |S|\le 1/\delta}} \mathop{\mathbb{E}}_{w\sim v}\left[\widehat{g_v^w}(S)^2\right]$$
$$= \mathop{\mathbb{E}}_{w\sim v}\left[\sum_{\substack{S\ni i \\ |S|\le 1/\delta}} \hat{f}_w(\sigma_{v\to w}^{-1}(S))^2\right]$$
$$= \mathop{\mathbb{E}}_{w\sim v} \sum_{T\ni\sigma_{v\to w}(i),\ |T|\le 1/\delta} \hat{f}_w(T)^2$$

$T \to \sigma_{v\to w}^{-1}(S)$.

All in all,
$$\tau \le \mathop{\mathbb{E}}_{w\sim v}\left[\mathrm{Inf}^{1/\delta}_{\sigma_{v\to w}(i)}(f_w)\right]$$

.

Averaging: for $\tau/2$ frac $w \sim v$,

$$\mathrm{Inf}^{\gamma/8}_{\sigma_{v\to w}(i)}(f_w) \geq \frac{\tau}{2}$$

GOOD NBRS

Here, we have that

1. $\frac{\eta}{2}$ v are GOOD, h_v has an inf. var.
2. for an inf var in h_v, $\frac{\tau}{2}$ nbrs of v, have an influential var $\sigma_{v\to w}(i)$.

$(i, \sigma_{v\to w}(i))$ is a good assignment for edge (v,w) in UG instance G.

So, "returning a random influential variable should be a good assignment if # inf vars were small.

This necessitates "low-deg influence".

<u>Decoding Scheme:</u>

for any v

Let $\mathrm{Cand}(v)$

$$= \left\{ i : \mathrm{Inf}_i^{1/\delta}(f_v) > \tau, \text{ or } \mathrm{Inf}_i^{1/\delta}(h_v) > \tau \right\}$$

<u>Then</u>, $|\mathrm{Cand}(v)| \leq \frac{2}{\tau \cdot \delta}$.

decoding:

$$\boxed{\text{for } v: \text{ assign } A(v) = \text{random element of } Cand(v).}$$

Analysis

For GOOD–GOOD nbrs, random choice agrees w.p $\geq \frac{1}{\left(\frac{2}{\tau\delta}\right)^2}$

$$= \frac{\tau^2\delta^2}{4}.$$

GOOD–GOOD Nbrs are $\geq \frac{\eta}{2} \cdot \frac{\tau}{2}$ frac

So, in all $\frac{\eta \cdot \tau^3 \delta^2}{16}$ frac constraints of UG are satisfied.

Choose p large enough so that

$$\lambda < \frac{\eta \cdot \tau^3 \cdot \delta^2}{16}.$$

Then we have a contradiction.

$\square$.

What did we use about the Max-Cut problem here?

① dictator test

② soundness for ① low-deg influence

& ② $[-1,1]$-valued functions.

Raghavendra '08 : Let P be any predicate, i.e., $P : \{-1,1\}^n \to \{0,1\}$

e.g. $P(x_1, x_2, x_3) = x_1 \vee x_2 \vee x_3 \hookrightarrow$ 3SAT pred.

There's a natural generalization of GW SDP for Max-P problems.

Thm: If there is a (c,s)-integrality gap for this SDP, then there is a dictator test with completeness $c-\eta$ & soundness $s+\eta$ $\forall \eta > 0$.