Computer Science 294 Lecture 20 Notes

Daniel Raban

March 23, 2023

1 Hardness of Approximation

1.1 Constraint statisfaction problems

Definition 1.1. A **constraint satisfaction problem (CSP)** over domain Ω is defined by a finite set of predicates Ψ where each $\psi \in Psi$ is some constraint $\psi : \Omega^r \to \{0, 1\}$. The **arity** of a CSP is the maximal arity r of a predicate in Ψ.

Example 1.1. Max-3SAT, Max-Cut, Max-3LIN, and Max-3Coloring are al CSPs.

All these problems are NP-hard.

Example 1.2. In Max-3SAT, $\Omega = \{T, F\}$, and a predicate could be something like $\psi(v_1, v_2, v_3) = v_1 \wedge v_2 \wedge v_3$ or $\psi(v_1, v_2, v_3) = \overline{v_1} \wedge v_2 \wedge \overline{v_3}$. The arity is 3.

Definition 1.2. An **instance** P for MAxCSP(Ψ) over variable set V is a list of tuples: (scope, predicate) $C = (s, \psi)$, where $\psi \in \Psi$ and $s = (v_1, \dots, v_r)$ is a typle of variables in V.

Definition 1.3. An **assignment** for an instance P is a labeling $F : V \to \Omega$. F **satisfies a constraint** (s, ψ) if $\psi(F(s)) = 1$, where $F(s) = (F(v_1), \dots, F(v_r))$.

Definition 1.4. The **value** pf F on P, is the fraction of constraints in P satisfied by F. That is,

$$\mathrm{Val}_p(F) = \mathbb{E}_{(S,\psi)\sim P}[\psi(F(S))].$$

The optimal value is

$$\mathrm{OPT}(P) := \max_F \mathrm{Val}_P(F).$$

Last time, we discussed string testers. The main insight is that CSPs are the same as string testers. Here is a dictionary between CSPs and string testing.

CSP instance P	string tester
Assignment $F : [n] \to \Omega$	$\omega \in \Omega^n$
Value of F	$\mathbb{P}$(tester accepts ω)
Ψ	predicates you apply in the string tester
number of constraints	$2^{\#\text{ random bits}}$

Example 1.3. Take Max-3SAT, for example. We think of the string as an assignment to all the variables. Our queries ask for the value of 3 bits in the string, and the predicates are the predicates of the DNF, things like $x_1 \wedge \overline{x_2} \wedge x_3$.

Definition 1.5. An **(α, β)-approximation algorithm** for MaxxCSP(Ψ) is an algorithm that on instances where whose best assignment has value $\geq \beta$, the algorithm is guaranteed to output and assignment with value $\geq \alpha$.[1]

Example 1.4. $(1, 1)$ approximating Max-Cut is easy because this is the case where the graph is bipartite.

Example 1.5. $(1/2, 0.51)$ approximating Max-Cut is easy because if we randomly cut each edge with probability $1/2$, we will cut half the edges on average.

Example 1.6. $(1, 1)$-approximating MaxSAT is NP-hard because if we can find a satisfying assignment for any formula which can be satisfied, then we can solve 3SAT.

Example 1.7. $(1, 1)$-approximating MaxColoring is NP-hard because solving this problem would allow us to solve 3Coloring.

Example 1.8. $(1, 1)$-approximating Max-3LIN is easy because because we can just use Gaussian elimination to see whether a system of linear equations has a solution.

However, the following theorem tells us that Gaussian elimination is not robust for solving this CSP.

Theorem 1.1 (Håstad)**.** *$(0.51, 0.99)$-approximating Max-3LIN is NP-hard.*

Theorem 1.2 (PCP theorem)**.** *$(1 - \delta_0, 1)$-approximating Max-3SAT is NP-hard.*

However, you can use a randomized algorithm (and then de-randomize it) to show the following.

Proposition 1.1. *$(7/8, 1)$-approximating Max-3SAT is easy.*

Theorem 1.3 (Håstad)**.** *For all constants $\delta > 0$, $(7/8 + \delta, 1)$-approximating Max-3SAT is NP-hard.*

[1] You can remember which variable is which by βest and αlgorithm.

1.2 Testing for dictators vs no notable coordinates

Håstad's idea was that to prove hardness for MaxCSP(Ψ), it suffices to design a "relaxed" dictator test that only uses predicates from Ψ. By relaxed, we mean that it is enough to make a test which rejects with "decent" probability when f is "very far" from being all dictators.

Recall that

$$\mathrm{Inf}_i^{(\rho)}(f) := \mathrm{Stab}_\rho(F_i f) = \sum_{S \ni i} \rho^{|S|-1} \widehat{f}(S)^2.$$

Qualitatively, we think of this as a "noisy influence."

Definition 1.6. The **total annotated influence** is

$$\begin{aligned}\mathbb{I}^{(\rho)} &:= \sum_{i=1}^{n} \mathrm{Inf}_o^{(\rho)}(f)\\ &= \sum_{S \neq \varnothing} |S| \rho^{|S|-1} \widehat{f}(S)^2\\ &= \sum_{k=1}^{n} k\rho^{k-1} W^k(f).\end{aligned}$$

Lemma 1.1. *For all* $0 < \rho < 1$ *and for all* k, $k\rho^{k-1} \leq \frac{1}{1-\rho}$.

This tells us that

$$\mathbb{I}^{(\rho)}(f) \leq \frac{\rho}{1-\rho}.$$

Definition 1.7. We sat that a coordinate j is ε**-notable** if $\mathrm{Inf}_j^{(1-\varepsilon)}(f) \geq \varepsilon$.

Example 1.9. If $f = \chi_i$, then $\mathrm{Inf}_i^{(\rho)}(f) = 1$.

Example 1.10. If $f = \mathrm{PARITY}_n = \chi_{[n]}$, then $\mathrm{Inf}_i^{(\rho)}(f) = \rho^{n-1}$.

Even though these are both characters, once we apply noise, χ_i has influence 1 and the parity function has exponentially small influence.

Example 1.11. If f is a $O(1)$-junta, then it has a notable coordinate.

Definition 1.8. An (α, β)**-Dictator-vs-No-Notable-Coordinates test using** Ψ is a function tester that for all $n \in \mathbb{N}$ can be applied to functions $f : \{\pm 1\}^n \to \{\pm 1\}$ and

- makes at most r queries to f and applies a preducate from Ψ.
- If f is a dictator, then $\mathbb{P}(\text{tester accepts } f) \geq \beta$.

- For all small enough $\varepsilon > 0$, if f has no ε-notable coordinates, then

$$\mathbb{P}(\text{tester accepts } f) \le \alpha + \lambda(\varepsilon), \qquad \text{where } \lambda(\varepsilon) \xrightarrow{\varepsilon \to 0} 0.$$

Theorem 1.4. *Suppose there exists an (α, β)-Dictator-vs-No-Notable-Coordinates test using Ψ. Then for all $\varepsilon > 0$, $(\alpha+\varepsilon, \beta-\varepsilon)$-approximating* $\mathrm{MaxCSP}(\Psi)$ *is Unique-Games-hard.*

In other words, there exists a polynomial time reduction so that a YES instance for the Unique Games problem is mapped to an Instance with $\mathrm{Val} \ge \beta - \varepsilon$ *and so that a NO instance for the Unique Games problem is mapped to an Instance with* $\mathrm{Val} \le \alpha + \varepsilon$*.*

Example 1.12. The **Unique Games**(q, δ) **problem** is a CSP with domain $\Omega = \{0, 1, \ldots, q-1\}$ and constraints like $x_7 - x_5 \equiv 3 \pmod 9$ or $x_5 - x_{11} \equiv 2 \pmod 9$. In YES instances, a $1 - \delta$ fraction of the constraints can be satisfied simultaneously. In a No instance, any assignment satisfies a $\le \delta$ fraction of the constraints.

Conjecture 1.1 (Unique games conjecture)**.** *For all $\delta > 0$, there exists a $q \in \mathbb{N}$ such that* $\mathrm{UG}(q, \delta)$ *is NP-hard.*

Håstad made a test based on the idea of BLR lienarity testing.

Theorem 1.5 (Håstad)**.** *For all $\delta > 0$, there is a $(1/2, 1-\delta)$-Dictator-vs-No-Notable-Coordinates test using 3LIN equations (such as $x_i \oplus x_j \oplus x_k \equiv 6 \pmod 2$). The test is*

- *Pick $X, Y \sim \{\pm 1\}^n$ uniformly at random.*
- *Pick a bit $B \sim \{\pm 1\}$ uniformly at random*
- *Let $Z \in \{\pm 1\}^n$ be defined as $Z_i = X_i \cdot Y_i \cdot B$.*
- *Take $Z' \sim N_{1-\delta}(Z)$*
- *Query $f(X), f(Y), f(Z)$ and accept iff $f(X) \cdot f(Y) \cdot f(Z') = B$.*

In particular,

$$\mathbb{P}(\textit{tester accepts } f) = \frac{1}{2} + \frac{1}{2} \sum_{|S| \text{ odd}} \widehat{f}(S)^3 (1-\delta)^{|S|}.$$

In $\mathbb{F}_2$ notation, the predicate we are checking is a linear equation: $F(x) + F(y) + f(z) \equiv 0 \pmod 2$, where $f(x) = (-1)^{F(x)}$. With z', this is a noisy linear equation.

Notice that if $f = \chi_i$ is a dictator, then

$$\mathbb{P}(\text{tester accepts}) = \frac{1}{2} + \frac{1}{2}(1-\delta) = 1 - \frac{\delta}{2}.$$

Proof. We want to calculate

$$\mathbb{P}(\text{tester accepts}) = \frac{1}{2} + \frac{1}{2}\,\mathbb{E}_{X,Y,Z',B}[f(X)f(Y)f(Z')B].$$

Look at

$$\begin{aligned}
\mathbb{E}_{X,Y,Z',B}[f(X)f(Y)f(Z')B \mid B = 1] &= \mathbb{E}[f(X)f(Y)\,\mathbb{E}_{Z'\sim N_{1-\delta}(X\cdot Y)}[f(Z')]] \\
&= \mathbb{E}[f(X)f(\underbrace{Y}_{X\cdot Z})T_{1-\delta}f(\underbrace{X\cdot Y}_{Z})] \\
&= \mathbb{E}_Z[\mathbb{E}_X[f(X)f(X\cdot Z)]T_{1-\delta}f(Z)] \\
&= \mathbb{E}_Z[(f*f)(Z)T_{1-\delta}f(Z)]
\end{aligned}$$

Using Plancherel's theorem,

$$\begin{aligned}
&= \sum_S \widehat{f*f}(S)(1-\delta)^{|S|}\widehat{f}(S) \\
&= \sum_S \widehat{f}(S)^3(1-\delta)^{|S|}.
\end{aligned}$$

The same calculation with $B = -1$ gives

$$\mathbb{E}_{X,Y,Z',B}[f(X)f(Y)f(Z')B \mid B = 1] = \sum_S \widehat{f}(S)^3(-(1-\delta))^{|S|}.$$

So we get We want to calculate

$$\begin{aligned}
\mathbb{P}(\text{tester accepts}) &= \frac{1}{2} + \frac{1}{4}\left(\mathbb{E}[f(X)f(Y)f(Z')B \mid B = 1] + \mathbb{E}[f(X)f(Y)f(Z')B \mid B = -1]\right) \\
&= \frac{1}{2} + \frac{1}{2}\sum_{S:|S|\text{ odd}} \widehat{f}(S)^3(1-\delta)^{|S|}.
\end{aligned}$$

Now we can check that if $\varepsilon < \delta$ and f has no ε-notable coordinates, then

$$\begin{aligned}
\mathbb{P}(\text{tester accepts } f) &= \frac{1}{2} + \frac{1}{2}\sum_{S:|S|\text{ odd}} \widehat{f}(S)^3(1-\delta)^{|S|} \\
&\le \frac{1}{2} + \frac{1}{2}\max_{S:|S|\text{ odd}} \cdot \sum_S \widehat{f}(S)^2 \\
&\le \frac{1}{2} + \frac{1}{2}\sqrt{\max_{S:|S|\text{ odd}} \widehat{f}(S)^2(1-\delta)^{|S|-1}} \\
&\le \frac{1}{2} + \frac{1}{2}\sqrt{\max_i \mathrm{Inf}_i^{(1-\delta)}(f)} \\
&\le \frac{1}{2} + \frac{1}{2}\sqrt{\max_i \mathrm{Inf}_i^{(1-\varepsilon)}(f)} \\
&\le \frac{1}{2} + \frac{1}{2}\sqrt{\varepsilon}.
\end{aligned}$$

□