

Lecture 15

Rounding SDPs for CSPs*

Recall the canonical SDP relaxation for an instance \mathcal{C} of $\text{CSP}(\Gamma)$:

$$\max \sum_{C=(R,S) \in \mathcal{C}} w_C \Pr_{L \sim \lambda_C} [L(S) \text{ satisfies } R],$$

subject to the following conditions:

- for all C , λ_C is a probability distribution on assignments $S \rightarrow D$;
- $(I_v[\ell])_{v \in V, \ell \in D}$ are joint random variables (which can be thought of as vectors), called “pseudoindicators,” that satisfy:

1. [optional] for all $C \in \mathcal{C}$, for all $v \in C$, and for all $\ell \in D$,

$$\mathbf{E}[I_v[\ell]] = \Pr_{L \sim \lambda_C} [L(v) = \ell];$$

2. for all $C \in \mathcal{C}$, for all $v \in C$, for all $\ell \in D$, for all $v' \in V$, and for all $\ell' \in D$,

$$\mathbf{E}[I_v[\ell] \cdot I_{v'}[\ell']] = \Pr_{L \sim \lambda_C} [L(v) = \ell \text{ and } L(v') = \ell'].$$

Also recall the following equivalent perspectives on the canonical SDP relaxation for a CSP:

- pseudoindicator random variables satisfying conditions 1 and 2 above;
- pseudoindicator random variables satisfying condition 2 above;
- a vector solution satisfying conditions 1 and 2 above (when viewed as a collection of pseudoindicator random variables);
- jointly Gaussian pseudoindicators satisfying conditions 1 and 2 above.

*Lecturer: Ryan O’Donnell. Scribe: Brian Kell.

15.1 Jointly Gaussian pseudoinicators

For this section, let $N = |V| \cdot |D|$. [Is this what N is supposed to mean, or is N arbitrary?]

Definition 15.1. Random variables Z_1, Z_2, \dots, Z_N are said to be *jointly Gaussian* if

$$\begin{bmatrix} Z_1 \\ Z_2 \\ \vdots \\ Z_N \end{bmatrix} = \begin{bmatrix} \mu_1 \\ \mu_2 \\ \vdots \\ \mu_N \end{bmatrix} + L \begin{bmatrix} X_1 \\ X_2 \\ \vdots \\ X_M \end{bmatrix},$$

where L is an $N \times M$ matrix and X_1, X_2, \dots, X_M are mutually independent standard normal random variables (that is, normal random variables with mean 0 and standard deviation 1). The *covariance matrix* Σ is the matrix $(\Sigma_{ij})_{i=1}^N_{j=1}^N$, where $\Sigma_{ij} = \mathbf{Cov}[Z_i, Z_j]$.

Fact 15.2. The covariance matrix is given by $\Sigma = LL^\top$.




Figure 15.1: Density of joint Gaussian random variables X and Y with $\mathbf{Var}(X) = 2$, $\mathbf{Var}(Y) = 1$, and $\mathbf{Cov}(X, Y) = -1$. [From PlanetMath.org]

Proposition 15.3. We can assume (computationally efficiently) that the pseudoinicators are jointly Gaussian. In other words, given a vector solution satisfying conditions 1 and 2, we can efficiently produce jointly Gaussian pseudoinicators satisfying conditions 1 and 2.

Proof. We start with a vector solution $\vec{y}_{v,\ell} \in \mathbb{R}^N$ satisfying conditions 1 and 2. Let Z_1, \dots, Z_N be mutually independent standard normal random variables.

As a first attempt, we define a collection of jointly distributed random variables $(G_v[\ell])_{v,\ell}$ as follows: To get one draw from (all of) the $G_v[\ell]$, pick $i \in [N]$ uniformly at random and output $G_v[\ell] = Z_i(\vec{y}_{v,\ell})_i$. Then the $G_v[\ell]$ are jointly Gaussian. [Why?] Moreover, they satisfy condition 2:

$$\mathbf{E}_{Z,i}[G_v[\ell]G_{v'}[\ell']] = \mathbf{E}_Z \left[\sum_i \frac{1}{N} Z_i(\vec{y}_{v,\ell})_i \cdot Z_i(\vec{y}_{v',\ell'})_i \right] = \frac{1}{N} \sum_i (\vec{y}_{v,\ell})_i (\vec{y}_{v',\ell'})_i \cdot \mathbf{E}[Z_i^2] =: \langle \vec{y}_{v,\ell}, \vec{y}_{v',\ell'} \rangle.$$

However, when we check to see that they satisfy condition 1, we find that we have

$$\mathbf{E}_{Z,i}[G_v[\ell]] = \mathbf{E}_Z \left[\sum_i \frac{1}{N} Z_i(\vec{y}_{v,\ell})_i \right] = \frac{1}{N} \sum_i (\vec{y}_{v,\ell})_i \cdot \mathbf{E}[Z_i] = 0.$$

To fix this problem, we instead output

$$G_v[\ell] = \text{avg}_j(\vec{y}_{v,\ell})_j + Z_i((\vec{y}_{v,\ell})_i - \text{avg}_j(\vec{y}_{v,\ell})_j).$$

This definition of $G_v[\ell]$ satisfies both conditions 1 and 2. [Why? And why are these new $G_v[\ell]$ jointly Gaussian?] \square

15.2 Goemans–Williamson for Max-Cut, redux

The Goemans–Williamson algorithm for Max-Cut can be viewed as follows:

1. Solve the SDP and get jointly Gaussian pseudoindicators $(G_v[0], G_v[1])_{v \in V}$.
2. Draw once from them to obtain numbers $(g_v[0], g_v[1])_{v \in V}$.
3. Output the assignment $F(v) = \text{argmax}_{\ell \in \{0,1\}} \{g_v[\ell]\}$.

[Why can GW be viewed this way?]

The analysis of this algorithm uses the following fact, proved by Sheppard in 1899 [She99] [Is this the right reference?]: If Z and Z' are standard normal random variables and $\mathbf{E}[ZZ'] = \rho$, then

$$\Pr[\text{sgn}(Z) = \text{sgn}(Z')] = \frac{1}{2} + \frac{1}{\pi} \cos^{-1} \rho.$$

15.3 The Unique Games Conjecture

Definition 15.4. The problem *Unique-Games* $_q$, abbreviated UG_q , is a 2-CSP over the domain $D = \{0, 1, 2, \dots, q-1\}$, for which

$$\Gamma = \{ R : D^2 \rightarrow \{0, 1\} \mid \exists \text{ permutation } \pi \text{ of } D \text{ s.t. } R(a, b) = 1 \text{ iff } \pi(a) = b \};$$

in other words, the only allowable constraints in an instance of UG_q are bijective constraints of arity 2.

Example 15.5. Max-Cut is a subproblem of UG_2 . In an instance of Max-Cut, every edge (i, j) in the graph has a corresponding constraint $x_i \neq x_j$, that is, the allowable assignments are $\{\{x_i = 0, x_j = 1\}, \{x_i = 1, x_j = 0\}\}$. If the graph is bipartite, then all of these constraints are simultaneously satisfiable; otherwise, the objective of Max-Cut is to maximize the fraction of constraints that are satisfied.

Remark 15.6. For UG_2 , the greedy algorithm (local propagation) is a $(1, 1)$ -approximation algorithm. This is like 2-coloring a bipartite graph (2-coloring is a subproblem of UG_2). In an instance of UG_2 , the domain is $D = \{0, 1\}$, and all constraints are of the form $x_i = x_j$ or of the form $x_i \neq x_j$. If it is possible to simultaneously satisfy all of these constraints (that is, if $\text{Opt} = 1$), then we can easily find a satisfying assignment by arbitrarily choosing a starting variable and arbitrarily assigning it either 0 or 1; from then on, the value of any unassigned variable that shares a constraint with a variable that has already been assigned is determined by the constraint. The arbitrary choices we made at the beginning are not important, because the “complement” assignment of any satisfying assignment is also a satisfying assignment.

Conjecture 15.7. Unique Games Conjecture [Kho02]. For every $\epsilon > 0$ there exists a value of q such that $(1/2, 1 - \epsilon)$ -approximating UG_q is NP-hard. [Any $\alpha \in (0, 1)$ can be substituted for $1/2$ here.]

Remark 15.8. The Unique Games Conjecture implies many other optimal inapproximability results.

On the other hand, we can get a $(1/2, 1 - \epsilon)$ -approximation by solving an SDP of size $\exp(qn^{\epsilon^{\Theta(1)}})$ [ABS10], [BRS11].

We can $(1 - O(\sqrt{\log q}\sqrt{\epsilon}), 1 - \epsilon)$ -approximate UG_q by SDP rounding [CMM06]. We can do the following:

1. Solve the canonical SDP relaxation to get a collection of jointly Gaussian pseudoindicators $(G_v[\ell])_{v \in V, \ell \in \{0, \dots, q-1\}}$.
2. Draw once from them to obtain numbers $(g_v[0], g_v[1], \dots, g_v[q-1])_{v \in V}$.
3. Output the assignment $F(v) = \text{argmax}_{\ell} \{g_v[\ell]\}$.

Note that this is a randomized algorithm.

“Theorem” 15.9. If $\text{SDPOpt}(\mathcal{C}) \geq 1 - \epsilon$, then

$$\mathbf{E}_F[\text{Val}_{\mathcal{C}}(F)] \geq 1 - O(\sqrt{\log q}\sqrt{\epsilon}).$$

The proof of this “theorem” is left as an exercise (and could probably be published as a paper). It should be noted that CMM did something slightly different.

15.4 (α, β) -decision algorithms

Recall that an algorithm is said to (α, β) -approximate $\text{CSP}(\Gamma)$ if whenever $\text{Opt}(\mathcal{C}) \geq \beta$ the algorithm produces a solution with value at least α . [For a proposed rounding algorithm, we often prove a theorem that looks something like: If $\text{SDPOpt}(\mathcal{C}) \geq \beta$, then our rounding algorithm produces a solution with value at least α . In particular, this kind of theorem shows that there exists a solution with value at least α .]

Definition 15.10. An algorithm (α, β) -decides $\text{CSP}(\Gamma)$ if:

- on instances \mathcal{C} with $\text{Opt}(\mathcal{C}) \geq \beta$, the algorithm outputs YES;
- on instances \mathcal{C} with $\text{Opt}(\mathcal{C}) < \alpha$, the algorithm outputs NO.

Note that (α, β) -deciding $\text{CSP}(\Gamma)$ is strictly easier than (α, β) -approximating it. If we have an algorithm to (α, β) -approximate $\text{CSP}(\Gamma)$, then we can (α, β) -decide it by doing the following:

Given an instance \mathcal{C} , run the (α, β) -approximation algorithm on \mathcal{C} to get a solution \mathcal{S} . If $\text{Val}(\mathcal{S}) \geq \alpha$, output YES; if $\text{Val}(\mathcal{S}) < \alpha$, output NO.

This is an (α, β) -decision algorithm: If $\text{Opt}(\mathcal{C}) \geq \beta$, then $\text{Val}(\mathcal{S}) \geq \alpha$, so we output YES. If $\text{Opt}(\mathcal{C}) < \alpha$, then $\text{Val}(\mathcal{S}) \leq \text{Opt}(\mathcal{C}) < \alpha$, so we output NO. On the other hand, there is no way to make an (α, β) -approximation algorithm from an (α, β) -decision algorithm, because if $\text{Opt}(\mathcal{C}) \geq \beta$ then the (α, β) -decision algorithm will just say YES and will give no useful information.

A potentially good (α, β) -decision algorithm for $\text{CSP}(\Gamma)$ is the following:

- (\star) Given an instance \mathcal{C} , solve the canonical SDP relaxation. If $\text{SDPOpt} \geq \beta$, output YES; if $\text{SDPOpt} < \beta$, output NO.

Certainly if $\text{Opt}(\mathcal{C}) \geq \beta$, then $\text{SDPOpt}(\mathcal{C}) \geq \text{Opt}(\mathcal{C}) \geq \beta$, so this algorithm outputs YES. If we additionally know that whenever $\text{Opt}(\mathcal{C}) < \alpha$ we have $\text{SDPOpt}(\mathcal{C}) < \beta$, then this algorithm will correctly output NO whenever $\text{Opt}(\mathcal{C}) < \alpha$; however, if there is an instance \mathcal{C} with $\text{Opt}(\mathcal{C}) < \alpha$ but $\text{SDPOpt}(\mathcal{C}) \geq \beta$, then this algorithm will incorrectly output YES for that instance.

Definition 15.11. An instance \mathcal{C} of $\text{CSP}(\Gamma)$ is an (α, β) -SDP-gap instance if $\text{SDPOpt}(\mathcal{C}) \geq \beta$ but $\text{Opt}(\mathcal{C}) < \alpha$.

As observed above, the existence of such an instance is a barrier to the SDP algorithm (\star) being an (α, β) -decision algorithm. In other words, there exists an (α, β) -SDP-gap instance if and only if the algorithm (\star) is not an (α, β) -decision algorithm.

Definition 15.12.

$$\text{SDPGap}_\Gamma(\beta) = \inf\{\text{Opt}(\mathcal{C}) \mid \mathcal{C} \text{ is an instance of } \text{CSP}(\Gamma) \text{ with } \text{SDPOpt}(\mathcal{C}) \geq \beta\}.$$

Remark 15.13. Tautologically, for every β , the SDP algorithm (\star) is an $(\text{SDPGap}_\Gamma(\beta), \beta)$ -decision algorithm.

Theorem 15.14. *[Rag08], [RS09] [Are these the right references?] Assume the Unique Games Conjecture. Then for every $\epsilon > 0$, $(\text{SDPGap}_\Gamma(\beta - \epsilon) + \epsilon, \beta - \epsilon)$ -deciding $\text{CSP}(\Gamma)$ is NP-hard.*

Essentially, doing better than the algorithm (\star) is NP-hard. Also, in particular, we see that $(\text{SDPGap}_\Gamma(\beta - \epsilon) + \epsilon, \beta - \epsilon)$ -approximating is NP-hard. However, the following theorem says that this is the exact threshold, because $(\text{SDPGap}_\Gamma(\beta - \epsilon) - \epsilon, \beta)$ -approximating is in P.

Theorem 15.15. *[Rag08], [RS09] [Are these the right references?] For every constraint satisfaction problem $\text{CSP}(\Gamma)$, for every $\epsilon > 0$, and for every β , there exists an SDP rounding algorithm which $(\text{SDPGap}_\Gamma(\beta - \epsilon) - \epsilon, \beta)$ -approximates the CSP and is $\text{poly}(n)$ -time. [However, it is $\exp(\exp(\text{poly}(kq/\epsilon)))$ -time, where k is the maximum arity of the relations in the constraints and q is the size of the domain.]*

Proof. *[This needs to be cleaned up and more fully explained.]* An outline of the algorithm:

1. Given \mathcal{C} with $\text{Opt}(\mathcal{C}) \geq \beta$, find an optimal SDP solution \mathcal{S} .
2. Drop an ϵ -fraction of the constraints to get \mathcal{C}' , and get an SDP vector solution \mathcal{S}' for \mathcal{C}' with

$$\text{SDPVal}_{\mathcal{C}'}(\mathcal{S}') \geq \text{SDPVal}(\mathcal{S}) - O(\epsilon) = \text{SDPOpt}(\mathcal{C}) - O(\epsilon) \geq \text{Opt}(\mathcal{C}) - O(\epsilon) \geq \beta - O(\epsilon),$$
 with vectors in \mathbb{R}^d , where $d = \text{poly}(kq/\epsilon)$.
3. Losing another $O(\epsilon)$ on $\text{SDPVal}(\mathcal{S}')$, discretize the coordinates of these vectors to $\text{poly}(\epsilon/kq)$. (From now on, \mathcal{S}' will refer to these discretized vectors.) Now for every $v \in V$, there are only $\exp(\text{poly}(kq/\epsilon))$ many possible vectors $\vec{y}_{v,\ell} \in \mathcal{S}'$.
4. For any $v, v' \in \mathcal{C}'$ with identical vectors $(\vec{y}_{v,\ell})_{\ell \in D}$, identify them, forming \mathcal{C}'' , on $\exp(\text{poly}(kq/\epsilon))$ many variables. [There is a potential problem here: Is \mathcal{C}'' really an instance of $\text{CSP}(\Gamma)$? There is a fix for this problem.]
5. Solve \mathcal{C}'' in time $\exp(\exp(\text{poly}(kq/\epsilon)))$ to get $F'' : V'' \rightarrow D$ with $\text{Val}_{\mathcal{C}''}(F'') = \text{Opt}(\mathcal{C}'')$. Unfold this to $F' : V \rightarrow D$ for \mathcal{C}' . Then $\text{Val}_{\mathcal{C}'}(F') = \text{Opt}(\mathcal{C}'')$, so $\text{Val}_{\mathcal{C}}(F') \geq \text{Opt}(\mathcal{C}'') - O(\epsilon)$. Hence

$$\text{SDPOpt}(\mathcal{C}'') \geq \text{SDPVal}_{\mathcal{C}''}(\mathcal{S}') = \text{SDPVal}_{\mathcal{C}'}(\mathcal{S}') \geq \text{SDPVal}_{\mathcal{C}}(\mathcal{S}) - O(\epsilon) \geq \dots \geq \beta - \epsilon.$$

Therefore, by definition, $\text{Opt}(\mathcal{C}'') \geq \text{SDPGap}_\Gamma(\beta - \epsilon)$. □

Bibliography

- [ABS10] Sanjeev Arora, Boaz Barak, and David Steurer. Subexponential algorithms for Unique Games and related problems. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, pages 563–572, 2010. [15.3](#)
- [BRS11] Boaz Barak, Prasad Raghavendra, and David Steurer. Rounding semidefinite programming hierarchies via global correlation. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science*, 2011. [15.3](#)
- [CMM06] Moses Charikar, Konstantin Makarychev, and Yury Makarychev. Near-optimal algorithms for Unique Games. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 205–214, 2006. [15.3](#)
- [Kho02] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 767–775. ACM Press, 2002. [15.7](#)
- [Rag08] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pages 245–254, 2008. [15.14](#), [15.15](#)
- [RS09] P. Raghavendra and D. Steurer. Integrality gaps for strong SDP relaxations of Unique Games. In *Foundations of Computer Science, 2009. FOCS '09. 50th Annual IEEE Symposium on*, pages 575–585, oct. 2009. [15.14](#), [15.15](#)
- [She99] W. F. Sheppard. On the application of the theory of error to cases of normal distribution and normal correlation. *Royal Society of London Philosophical Transactions Series A*, 192:101–167, 1899. [15.2](#)