Semidefinite Programming

Lecture 9

Outline

From linear to semidefinite programming

Max-Cut and a random cut

The Goemans–Williamson SDP

Near-bipartite graphs

Planted clique and the Lovász theta function

Outline

From linear to semidefinite programming

Max-Cut and a random cut

The Goemans–Williamson SDP

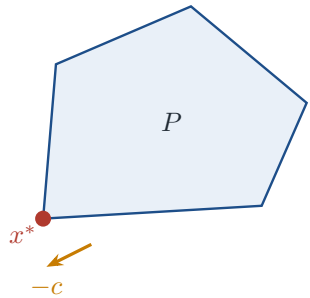
Near-bipartite graphs

Planted clique and the Lovász theta function

Last lecture: linear programming

$$\begin{aligned} \text{minimize} \quad & \langle c, x \rangle \\ \text{s.t.} \quad & \langle a_i, x \rangle \geq b_i \quad \forall i \end{aligned}$$

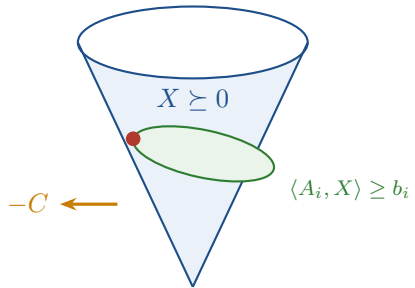
- $x \in \mathbb{R}^d$, $c, a_i \in \mathbb{R}^d$, $b_i \in \mathbb{R}$
- feasible set: a polyhedron
- solvable in polynomial time



Semidefinite programming: the same, with matrices

$$\begin{aligned} \text{minimize} \quad & \langle C, X \rangle \\ \text{s.t.} \quad & \langle A_i, X \rangle \geq b_i \quad \forall i \\ & X \succeq 0 \end{aligned}$$

- X, C, A_i symmetric $d \times d$
- $\langle C, X \rangle = \sum_{i,j} C_{ij} X_{ij} = \operatorname{tr}(CX)$
- linear objective, linear constraints, plus the PSD cone $X \succeq 0$



feasible set: a *spectrahedron* — convex, but curved

Two facts we will use

Gram matrices

$$X \succeq 0 \quad \iff \quad X = V^\top V \quad \iff \quad X_{ij} = \langle v_i, v_j \rangle \text{ for some vectors } v_1, \ldots, v_d$$

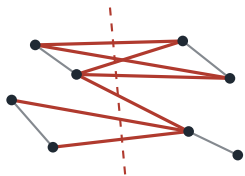
Spectral theorem: $X = U^\top \Lambda U$ with $\Lambda \succeq 0$; take $V = \sqrt{\Lambda}\, U$ and let v_i be its columns.

Efficiency

SDPs can be solved to accuracy ε in time polynomial in the input size and $\log(1/\varepsilon)$.

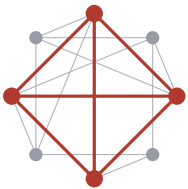
Ellipsoid method: a violated constraint for $X \not\succeq 0$ is an eigenvector of a negative eigenvalue.

Today: two examples



Max-Cut

Goemans–Williamson: 0.878-approximation



Planted clique

Lovász theta: a certificate robust to a monotone adversary

Outline

From linear to semidefinite programming

Max-Cut and a random cut

The Goemans–Williamson SDP

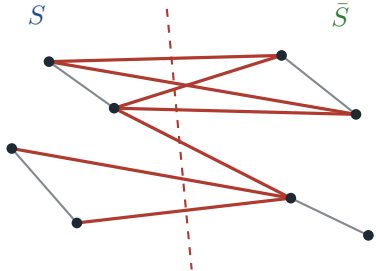
Near-bipartite graphs

Planted clique and the Lovász theta function

The Max-Cut problem

Find a partition $(S, \bar{S})$ of V with as many crossing edges as possible:

$$\mathrm{OPT} = \max_{S \subseteq V} |E(S, \bar{S})|$$



$|E(S, \bar{S})| = 7$ of $|E| = 11$

- **multiplicative:** find a cut of size $\geq \alpha \cdot \mathrm{OPT}$ — how large can α be?
- **near-bipartite:** if $\mathrm{OPT} \geq (1 - \varepsilon)|E|$, find a cut of size $\geq (1 - f(\varepsilon))|E|$

A random cut already gives $\frac{1}{2}$

Put each vertex into S independently with probability $\frac{1}{2}$:

$$\mathbb{E}\,|E(S,\bar{S})| = \sum_{e \in E} \Pr[e \text{ is cut}] = \frac{|E|}{2}$$

whatever happens to the first endpoint, the second one lands on the other side w.p. $\frac{1}{2}$

$$\mathbb{E}\,|E(S,\bar{S})| = \frac{|E|}{2} \;\geq\; \frac{\mathrm{OPT}}{2} \qquad (\text{since } \mathrm{OPT} \leq |E|)$$

Derandomization: conditional expectations

Fix an order $v_1, \dots, v_n$ and decide the sides one by one:

$$\mathbb{E}\big[|E(S,\bar{S})|\big] = \tfrac{1}{2}\,\mathbb{E}\big[|E(S,\bar{S})| \,\big|\, v_1 \in S\big] + \tfrac{1}{2}\,\mathbb{E}\big[|E(S,\bar{S})| \,\big|\, v_1 \in \bar{S}\big]$$

- the larger of the two branches is $\geq \mathbb{E}|E(S,\bar{S})|$: put v_1 there
- each conditional expectation is computable in $O(|E|)$ time
- repeat for $v_2, v_3, \dots$: the expectation never drops

Deterministic cut with at least $|E|/2 \;\geq\; \mathrm{OPT}/2$ edges.

Outline

From linear to semidefinite programming

Max-Cut and a random cut

The Goemans–Williamson SDP

Near-bipartite graphs

Planted clique and the Lovász theta function

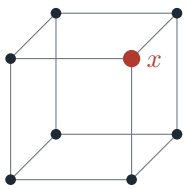
Max-Cut as a quadratic program

Encode a cut by $x \in \{\pm 1\}^n$: $x_i = -1$ for $i \in S$, $+1$ otherwise.

An edge $\{i, j\}$ is cut $\iff \dfrac{1 - x_i x_j}{2} = 1$:

$$\mathrm{OPT} = \max_{x \in \{\pm 1\}^n} \sum_{\{i,j\} \in E} \frac{1 - x_i x_j}{2}$$

maximizing a quadratic over the hypercube: NP-hard



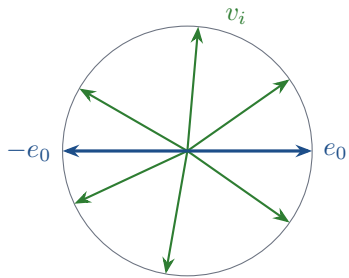
$\{\pm 1\}^n$

Relaxation: unit vectors instead of signs

$$\mathrm{SDP} := \max_{v_1,\dots,v_n \in S^{n-1}} \sum_{\{i,j\}\in E} \frac{1 - \langle v_i, v_j \rangle}{2}$$

It is a relaxation: given signs x_i, fix any $e_0 \in S^{n-1}$ and put $v_i := x_i\, e_0$. Then

$$\langle v_i, v_j \rangle = x_i x_j \quad \Longrightarrow \quad \mathrm{SDP} \;\geq\; \mathrm{OPT}.$$



signs ± 1: two antipodal points;
the relaxation allows all unit vectors

The vector program is an SDP

- vectors $\to$ matrix: $X_{ij} := \langle v_i, v_j \rangle$ gives $X \succeq 0$, $X_{ii} = 1$
- matrix $\to$ vectors: $X \succeq 0$ gives $X = V^\top V$; the columns v_i are unit vectors

$$\begin{aligned} \max \quad & \frac{|E|}{2} - \frac{\langle A, X \rangle}{4} \\ \text{s.t.} \quad & X_{ii} = 1 \quad \forall i \\ & X \succeq 0 \end{aligned}$$

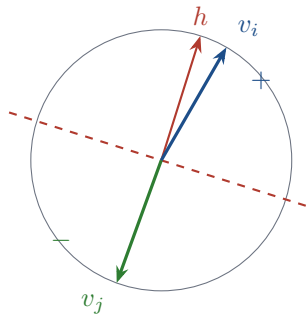
A: adjacency matrix, $\langle A, X \rangle = 2 \sum_{\{i,j\} \in E} X_{ij}$; equivalently $\min \langle A, X \rangle$ s.t. $X_{ii} = 1$, $X \succeq 0$.

Rounding: cut the sphere by a random hyperplane

Pick $h \in S^{n-1}$ uniformly and set

$$x_i := \operatorname{sgn}\langle v_i, h\rangle.$$

The hyperplane $h^\perp$ splits the sphere in two; $S := \{i : x_i = -1\}$.

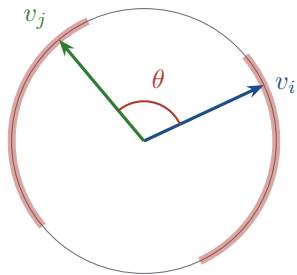


One edge at a time

Let $\theta_{ij} := \angle(v_i, v_j)$.

- **rounding:** $\Pr[\{i,j\} \text{ cut}] = \dfrac{\theta_{ij}}{\pi}$
- **SDP:**
  $$\frac{1 - \langle v_i, v_j \rangle}{2} = \frac{1 - \cos\theta_{ij}}{2} = \sin^2 \frac{\theta_{ij}}{2}$$

everything happens in the plane spanned by v_i, v_j



shaded: the h that separate v_i, v_j
(two arcs, total measure 2θ)

The Goemans–Williamson constant

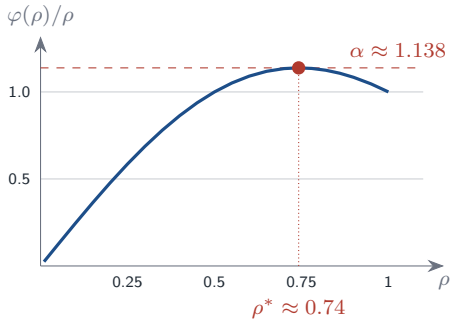
Write $\rho := \theta/\pi \in (0, 1]$ and

$$\varphi(\rho) := \sin^2\Big(\frac{\rho\pi}{2}\Big)$$

- rounding gains ρ
- the SDP charged $\varphi(\rho)$

$$\alpha := \max_{\rho\in(0,1]} \frac{\varphi(\rho)}{\rho} \approx 1.1382$$

$$\Rightarrow \quad \rho \;\geq\; \alpha^{-1}\varphi(\rho)$$



Putting it together

$$\mathbb{E}\big[|E(S,\bar{S})|\big] \underset{\text{per edge}}{=} \sum_{\{i,j\}\in E} \rho_{ij} \underset{\rho \geq \alpha^{-1}\varphi(\rho)}{\geq} \alpha^{-1} \sum_{\{i,j\}\in E} \varphi(\rho_{ij})$$

$$\underset{\text{SDP value}}{=} \alpha^{-1}\,\mathrm{SDP} \underset{\text{relaxation}}{\geq} \alpha^{-1}\,\mathrm{OPT}$$

Theorem (Goemans–Williamson 1995)

Solving the SDP and rounding with a random hyperplane gives a cut of expected size at least $\alpha^{-1}\mathrm{OPT} \approx 0.878 \cdot \mathrm{OPT}$.

Optimal among polynomial-time algorithms, assuming the Unique Games Conjecture.

Outline

From linear to semidefinite programming

Max-Cut and a random cut

The Goemans–Williamson SDP

Near-bipartite graphs

Planted clique and the Lovász theta function

The $(1-\varepsilon)$ regime

Assume the promise $\mathrm{OPT} \geq (1-\varepsilon)|E|$. Since the SDP is a relaxation,

$$\frac{1}{|E|} \sum_{\{i,j\} \in E} \frac{1 - \langle v_i, v_j \rangle}{2} \;\geq\; \frac{\mathrm{OPT}}{|E|} \;\geq\; 1 - \varepsilon.$$

With $\gamma_{ij} := 1 + \langle v_i, v_j \rangle \geq 0$ this says exactly

$$\frac{1}{|E|} \sum_{\{i,j\} \in E} \gamma_{ij} \;\leq\; 2\varepsilon$$

so on average the endpoints of an edge are *nearly antipodal*

Uncut edges: how close are v_i and $-v_j$?

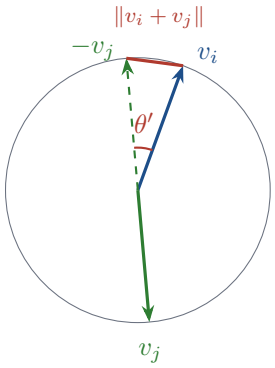
Let $\rho'_{ij} := \dfrac{\angle(v_i, -v_j)}{\pi}$, so

$$\Pr[\{i,j\} \textit{ not } \text{cut}] = \rho'_{ij}.$$

Chord versus angle:

$$\|v_i + v_j\| = \sqrt{2 + 2\langle v_i, v_j\rangle} = \sqrt{2\gamma_{ij}}$$

$$\rho'_{ij} \;\le\; \frac{\|v_i + v_j\|}{2} \;=\; \sqrt{\frac{\gamma_{ij}}{2}}$$



$\theta' \le \frac{\pi}{2}\|v_i + v_j\|$ for $\theta' \in [0, \pi]$

Averaging: the uncut fraction is at most $\sqrt{\varepsilon}$

$$\frac{1}{|E|}\sum_{\{i,j\}\in E}\rho'_{ij} \underset{\text{previous slide}}{\leq} \frac{1}{|E|}\sum_{\{i,j\}\in E}\sqrt{\frac{\gamma_{ij}}{2}} \underset{\text{concavity of }\sqrt{\cdot}}{\leq} \sqrt{\frac{1}{2|E|}\sum_{\{i,j\}\in E}\gamma_{ij}} \underset{\text{promise}}{\leq} \sqrt{\varepsilon}$$

Theorem

If $\mathrm{OPT} \geq (1-\varepsilon)|E|$, the same SDP and rounding return a cut of expected size at least $(1-\sqrt{\varepsilon})\,|E|$.

$1 - O(\sqrt{\varepsilon})$ is again optimal for polynomial-time algorithms under the Unique Games Conjecture.

Outline

From linear to semidefinite programming

Max-Cut and a random cut

The Goemans–Williamson SDP

Near-bipartite graphs

Planted clique and the Lovász theta function

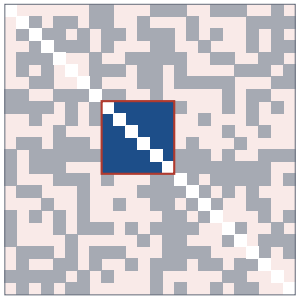
Recall: planted clique (Lecture 6)

Distinguish $G(n, \frac{1}{2})$ from $G(n, \frac{1}{2})$ with a planted k-clique.

Signed adjacency matrix A_G: $+1$ on edges, -1 on non-edges, 0 on the diagonal.

- clique S, $|S| = k$: $\mathbf{1}_S^\top A_G \mathbf{1}_S = k(k-1)$, so $\lambda_{\max}(A_G) \geq k - 1$
- pure random: $\|A_G\| \lesssim \sqrt{n}$ w.h.p.

distinguishes when $k \gtrsim \sqrt{n}$



A_G with a planted 6-clique:
an all-$(+1)$ block hidden in noise

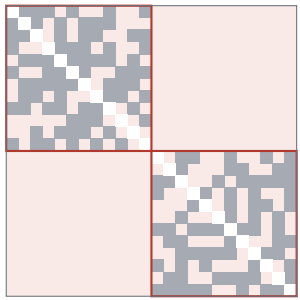
A "helpful" adversary breaks the spectral test

Semirandom model: after sampling, an adversary (unbounded, knows the algorithm) may

- **planted case:** add edges
- **null case:** delete edges

✓ the k-clique survives: $\lambda_{\max}(A_G) \geq k - 1$

× deleting all edges between two halves makes $\lambda_{\max}(A_G) = \Omega(n)$



null case after deletions: two all-(-1) off-diagonal blocks, $v = \frac{1}{\sqrt{n}}(\mathbf{1}, -\mathbf{1})$ gives $v^\top A_G v \approx n/2$

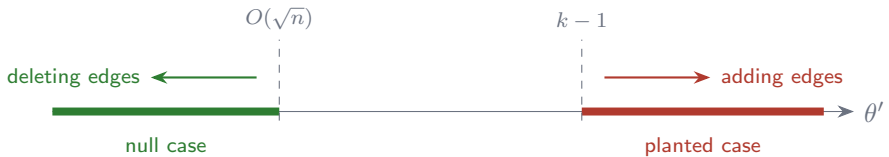
Wanted: a monotone certificate

A polynomial-time $\theta'(G)$ with

(1) G has a k-clique $\Rightarrow$ $\theta'(G) \geq k-1$

(2) $G_1 \subseteq G_2 \Rightarrow \theta'(G_1) \leq \theta'(G_2)$ (monotone)

(3) $\theta'(G) \leq \lambda_{\max}(A_G)$, hence $O(\sqrt{n})$ on $G(n, \frac{1}{2})$ w.h.p.



the gap survives the adversary: the two moves push θ' the *safe* way

The certificate: minimize the top eigenvalue

$$\begin{aligned} \theta'(G) := \quad & \min \quad \lambda_{\max}(B) \\ & \text{s.t.} \quad B_{ij} = 1 \quad \forall \{i,j\} \in E \\ & \qquad B_{ii} = 0 \quad \forall i \end{aligned}$$

B symmetric; entries on non-edges are free

0	1	1	?	?	?
1	0	1	?	?	1
1	1	0	1	?	?
?	?	1	0	1	?
?	?	?	1	0	1
?	1	?	?	1	0

the pattern of B

(1) Cliques are detected

Let S be a k-clique and B any feasible matrix. All pairs inside S are edges, so

$$\mathbf{1}_S^\top B\,\mathbf{1}_S = \sum_{i\neq j\in S} \underbrace{B_{ij}}_{=1} + \sum_{i\in S} \underbrace{B_{ii}}_{=0} = k(k-1)$$

$$\lambda_{\max}(B) \;\geq\; \frac{\mathbf{1}_S^\top B \mathbf{1}_S}{\|\mathbf{1}_S\|^2} \;=\; \frac{k(k-1)}{k} \;=\; k-1 \qquad \Longrightarrow \qquad \theta'(G) \;\geq\; k-1$$

(2) Monotone and (3) small on random graphs

(2) Monotonicity

More edges = more constraints = a *smaller* feasible set, so the minimum can only grow:

$$G_1 \subseteq G_2 \quad \Longrightarrow \quad \theta'(G_1) \leq \theta'(G_2).$$

(3) Random graphs

A_G is feasible ($+1$ on edges, 0 on the diagonal), hence

$$\theta'(G) \;\leq\; \lambda_{\max}(A_G) \;\leq\; \|A_G\| \;\lesssim\; \sqrt{n} \quad \text{w.h.p. for } G \sim G(n, \tfrac{1}{2}).$$

Why is this an SDP?

Feasibility form: is there B in the polyhedron
$P(G) = \{B \text{ symmetric} : B_{ij} = 1 \text{ on edges},\ B_{ii} = 0\}$ with $\lambda_{\max}(B) \le \gamma$?

$$\lambda_{\max}(B) \le \gamma \quad \Longleftrightarrow \quad \gamma I - B \succeq 0. \qquad \text{Put } X := \gamma I - B.$$

find $X \succeq 0$ with

$$X_{ij} = -1 \quad \forall \{i,j\} \in E$$
$$X_{ii} = \gamma \quad \forall i$$

or, as an optimization,

$$\begin{aligned} \min \quad & X_{11} \\ \text{s.t.} \quad & X_{ij} = -1 \quad \forall \{i,j\} \in E \\ & X_{ii} = X_{11} \quad \forall i, \qquad X \succeq 0 \end{aligned}$$

This is the Lovász theta function

$$\theta'(G) = \vartheta(\bar{G}) - 1 \qquad (\vartheta\text{: Lovász theta, } \bar{G}\text{: complement})$$

- $\omega(G) \le \vartheta(\bar{G})$: the certificate upper-bounds the clique number
- $\mathbb{E}\,\vartheta\big(G(n, \frac{1}{2})\big) = \Theta(\sqrt{n})$
- monotone $\Rightarrow$ the planted-clique algorithm survives the adversary for $k \gtrsim \sqrt{n}$

Summary

- **SDP** = linear objective and constraints over the PSD cone; $X \succeq 0 \iff X$ is a Gram matrix
- **Max-Cut:** vectors on a sphere + random hyperplane: $0.878 \cdot \mathrm{OPT}$, and $(1 - \sqrt{\varepsilon})|E|$ when $\mathrm{OPT} \geq (1 - \varepsilon)|E|$
- **Planted clique:** $\theta'(G) = \vartheta(\bar{G}) - 1$ is a *monotone* certificate, robust to a semirandom adversary

Both are LP-style relaxation + rounding — with the PSD cone doing the geometric work.