Notes 1: Semidefinite Programs and Max-Cut

1. Semidefinite programs

Semidefinite programs generalize linear programs (LP). Recall that a linear program looks like the following:

$$(1) \qquad \max \quad 2x_1 + 3x_2 - 4x_3$$
$$5x_1 - 8x_2 + 4x_3 \leqslant 10$$
$$4x_1 + 3x_2 - x_3 \leqslant 5$$
$$x_1, x_2, x_3 \geqslant 0$$

More generally, a linear program (in canonical form) takes the form

$$\max \quad c^T x$$
$$a_1^T x \leqslant b_1$$
$$\vdots$$
$$a_m^T x \leqslant b_m$$
$$x \geqslant 0$$

where $x, c, a_1, \ldots, a_m \in \mathbb{R}^n$ are all n-dimensional real vectors, and $b_1, \ldots, b_m \in \mathbb{R}$ are real scalars. Here x represents our LP variables, c represents our linear objective function, and $a_1, \ldots, a_m$ are the linear constraints. The last inequality constraint $x \geqslant 0$ means that x has to be entry-wise nonnegative.

By contrast, a semidefinite program (SDP) looks like the following:

$$\max \quad \begin{pmatrix} 2 & 3/2 \\ 3/2 & -4 \end{pmatrix} \bullet \begin{pmatrix} x_1 & x_2 \\ x_2 & x_3 \end{pmatrix}$$
$$\begin{pmatrix} 5 & -4 \\ -4 & 4 \end{pmatrix} \bullet \begin{pmatrix} x_1 & x_2 \\ x_2 & x_3 \end{pmatrix} \leqslant 10$$
$$\begin{pmatrix} 4 & 3/2 \\ 3/2 & -1 \end{pmatrix} \bullet \begin{pmatrix} x_1 & x_2 \\ x_2 & x_3 \end{pmatrix} \leqslant 5$$
$$\begin{pmatrix} x_1 & x_2 \\ x_2 & x_3 \end{pmatrix} \succcurlyeq 0$$

Here $\bullet$ denotes the Frobenius/Hadamard inner product between two matrices, defined as the entry-wise inner product between two n-by-n matrices

$$A \bullet B \stackrel{\text{def}}{=} \sum_{1 \leqslant i,j \leqslant n} A_{ij} B_{ij} .$$

The above semidefinite program has exactly the same objective function as the linear program (1) above, because

$$\begin{pmatrix} 2 & 3/2 \\ 3/2 & -4 \end{pmatrix} \bullet \begin{pmatrix} x_1 & x_2 \\ x_2 & x_3 \end{pmatrix} = 2x_1 + 3x_2 - 4x_3 .$$

Compared to (1), the main difference is that the final nonnegative constraint is replaced with a *positive semidefinite* constraint, as defined now.

Definition 1.1. A symmetric n-by-n matrix M is *positive semidefinite* if for every $y \in \mathbb{R}^n$, the quadratic form $y^T M y \geqslant 0$.

A general semidefinite program takes the form

$$\max \quad C \bullet X$$
$$A_1 \bullet X \leqslant b_1$$
$$\vdots$$
$$A_m \bullet X \leqslant b_m$$
$$X \succcurlyeq 0$$

where $X, C, A_1, \ldots, A_m$ are all n-by-n real symmetric matrices, and $b_1, \ldots, b_m \in \mathbb{R}$ are real scalars. The matrix X represents our SDP variables.

2. Positive semidefiniteness

The positive semidefinite (PSD) condition has a number of equivalent definitions.

Proposition 2.1. *Given a symmetric n-by-n matrix M, the following are equivalent:*
  (a) For every $y \in \mathbb{R}^n$, we have $y^T M y \geqslant 0$
  (b) All eigenvalues of M are nonnegative
  (c) $X = U^T U$ for some n-by-n matrix U (U is not necessarily symmetric)

The condition (c) is equivalent to saying that there are n vectors $u_1, \ldots, u_n \in \mathbb{R}^n$ such that X encodes the inner products between them. More precisely, $X_{ij} = u_i^T u_j$, namely the ij-entry of X equals the inner product between the i- and the j-vectors. To see this, simply define $u_1, \ldots, u_n$ as the column vectors of U, and condition (c) becomes

To prove the proposition, we will use the following theorem.

Theorem 2.2 (Spectral theorem for real symmetric matrices)**.** *Any n-by-n real symmetric matrix M has n real eigenvalues $\lambda_1, \ldots, \lambda_n$ and n orthonormal eigenvectors $v_1, \ldots, v_n$. Equivalently, we can express any such an M as*

$$(2) \qquad\qquad M = V \Lambda V^T \, ,$$

where V is an n-by-n matrix whose columns are precisely the eigenvectors $v_1, \ldots, v_n$, and Λ is a diagonal matrix with eigenvalues $\lambda_1, \ldots, \lambda_n$ on its diagonal ($\Lambda_{ii} = \lambda_i$). Since the eigenvectors are orthonormal, we also have $V^T V = V V^T = I$. Decomposition (2) can be represented in picture as

or as a sum of outer products

$$M = \sum_{1 \leqslant i \leqslant n} \lambda_i v_i v_i^T \, .$$

Proof of the proposition. (a) $\Rightarrow$ (b): We prove the contrapositive; if (b) is violated, then so must (a). If M has a negative eigenvalue, say λ_i, we simply take y to be its corresponding eigenvector v_i, and the quadratic form becomes

$$y^T M y = v_i^T M v_i = \lambda_i < 0 \, ,$$

violating condition (a).

(b) $\Rightarrow$ (c): Let $\sqrt{\Lambda}$ be the diagonal matrix with $\sqrt{\lambda_1}, \ldots, \sqrt{\lambda_n}$ on its diagonal ($\sqrt{\Lambda}_{ii} = \sqrt{\lambda_i}$), and let $U = \sqrt{\Lambda} V^T$. Since M has only nonnegative eigenvalues (and they lie on the diagonal),

$\sqrt{\Lambda}$ has only real entries. Also $(\sqrt{\Lambda})^T = \sqrt{\Lambda}$ because it is a diagonal matrix. Then the spectral decomposition (2) becomes

$$M = V\sqrt{\Lambda}^T\sqrt{\Lambda}V^T = U^TU .$$

(c) $\Rightarrow$ (a): For any $y \in \mathbb{R}^n$,

$$y^T My = y^T U^T U y = \|Uy\|_2^2 \geqslant 0 . \qquad \square$$

3. LP vs SDP

Every linear program can be represented as a semidefinite program. For example, the linear program in (1) is equivalent to a semidefinite program with a 3-by-3 symmetric matrix X as its variables, and we may add additional constraints to force X to be diagonal (by making every off-diagonal entry to be zero). We can also make the objective matrix C diagonal, that is, the SDP objective function becomes

$$C \bullet X = \begin{pmatrix} 2 & & \\ & 3 & \\ & & -4 \end{pmatrix} \bullet \begin{pmatrix} x_1 & & \\ & x_2 & \\ & & x_3 \end{pmatrix} = 2x_1 + 3x_2 - 4x_3 ,$$

exactly the same as the linear program objective. The positive semidefinite constraint $X \succcurlyeq 0$ is equivalent to requiring all diagonal entries of X to be nonnegative, and is the same as the nonnegative constraints $x_1, x_2, x_3 \geqslant 0$ in the linear program (1).

In general, a semidefinite program reduces to a linear program when the objective matrix C and constraint matrices $A_1, \ldots, A_m$ are all diagonal matrices.

Recall that a linear program amounts to maximizing a linear objective function in a polyhedron (a region that is the intersection of halfspaces of the form $\{x \in \mathbb{R}^n \mid a_i^T x \leqslant b_i\}$).

A semidefinite program amounts to maximizing a linear objective function in a feasible region whose boundary may involve curves. For example, its feasible region may be the unit circle.

This can happen with the positive semidefinite constraint

(3)
$$\begin{pmatrix} 1 & x & y \\ x & 1 & 0 \\ y & 0 & 1 \end{pmatrix} \succcurlyeq 0$$

and certain matrix entries have been forced to be either 0 or 1 as indicated. Since a matrix is positive semidefinite precisely when all its principal minors are nonnegative (look up "Sylvester's criterion" on Wikipedia), (3) is equivalent to the following three constraints:

$$1 \geqslant 0$$
$$1 - x^2 \geqslant 0$$
$$1 - x^2 - y^2 \geqslant 0 .$$

The first constraint is trivial, and the last constraint dominates the second. This shows that (3) is equivalent to (x, y) being inside the unit circle.

If a linear program involves only rational entries, then its optimal solution must also have only rational entries. The same is not necessarily true for a semidefinite program, as the previous circle example demonstrates. For instance, if we maximize along the direction $(1, 1)$ over the unit circle, the optimal solution $(1/\sqrt{2}, 1/\sqrt{2})$ has irrational entries. We cannot hope to represent a SDP

optimal solution exactly, partly explaining why the algorithms for solving SDP we describe later can only look for an approximate solution.

Finally, if a linear program has finite optimum value, then the optimum value is achieved by some solution. This is again not always true for a semidefinite program. Here is a counterexample:

$$\min \quad x$$

$$\begin{pmatrix} x & 1 \\ 1 & y \end{pmatrix} \succcurlyeq 0$$

We are minimizing x, but it's equivalent to maximizing $-x$. The PSD constraint is satisfied when $x, y \geqslant 0$ and $xy \geqslant 1$. Equivalently, when $y > 0$ and $x \geqslant 1/y$. The SDP optimum value is zero, by choosing smaller and smaller $x > 0$, but $x = 0$ is not feasible. This shows that the "minimum" is in fact an infimum.

4. Max-Cut and Goemans–Williamson rounding

A classical problem is to find a subset $S \subseteq V(G)$ maximizing the number of edges going between S and $V(G) \setminus S$. This Max-Cut problem is known to be NP-hard to solve exactly, so we will settle for an approximation algorithm.

Theorem 4.1 (Goemans–Williamson [2]). *There is a polynomial-time algorithm to approximate Max-Cut to within* $0.878\ldots$

In other words, if the best cut in a graph with 10000 cross edges, then the algorithm will return a cut with at least 8780 cross edges. The algorithm is based on a semidefinite program first proposed by Delorme and Poljak [1]. The rounding algorithm came from a seminal work by Goemans and Williamson, where they interpret the SDP as a vector optimization program. To motivate the vector program, first formulate Max-Cut as the following quadratic integer program:

$$\max \quad \sum_{(u,v) \in E} \frac{1 - x_u x_v}{2}$$

$$x_u \in \{+1, -1\} \quad \text{for every } u \in V(G)$$

Given any candidate solution S to Max-Cut, if we assign -1 to vertices in S and $+1$ to vertices outside S, then the term $(1 - x_u x_v)/2$ will be either 1 or 0, depending on whether the edge (u, v) crosses the cut. Therefore the objective function measures exactly the total number of cut edges. We are assigning ± 1 rather than $\{0, 1\}$ because it will simplify the calculations to come.

The above quadratic integer program is equivalent to Max-Cut and is NP-hard to solve exactly, so we relax the problem by instead assigning a *vector* to every vertex. We arrive at the following vector program.

$$(4) \qquad\qquad \max \quad \sum_{(u,v) \in E} \frac{1 - \langle y_u, y_v \rangle}{2}$$

$$\|y_u\|_2^2 = 1 \quad \text{for every } u \in V(G)$$

$$y_u \in \mathbb{R}^n \quad \text{for every } u \in V(G)$$

The norm constraint $\|y_u\|_2^2 = 1$ says that we are assigning a unit vector to every vertex. Here the dimension n of the space containing the vectors equals the number of vertices of G.

Claim 4.2. *The vector program optimum value is always at least the Max-Cut optimum value.*

Proof. Pick any unit vector $y_0 \in \mathbb{R}^n$. Given a candidate solution S to Max-Cut, we assign $-y_0$ to vertices in S and y_0 to vertices outside S. Then the term $(1 - \langle y_u, y_v \rangle)/2$ will be either 1 or 0, depending on whether the edge (u, v) crosses the cut. Therefore there is always a vector assignment with value at least as large as the number of edges crossing S. $\square$

When y_u and y_v are both unit vectors, the objective function is a sum of terms of the form

$$(1 - \langle y_u, y_v \rangle)/2 = (\|y_u\|_2^2 - 2\langle y_u, y_v \rangle + \|y_v\|_2^2)/4 = \|y_u - y_v\|_2^2/4 \ .$$

Therefore the vector program is seeking a mapping of vertices to the high dimensional unit sphere that maximizes the sum of squared distances between endpoints of an edge.

The vector program (4) is equivalent to the following semidefinite program (that can be solved efficiently):

$$(5) \qquad \max \quad \sum_{(u,v)\in E} \frac{1}{4} L^{uv} \bullet Y$$
$$Y_{uu} = 1 \quad \text{for every } u \in V(G)$$
$$Y \succcurlyeq 0 \,,$$

where L^{uv} is the matrix that is zero everywhere, except that it is 1 at uu and vv entries, and is -1 at uv and vu entries. Recall that the PSD constraint means Y encodes the inner products between vectors $\{y_u\}_{u\in V(G)}$, so that $Y_{uv} = \langle y_u, y_v \rangle$. One can check that $L^{uv} \bullet Y/4 = \|y_u - y_v\|_2^2/4$, which equals $(1 - \langle y_u, y_v \rangle)/2$ for unit vectors y_u and y_v, so (5) and (4) have the same objective functions. The constraint $Y_{uu} = 1$ in the SDP (5) is identical to $\|y_u\|_2^2 = 1$ in the vector program (4). In general, any SDP is equivalent to a vector optimization problem with linear objective function and linear constraints on inner products.


FIGURE 1. Random hyperplane rounding (picture from [3])

Given a solution to the vector program, Goemans–Williamson proposed to randomly round them into a cut, by picking a random hyperplane, and putting all vectors on one side of the hyperplane to S. In other words, one pick a random unit vector $g \in \mathbb{R}^n$ (a normal vector to the hyperplane), and set $S = \{u \in V \mid \langle y_u, g \rangle \geqslant 0\}$. We now show that, in expectation, this random subset S cuts at least $0.878\ldots$ fraction of edges in the optimal cut.

Proposition 4.3.
$$\mathbb{E}[\text{number of edges cut}] \geqslant 0.878\cdots \times \text{SDP-OPT} \,,$$
*where* SDP-OPT *denotes the optimum value of the SDP (or the vector program).*

Note that the proposition together with Claim 4.2 implies Theorem 4.1 (at least in expectation).

Proof of proposition. Let $\{y_u\}_{u\in V}$ be an optimal vector assignment to (4). Expand
$$\mathbb{E}[\text{number of edges cut}] = \sum_{(u,v)\in E} \mathbb{E}[\mathbb{1}((u,v) \text{ is cut by } S)]$$

by linearity of expectation. Focusing on each summand, each edge (u,v) is cut with probability θ_{uv}/π, where θ_{uv} is the angle between y_u and y_v. In other words, $\theta_{uv} = \arccos\langle y_u, y_v \rangle$.

On the other hand, each edge (u,v) contributes $(1 - \langle y_u, y_v \rangle)/2 = (1 - \cos\theta_{uv})/2$ to the vector program. We now compare the contribution to the cutting probability and the contribution to the vector program.

Lemma 4.4 (Goemans–Williamson). *For every $0 \leqslant \theta \leqslant \pi$,*

$$\frac{\theta}{\pi} \geqslant 0.878\ldots \frac{1 - \cos\theta}{2}.$$

Therefore

$$\mathbb{E}[\text{number of edges cut}] = \sum_{(u,v) \in E} \mathbb{P}[(u,v) \text{ is cut by } S] = \sum_{(u,v) \in E} \frac{\theta_{uv}}{\pi}$$

$$\geqslant 0.878 \cdots \sum_{(u,v) \in E} \frac{1 - \langle y_u, y_v \rangle}{2} = 0.878 \ldots \text{SDP-OPT}. \qquad \square$$

We remark that Goemans–Williamson algorithm also works with the weighted version of the problem, where each edge (u, v) has a nonnegative weight w_{uv}, and we seek to maximizing the sum of weights of edges crossing the cut.

REFERENCES

[1] Charles Delorme and Svatopluk Poljak. Laplacian eigenvalues and the maximum cut problem. *Mathematical Programming*, 62(3, Ser. A):557–574, 1993.

[2] Michel Xavier Goemans and David Paul Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, November 1995.

[3] László Lovász. Semidefinite programs and combinatorial optimization. In *Recent advances in algorithms and combinatorics*, volume 11 of *CMS Books Math./Ouvrages Math. SMC*, pages 137–194. Springer, New York, 2003.