APPROXIMATING SPARSE SEMIDEFINITE PROGRAMS

Kevin Shu[1]

Joint with Grigoriy Blekherman[1], based on [1]

[1]Georgia Institute of Technology, Department of Mathematics

Introduction

Can we approximate sparse semidefinite programs using fewer variables?
Semidefinite programs are convex optimization problems of the following form:

$$\begin{aligned} \text{minimize} \quad & \langle B^0, X\rangle \\ \text{such that} \quad & \langle B^1, X\rangle = b_1 \\ & \cdots \\ & \langle B^k, X\rangle = b_k \\ & X \succeq 0 \end{aligned}$$

$X \succeq 0$ means that X is positive semidefinite (PSD); all eigenvalues of X are nonnegative.

These programs are often used to approximate polynomial optimization problems, with applications in control theory, power systems, and combinatorial optimization.[2]

General semidefinite programs are very memory intensive, so it is of interest to reduce the number of variables involved in the program. We say that the above semidefinite program is G-sparse for a graph G if each B^ℓ has the property that $B^\ell_{ij} = 0$ if $\{i, j\}$ is not an edge of G. For G-sparse programs, the variables corresponding to nonedges of G are not used except to define the positive semidefiniteness condition.

G-sparse programs appear in natural problems. For example, there is a well known semidefinite program of Goemans-Williamson for approximating the NP-complete MAX-CUT problem, which is G-sparse for input graphs G.

Our aim is to approximately solve such sparse semidefinite programs while only keeping track of the variables needed to define these linear constraints and the linear objective. We do this by considering the cone of PSD-completable matrices, and a well known relaxation which we call the locally-PSD cone.

Positive Semidefinite Completion

For any graph G, there is a vector space of G-partial matrices, denoted

$$\mathcal{H}(G) = \mathbb{R}^{V(G) \cup E(G)}$$

We can project an $n \times n$ matrix, X, onto the space of partial matrices, by 'forgetting' the entries of X which correspond to nonedges of G.

$$\begin{pmatrix} 1 & -1 & x & -1 \\ -1 & 1 & -1 & y \\ x & -1 & 1 & -1 \\ -1 & y & -1 & 1 \end{pmatrix} \to \begin{pmatrix} 1 & -1 & ? & -1 \\ -1 & 1 & -1 & ? \\ ? & -1 & 1 & -1 \\ -1 & ? & -1 & 1 \end{pmatrix}$$

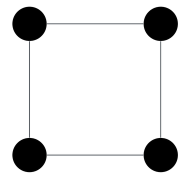

This corresponds to the natural projection onto partial matrices for a 4 cycle.

The projection of the PSD cone onto the space of G-partial matrices will be referred to as the PSD completable cone.

$$\Sigma(G) = \{X \in \mathcal{H}(G) : \exists X' \succeq 0 : \pi(X') = X\}$$

The positive semidefinite completion problem is to determine whether a G-partial matrix has a PSD completion. This has been studied extensively, and results are known for chordal graphs and series parallel graphs. [3, 4] Solving G-sparse semidefinite program is equivalent to solving conical programs over $\Sigma(G)$.

Locally-PSD Matrices

Cliques in the graph G correspond to complete submatrices of G-partial matrices:

$$\begin{pmatrix} 1 & -1 & 1 & -1 \\ -1 & 1 & -1 & ? \\ 1 & -1 & 1 & -1 \\ -1 & ? & -1 & 1 \end{pmatrix}$$

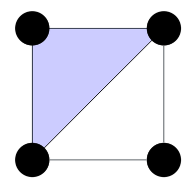

A locally-PSD matrix with one fully specified submatrix and its corresponding clique highlighted. This submatrix is PSD.

Principal submatrices of PSD matrices are PSD, so for a G-partial matrix to be PSD completable, all of the submatrices corresponding to cliques must be PSD.

$$\mathcal{P}(G) = \{X \in \mathcal{H}(G) : \forall \text{ cliques } C \subseteq V(G), X|_C \succeq 0\}$$

The locally-PSD cone is the cone of matrices where for all cliques of G, the corresponding submatrices are PSD. It is clear that $\Sigma(G) \subseteq \mathcal{P}(G)$.

If the number of cliques of G is small, then checking that a partial matrix is locally-PSD is easy without introducing auxiliary variables.

Remark: More generally, we might consider G-partial matrices where only a subset of the cliques correspond to PSD matrices. This leads to the more general setting of considering simplicial complexes instead of graphs, and will be explored more in other work.

A Quantitative Measure

Given $X \in \mathcal{P}(G)$, we define the additive distance of X to be

$$\epsilon_G(X) = \min\{\epsilon : X + \epsilon I \in \Sigma(G)\}$$

- This is the smallest ϵ so that $X + \epsilon I$ is PSD completable.
- This is related to the largest minimum eigenvalue of any completion of X.
- It is concave on $\mathcal{P}(G)$.
- For fixed X, this can be computed using a semidefinite program.

We then say that the additive distance of the graph G is

$$\epsilon(G) = \min\{\epsilon_G(X) : X \in \mathcal{P}(G),\ \operatorname{tr}(X) = 1\}$$

We can find lower bounds on this quantity by finding completions for matrices in $\mathcal{P}(G)$ whose minimum eigenvalue is close to 0.

PSD-Completable v.s. Locally-PSD

It is well known that if the graph G is **chordal**, then $X \in \mathcal{H}(G)$ is PSD-completable if and only if it is locally-PSD, so that $\epsilon(G) = 0$. On the other hand for all nonchordal graphs, $\epsilon(G) > 0$.[5] This is related to the fact that there exist nonnegative polynomials on some varieties which are not sum-of-squares.[6]

It was shown in [7] that if G is series parallel, then $X \in \mathcal{P}(G)$ is PSD completable if and only if for all cycles $C \subseteq G$, $X|_C$ is PSD completable. This is a fact we exploit to reduce the computation of $\epsilon(G)$ for G series parallel to the case when G is a cycle.

Results

Approximation

This additive distance is useful for bounding the quality of the approximation resulting from replacing the PSD cone with the locally-PSD cone for general programs, but the result is easiest to state when the program satisfies some conditions.

We say that a semidefinite program is of **Goemans-Williamson type** if it is G-sparse, the objective has 0 trace, the trace of any feasible point is at most 1, and if the rescaled identity $\frac{1}{n}I$ is a feasible point of the program.

Approximation Guarantee: If α is the value of a semidefinite of Goemans-Williamson type, and α' is the value of the program obtained by replacing the PSD cone with the locally-PSD cone, then

$$\alpha' \le \alpha \le \frac{1}{1 + n\epsilon(G)}\alpha'.$$

Additive Distance Computations

The **clique sum** of two graphs G_1 and G_2 is the result of gluing G_1 and G_2 along a common clique, and we denote it by $G_1 \oplus G_2$.

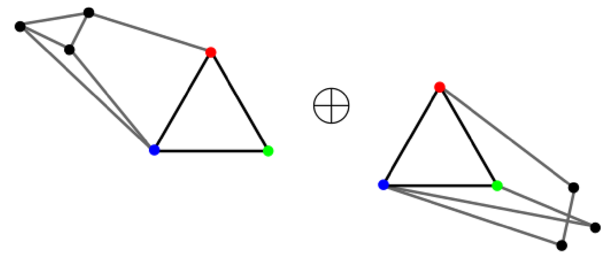

Clique Sume and Additive Distance:

$$\epsilon(G_1 \oplus G_2) = \max\{\epsilon(G_1), \epsilon(G_2)\}$$

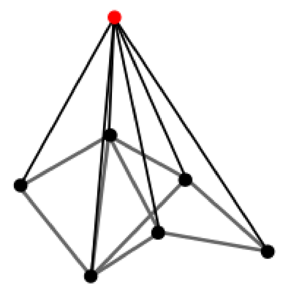

The **cone** over a graph G is the graph $\hat{G}$ obtained by introducing a new vertex to G and adding in all possible edges out of that new vertex.

Coning and Additive Distance:

$$\epsilon(\hat{G}) = \epsilon(G)$$

The **chordal girth** of a graph G, $g(G)$, is the size of the smallest induced cycle in G of size at least 4. A graph is series-parallel if has no K_4 minors.

Series Parallel Graphs: If G is a series-parallel graph, then

$$\epsilon(G) = \frac{1}{g(G)}\left(\frac{1}{\cos(\frac{\pi}{g(G)})} - 1\right) = O\left(\frac{1}{g(G)^3}\right)$$

We can also consider the ranks of optimal PSD completions:

If G is any graph, and $X \in \mathcal{P}(G) \setminus \Sigma(G)$, then $X + \epsilon_G(X)I$ can be completed to a rank $n - g(G) + 2$ PSD matrix.

Thickened Graphs

More recent results use ideas from algebraic geometry to characterize the extreme rays of $\mathcal{P}(G)$ for graphs which can be written as quotients of chordal graphs. We give a simplified example of this kind of result here.

A **thickened graph** is a graph H obtained from a **base graph**, G (which we will allow to have multi-edges and loops), and replacing the edges with chordal graphs with marked starting and terminal vertices.

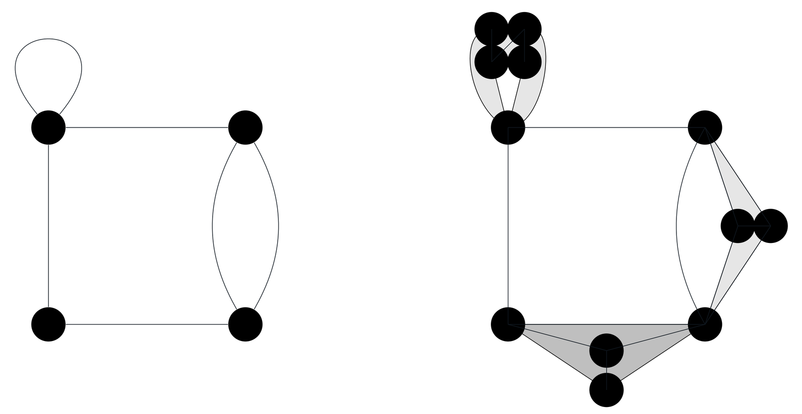

Thickened Graphs: If H is a thickened graph with a triangle free, simple, base graph G, then $\epsilon(H) \le \epsilon(G)$. (unpublished)

Thickened Cycles: Suppose that G is a chordal graph, and H is a graph obtained by identifying two vertices in G which are of distance at least 4 apart. Then $\epsilon(H) = \epsilon(C_g)$, where g is the chordal girth of H, and C_g is a cycle of length g. (unpublished)

Conclusions

It is possible to solve sparse semidefinite programs as long as the sparsity pattern has certain structural properties. It is often enough for the chordal girth of the graph to be $o(n^{\frac{1}{3}})$ to get a good approximation.

References

[1] Grigoriy Blekherman and Kevin Shu. *Sums of Squares and Sparse Semidefinite Programming*. 2020. arXiv: 2010.11311 [math.AG].

[2] Grigoriy Blekherman, Pablo A Parrilo, and Rekha R Thomas. *Semidefinite optimization and convex algebraic geometry*. SIAM, 2012.

[3] Wayne W. Barrett, Charles R. Johnson, and Raphael Loewy. "The real positive definite completion problem: cycle completability". In: *Mem. Amer. Math. Soc.* 122.584 (1996), pp. viii+69. ISSN: 0065-9266. DOI: 10.1090/memo/0584. URL: https://doi.org/10.1090/memo/0584.

[4] L. Vandenberghe and M.S. Andersen. *Chordal Graphs and Semidefinite Optimization*. Foundations and trends in optimization. Now Publishers, 2015. ISBN: 9781680830392. URL: https://books.google.com/books?id=IpAVjwEACAAJ.

[5] Robert Grone et al. "Positive definite completions of partial Hermitian matrices". In: *Linear Algebra and its Applications* 58 (1984), pp. 109 –124. ISSN: 0024-3795. DOI: https://doi.org/10.1016/0024-3795(84)90207-6. URL: http://www.sciencedirect.com/science/article/pii/0024379584902076.

[6] Grigoriy Blekherman, Rainer Sinn, and Mauricio Velasco. "Do Sums of Squares Dream of Free Resolutions?" In: *SIAM J. Appl. Algebra Geom.* 1 (2016), pp. 175–199.

[7] Monique Laurent. "The real positive semidefinite completion problem for series-parallel graphs". In: *Linear Algebra Appl.* 252 (1997), pp. 347–366. ISSN: 0024-3795. DOI: 10.1016/0024-3795(95)00741-5. URL: https://doi.org/10.1016/0024-3795(95)00741-5.