

MIT Open Access Articles

Inference on graphs via semidefinite programming

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Bandeira, Afonso S. "Inference on Graphs via Semidefinite Programming." Proceedings of the National Academy of Sciences 113, no. 16 (April 8, 2016): 4238–4239. © 2017 National Academy of Sciences

As Published: <http://dx.doi.org/10.1073/pnas.1603405113>

Publisher: National Academy of Sciences (U.S.)

Persistent URL: <http://hdl.handle.net/1721.1/109985>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Inference on graphs via semidefinite programming

Afonso S. Bandeira^{a,1}

Inference problems on graphs arise naturally when trying to make sense of network data. Oftentimes, these problems are formulated as intractable optimization programs. This renders the need for fast heuristics to find adequate solutions and for the study of their performance. For a certain class of problems, Javanmard et al. (1) successfully use tools from statistical physics to analyze the performance of semidefinite programming relaxations, an important heuristic for intractable problems.

Inverse Problems On Graphs

A particularly interesting class of inverse problems in graphs is that of synchronization problems (2, 3). These include registration of multiple pictures of the same scene, alignment of signals, community detection, and many others. After associating objects (for example, images or signals) to nodes of a graph, the goal is to estimate labels of the nodes (viewing direction of the image, shift of the signal, or community membership) from pairwise information about labels of nodes sharing edges (obtained, for example, by comparing two images). It is productive to think of the labels as being in a group (for example, of transformations) and the pairwise information being about group ratios (or relative transformations). An important example is synchronization of 3D rotations, which is a crucial step in the reconstruction problem in cryoelectron microscopy (4).

Community detection under the binary stochastic block model has recently received significant attention (5–10). In this model, a random graph is drawn on vertices belonging to two communities of the same size. Each pair of nodes share an edge, independently, with a certain probability: p if they belong to the same community and q otherwise (oftentimes $p > q$). The goal is to recover the community memberships from the observed network. This is an instance of a synchronization problem over \mathbb{Z}_2 (11), the group of two elements. Indeed, the existence or not of an edge between two nodes provides information about the relative community membership of the nodes.

A natural approach to efficiently estimate the memberships of the nodes of a graph generated from

this model is belief propagation. In a nutshell, the idea is to have each node record its confidence of belonging to one of the communities and, in each iteration, have each node observe the confidences of all its neighbors and update its own confidence accordingly (through Bayes rule). In fact, Decelle et al. (5) used tools from statistical physics to make remarkably precise predictions of the performance of belief propagation in community detection and, based on those, conjectured the exact location of a phase transition below which it is information theoretically impossible to make an estimate of the community memberships that outperforms an uninformed random guess. This conjectured phase transition was mathematically proved in a recent series of works (6–8) that also proposed ingenious estimation procedures that provably outperform a random guess above the phase transition. However, the exact performance of belief propagation has yet to be fully rigorously understood [there has, nevertheless, been remarkable progress (12–14)].

Semidefinite Programming

From the point of view of worst-case approximation in theoretical computer science, semidefinite programming [and the so-called sum-of-squares hierarchy (15–18)] seems to provide an “off the shelf” hierarchy of optimal algorithms for a large class of problems (19) [this is related to the famous unique games conjecture (20)]. In fact, semidefinite programming has been used to highlight another phase transition (10) in the binary stochastic block model, below which it is information theoretically impossible to exactly recover the community memberships of all of the nodes (although it may be possible to outperform a random guess). Using estimates in random matrix theory, it was recently shown that the semidefinite programming heuristic succeeds at exact recovery above this threshold (21, 22). It is worth noting that belief propagation can be viewed as attempting to approximate the marginal distributions of the node memberships and not necessarily directly solving an optimization problem, whereas the semidefinite program is obtained via the relaxation of an optimization problem.

^aDepartment of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02142

Author contributions: A.S.B. wrote the paper.

The author declares no conflict of interest.

See companion article on page E2218.

¹Email: bandeira@mit.edu.

This raises the natural question of understanding the performance of the semidefinite programming approach in the sparse regime, where exact recovery is impossible. A first step toward this was taken in ref. 23, and the results were further improved in ref. 24. However, a precise rigorous understanding of its performance seems outside the scope of our current techniques. Remarkably, Javanmard et al. (1) use nonrigorous tools from statistical physics to make very precise predictions, confirmed with numerical simulations, of the performance of this approach. These predictions suggest that semidefinite programming performs extremely well in this regime. In my view, this is particularly exciting for a few reasons: Firstly, it suggests that algorithms from the sum-of-squares hierarchy (the hierarchy of semidefinite programs believed to be optimal from the worst-case approximation point of view) may be nearly optimal also for a certain class of inference problems. Secondly, the results in Javanmard et al. (1) suggest that these tools may be used to produce precise predictions for the performance of algorithms in this hierarchy. Finally, this makes a fascinating connection between hierarchies of semidefinite programs (which are often believed, in theoretical computer science, to yield optimal algorithms for a certain class of problems) and belief propagation [which is tightly connected with the analysis in Javanmard et al. (1) and is often believed, in statistical physics, to be optimal for certain inference problems]. It is worth noting that the predictions suggest that, for a fixed average degree, this algorithm does not fully reach the described phase transition

(although getting very close to it). In fact, a recent paper of Moitra et al. provides another argument suggesting that this is the case (25).

Javanmard et al. (1) address not only the community detection problem but also a \mathbb{Z}_2 synchronization problem with Gaussian noise (which can be seen as a surrogate of community detection) and synchronization of rotations of the plane, $U(1)$. There is a semidefinite programming-based framework for general synchronization problems (3), and the success of the techniques used to analyze the \mathbb{Z}_2 and $U(1)$ suggest that statistical physics tools may be useful for understanding the performance of this approach to more general settings.

Although semidefinite programming is solvable in polynomial time (26), it tends to be computationally expensive. Javanmard et al. (1) used a low-rank factorization heuristic, proposed over a decade ago in ref. 27, to solve fairly large problem instances. Interestingly, they also made use of this factorization in their analysis. With the factorization, the semidefinite program can be seen as an optimization problem where each node is associated with a vector variable, which makes it more suitable for the statistical physics tools that are used. This suggests that the predictions obtained are valid for the heuristic itself, rendering it extremely effective at solving these types of semidefinite programs. In fact, this was recently backed up by rigorous guarantees for this heuristic in refs. 28 and 29.

Acknowledgments

A.S.B. was supported by National Science Foundation Grant DMS-1317308.

- 1 Javanmard A, Montanari A, Ricci-Tersenghi F (2015) Phase transitions in semidefinite relaxations. *Proc Natl Acad Sci USA* 113:E2218–E2223.
- 2 Singer A (2011) Angular synchronization by eigenvectors and semidefinite programming. *Appl Comput Harmon Anal* 30(1):20–36.
- 3 Bandeira AS, Chen Y, Singer A (2015) Non-unique games over compact groups and orientation estimation in cryo-em. arXiv:1505.03840.
- 4 Singer A, Shkolnisky Y (2011) Three-dimensional structure determination from common lines in cryo-em by eigenvectors and semidefinite programming. *SIAM J Imaging Sci* 4(2):543–572.
- 5 Decelle A, Krzakala F, Moore C, Zdeborová L (2011) Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Phys Rev E Stat Nonlin Soft Matter Phys* 84(6 Pt 2):066106.
- 6 Mossel E, Neeman J, Sly A (2012) Stochastic block models and reconstruction. arXiv:1202.1499.
- 7 Mossel E, Neeman J, Sly A (2013) A proof of the block model threshold conjecture. arXiv:1311.4115.
- 8 Massoulié L (2014) Community detection thresholds and the weak Ramanujan property. *Proceedings of the 46th Annual ACM Symposium on Theory of Computing* (Assoc Comput Machinery, New York), pp 694–703.
- 9 Mossel E, Neeman J, Sly A (2014) Consistency thresholds for the planted bisection model. arXiv:1407.1591v2.
- 10 Abbe E, Bandeira AS, Hall G (2014) Exact recovery in the stochastic block model. arXiv:1405.3267.
- 11 Abbe E, Bandeira AS, Bracher A, Singer A (2014) Decoding binary node labels from censored edge measurements: Phase transition and efficient recovery. *IEEE Trans Network Sci Eng* 1(1):10–22.
- 12 Mossel E, Neeman J, Sly A (2013) Belief propagation, robust reconstruction, and optimal recovery of block models. arXiv:1309.1380.
- 13 Bordenave C, Lelarge M, Massoulié L (2015) Non-backtracking spectrum of random graphs: Community detection and non-regular Ramanujan graphs. arXiv:1501.06087.
- 14 Abbe E, Sandon C (2015) Detection in the stochastic block model with multiple clusters: Proof of the achievability conjectures, acyclic BP, and the information-computation gap. arXiv:1512.09080.
- 15 Parrilo PA (2000) Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization. PhD thesis (Calif Inst Technol, Pasadena, CA).
- 16 Lasserre JB (2001) Global optimization with polynomials and the problem of moments. *SIAM J Optim* 11(3):796–817.
- 17 Shor N (1987) An approach to obtaining global extremums in polynomial mathematical programming problems. *Cybern Syst Anal* 23(5):695–700.
- 18 Nesterov Y (2000) Squared functional systems and optimization problems. *High Performance Optimization*, Applied Optimization Series (Springer, New York), Vol 33, pp 405–440.
- 19 Barak B, Steurer D (2014) Sum-of-squares proofs and the quest toward optimal algorithms. arXiv:1404.5236.
- 20 Khot S (2010) On the unique games conjecture (invited survey). *Proceedings of the 2010 IEEE 25th Annual Conference on Computational Complexity* (IEEE Comput Soc, Washington, DC), pp 99–121.
- 21 Bandeira AS (2015) Random Laplacian matrices and convex relaxations. arXiv:1504.03987.
- 22 Hajek B, Wu Y, Xu J (2014) Achieving exact cluster recovery threshold via semidefinite programming. arXiv:1412.6156.
- 23 Guedon O, Vershynin R (2014) Community detection in sparse networks via Grothendieck's inequality. arXiv:1411.4686.
- 24 Montanari A, Sen S (2015) Semidefinite programs on sparse random graphs. arXiv:1504.05910.
- 25 Moitra A, Perry W, Wein AS (2015) How robust are reconstruction thresholds for community detection? *Proceedings of the 48th Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York).
- 26 Vanderberghe L, Boyd S (1996) Semidefinite programming. *SIAM Rev* 38(1):49–95.
- 27 Burer S, Monteiro RDC (2003) A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Math Program* 95(2):329–357.
- 28 Bandeira AS, Boumal N, Voroninski V (2016) On the low-rank approach for semidefinite programs arising in synchronization and community detection. arXiv:1602.04426.
- 29 Montanari A (2016) A Grothendieck-type inequality for local maxima. arXiv:1603.04064.