Description of Work

Shuo Pang

ORCID
https://orcid.org/0000-0001-5652-9737

Google Scholar
https://scholar.google.com/citations?hl=en&user=kCQFJi8AAAAJ

6. *Truly Supercritical Trade-offs for Resolution, Cutting Planes, Monotone Circuits, and Weisfeiler-Leman*

   To write

5. *SoS lower bounds for Non-Gaussian Component Analysis*

   With Ilias Diakonikolas, Sushrut Karmalkar, and Aaron Potechin

   To appear at FOCS 2024

   Description. Non-Gaussian Component Analysis (NGCA) is the task of finding a non-Gaussian direction in a high-dimensional distribution—or in other words, finding a signal among noise. Specifically, given i.i.d. samples from a distribution P_v on $\mathbb{R}^n$ that behaves like a known distribution A in a hidden direction v and like a standard Gaussian in $v^\perp$, we want to find v. Of particular interest to learning and complexity theory is when the distribution A matches low-order moments with a true Gaussian, where it is known, under mild conditions, that current learning algorithms such as statistical-query and low-degree polynomial tests require a large number of samples to complete the task.

   This work studies the complexity of NGCA in the Sum-of-Squares (SoS) algorithmic framework, a state-of-the-art class of algorithms that broadly strengthens the aforementioned ones. We prove a super-constant degree SoS lower bound, which translates to super-polynomial time for the corresponding algorithms. Specifically, we show that if distribution A matches the first $(k-1)$ moments of $\mathcal{N}(0,1)$ and satisfies mild conditions, then with high probability, even given $< n^{\frac{(1-\epsilon)k}{2}}$ many Gaussian samples, degree $\widetilde{o}(\sqrt{\log n})$ SoS fails to refute the existence of direction v. This matches the known upper bound of $n^{k/2}$ [Dudeja-Hsu 22]. As a corollary, we obtain SoS lower bounds for several problems in robust statistics and learning of mixture models.

The proof introduces a new technique, along with a few improvements of the previous ones. We begin with the pseudo-calibration framework [Barak-Hopkins-Kelner-Kothari-Moitra-Potechin 2016] where, given the moment matrix M, we find an approximate factorization $M \approx LQL^T$ using minimum vertex separators, and show that with high probability Q is PSD while error terms are small. What's new is the following. First, instead of the minimum weight vertex separator, we use the minimum square separator. Second, analyzing the matrix Q requires a shift in approach, due to an intrinsic reason. In all prior works, Q, modulo a negligible error, is a "real matrix" whose entries are numbers. Here, however, it is (quite) a nontrivial linear combination of a class of equally dominating pseudo-random matrices (i.e., *graph matrices*, whose entries are polynomials of the input). To make the analysis, we introduce an algebraic method that we believe is of more general interest. Namely, we model the multiplications between the "important" pseudo-random matrices by an $\mathbb{R}$-algebra, construct all irreducible representations of this algebra, and use them to study the special element Q. Using this approach, we show that the PSDness of Q boils down to the multiplicative identities of Hermite polynomials.

4. *Graph Colouring Is Hard on Average for Polynomial Calculus and Nullstellensatz*
With Jonas Conneryd, Susanna de Rezende, Jakob Norström, Kilian Risse
FOCS 2023

Description. In this work, we prove that Polynomial Calculus, hence also Nullstellensatz, requires linear degree to refute the 3-colorability of sparse random graphs and random regular graphs. An optimal, exponential size lower bound follows via the known size-degree relation.

The proof goes by constructing an Alekhnovich-Razborov pseudo-reduction operator, where we use the technique of constructing "closure" that was introduced by [Romero-Tunçel 2021] in the study of large-girth graphs. We extend this technique to get rid of the large-girth assumption and deal with general sparse expanders.

3. *SoS Lower Bound for Exact Planted Clique*
CCC 2021

Description. This work proves Sum-of-Squares (SoS) degree lower bounds for the Exact Planted Clique problem on random graphs $G(n, 1/2)$. Here, the SoS algorithm wants to refute the existence of cliques of size ω in a sampled graph, where ω is so large that with probability $> 99.99\%$ the sample does not contain such a clique.

The word *exact* in the title means that we allow SoS algorithms to use the full set of axioms for its reasoning—including the axiom on clique size $\sum_{i=1}^{n} x_i = \omega$, the "global" axiom that the previous tight lower bound technique [Barak-Hopkins-Kelner-Kothari-Moitra-Potechin 2016] needs to weaken into a single objective function. In this work, we overcome this issue by proving a SoS degree lower bound $d = \Omega(\frac{\epsilon^2 \log n}{\log \log n})$ for $\omega = O(n^{\frac{1}{2}-\epsilon})$, almost optimal in d and ω.

Another motivation is about further developing average-case SoS lower bound techniques. To deal with the global axiom, we design the *pseudo-expectation* $\widetilde{E}(\cdot)$ differently from the pseudo-calibration framework [Barak et al. 2016]. Cost is, the moment matrix no longer has product-like entries, making it harder to analyze. To address this, we simplify the target matrix by using Hadamard product (one of the factors being a Johnson scheme, as inspired by [Feige-Krauthgamer 2003]) and a relativized matrix factorization (which conditions on subsets in the clique, whose analysis is based on binomial transform). The final positive semidefiniteness (PSD) proof relies on the analytical properties of a matrix family we call *factorial Hankel* matrices. In retrospect, the success of this approach relies on the initial design of $\widetilde{E}(\cdot)$ being "correct". Is there an a priori explanation for this design, which admittedly appears contingent?

2. *On CDCL-based Proof Systems with the Ordered Decision Strategy*
With Nathan Mull and Alexander Razborov
SAT 2020, SICOMP 2022

Description. In this work, we prove that conflict-driven clause learning (CDCL) SAT-solvers with the *ordered decision strategy* and *DECISION* *learning scheme*, are equivalent to ordered resolution. We also prove that, by replacing this learning scheme with its opposite—which stops backtracking right after the first non-conflict clause—the solvers become equivalent to general resolution. This is among the first theoretical studies of the interplay between specific decision strategies and clause learning.

For both results, we allow nondeterminism in the solver's ability to perform unit propagation, conflict analysis, and restarts, in a way similar to previous works. To aid the presentation of our results, and possibly future research, we define a model and language for discussing CDCL-based proof systems that allow for succinct and precise theorem statements.

1. *Large Clique Is Hard on Average for Resolution*
   CSR 2021

   Description. The main result of the paper is a $2^{\Omega\left(k^{1-o(1)}\right)}$ resolution size lower bound for the k-Clique problem on suitable random graph models, where $k < n^{1/3}$. This complements the result in [Beame-Impagliazzo-Subharwal 2007] which holds for $k > n^{5/6}$.

   Our proof is based on the bottleneck counting framework, where we use a variant of clause width. This width variant is defined by thresholding the 'density' of the vertex sets associated with a clause, using the same notion of density as in previous works [Beyesdorff-Galesi-Lauria 2013, Atserias-Bonacina-De Rezende-Lauria-Nordström-Razborov 2018]. We also extend the $n^{\Omega(k)}$ regular resolution lower bound [ABDLNR 2018] to a slightly stronger system that permits a certain degree of irregularity.