CONJURA · OPEN PROBLEM

Signed Low-Entropy Reweightings to Rank One on the Sphere

Entropy cost below the square-root barrier

Status. Statement: AI-written from Boaz Barak, Pravesh K. Kothari, David Steurer, *Quantum entanglement, sum of squares, and the log rank conjecture*, arXiv preprint (quant-ph), v2 submitted 9 July 2017, not yet checked by a human. Settled for $\delta \geq 1/2$ by the paper's own Theorem 2.3, even with the stronger requirement that r be nonnegative, and open for every $\delta < 1/2$; the paper additionally states that the nonnegative-r version of the question is false, so a positive answer must use cancellation.

Category. *Quantum.*

Conjecture (informal). *Take any probability distribution over unit vectors in n-dimensional space. Its second moment is the average of the outer products of those vectors, an n-by-n matrix, and in general it looks nothing like a rank one matrix. But you are allowed to reweight the distribution before averaging: multiply by a function that has average absolute value one, and pay a cost equal to the average of the absolute weight times its logarithm, which is the usual entropy price of tilting a distribution. The known result is that you can always buy an approximately rank one second moment for a price of about the square root of n, and that price is exactly right if the weights are required to be nonnegative, which is the case where reweighting means conditioning on an event. The question asked here is whether allowing the weights to be negative — so that cancellation, not just conditioning, is available — lets you buy the same thing for a price of n raised to an arbitrarily small positive power. The authors note that a positive answer below the square-root threshold may improve the best known bound for the log rank conjecture to* $\tilde{O}(n^{\delta})$ *— subject to the caveats of their footnotes 7 and 8 about which notion of "approximate" is needed and how the bound depends on* ε *— and that, if appropriately extended to pseudo-distributions, it would improve their algorithm's running time to* $\exp(\tilde{O}(n^{\delta}))$.

1 The setting

The Best Separable State problem $\mathrm{BSS}_{c,s}$ asks, given an $n^2 \times n^2$ quantum measurement operator, to distinguish the case that some separable (non-entangled) state is accepted with probability at least c from the case that every separable state is accepted with probability at most s. Equivalently, in its non-quantum form, it asks whether a given linear subspace of $n \times n$ matrices contains a rank one matrix or is ε-far from every rank one matrix. The paper's algorithm is the first for the problem that beats the brute-force $2^{O(n)}$ bound for general measurements; the earlier sos-based algorithm of Doherty, Parrilo and Spedalieri applies to general measurements but is not known to beat brute force, and Brandão–Christandl–Yard's quasi-polynomial guarantee is restricted to 1-LOCC measurements.

That algorithm runs the sum-of-squares hierarchy for $\tilde{O}(\sqrt{n})$ rounds and rounds the resulting pseudo-distribution using a *structure theorem*: every pseudo-distribution over the unit sphere admits a degree-$\tilde{O}(\sqrt{n})$ reweighting whose second moment is ε-close to rank one in Frobenius norm. The paper's Theorem 2.3 ("rank one reweighting") is stated for an arbitrary distribution over rank one $n \times n$ matrices X; in the setting of honest probability distributions, Theorem 2.3 specialises, on taking $X = vv^\top$, to the $\delta = 1/2$, nonnegative-r instance of the statement below. The paper also gives an equivalent combinatorial dual form (Theorem 2.4): every $N \times N$ real matrix of rank at most n has a principal submatrix on an index set of measure $\exp(-\sqrt{n}\,\mathrm{poly}(1/\varepsilon))$ that is ε-close to rank one.

The technique is a lift of Lovett's $\tilde{O}(\sqrt{n})$ bound on the deterministic communication complexity of rank-n Boolean matrices [Lov14] (following Rothvoß's proof [Rot14]), which is in turn the best known progress on the log rank conjecture of Lovász and Saks [LS88]; the equivalence between small monochromatic rectangles and low communication complexity is due to Nisan and Wigderson [NW94]. The paper's contribution is that Booleanity can be dropped if "exactly rank one" is relaxed to "approximately rank one", and that the whole argument can be embedded in the sum-of-squares proof system [BKS16].

The $\sqrt{n}$ in the exponent is therefore the same $\sqrt{n}$ that blocks progress on the log rank conjecture. The paper observes that its own Theorem 2.3 is tight as stated, so the exponent cannot be improved within the nonnegative-reweighting formulation, and that a black-box reduction from the log rank conjecture to a faster BSS algorithm looks unlikely because one needs statements that are not restricted to Boolean matrices. Question 8.1 is the authors' proposal for such a statement: relax nonnegativity of the reweighting and ask for entropy cost n^δ instead. Footnote 7 of the paper records a caveat, due to Gavinsky and Lovett [GL14], that the exact notion of "approximate" matters for the log rank application, and footnote 8 warns that the dependence of the bound on ε would need better control than the paper's own setting requires.

Progress to date. The paper proves the $\delta = 1/2$ case, in fact with a nonnegative r: its Theorem 2.3 gives, for every distribution μ over rank one n-by-n matrices and every $\varepsilon > 0$, a nonnegative reweighting of Kullback–Leibler cost at most $\sqrt{n}\,\mathrm{poly}(1/\varepsilon)$ whose mean is within ε (relative Frobenius) of a rank one matrix, and Theorem 5.1 lifts this to degree-$\tilde{O}(\sqrt{n})$ sum-of-squares reweightings of pseudo-distributions. The paper further asserts that as stated Theorem 2.3 is tight, and states outright that the answer to Question 8.1 is No if negative reweighting functions are disallowed. So the nonnegative branch is closed and every $\delta < 1/2$ requires cancellation. The paper reports no partial result for signed r at any $\delta < 1/2$.

Why it matters. A positive answer for any $\delta < 1/2$ would be the first improvement in a decade on the best known upper bound for the log rank conjecture, taking it from $\tilde{O}(\sqrt{n})$ to $\tilde{O}(n^\delta)$, and — if the argument can be run inside the sum-of-squares proof system, which is what the paper's whole machinery is designed to make possible — would simultaneously improve the running time of the only known sub-exponential algorithm for Best Separable State from $\exp(\tilde{O}(\sqrt{n}))$ to $\exp(\tilde{O}(n^\delta))$. Two problems that have resisted independently would move together. A negative answer would be almost as informative: it would say that the cancellation available to signed reweightings buys nothing over conditioning, closing off the route the authors identify as the most promising one and forcing the search for a genuinely different notion of closeness to rank one.

Why it is clean. The statement is three lines of elementary real analysis and linear algebra: a probability measure on a sphere, a signed function with a normalisation and an entropy bound, and a Frobenius-norm distance to a rank one matrix. It mentions no quantum mechanics, no communication complexity, no semidefinite programming and no pseudo-distributions, even though it was extracted from a paper about all four. There is nothing to install and nothing unpublished to rely on; a reader can start attacking it from the statement alone. The known boundary case is equally clean, which makes it easy to check whether a proposed construction is actually doing anything new.

2 Notation and parameters

- $n \in \mathbb{N}$ is the ambient dimension, and $\mathbb{S}^{n-1} = \{v \in \mathbb{R}^n : \|v\|_2 = 1\}$ is the real unit sphere.
- For $A \in \mathbb{R}^{n\times n}$, $\|A\|_F = \left(\sum_{i,j} A_{i,j}^2\right)^{1/2}$ denotes the Frobenius norm.
- A matrix $L \in \mathbb{R}^{n\times n}$ is *rank one* if $L = ab^\top$ for some $a, b \in \mathbb{R}^n$; it is *nonzero* if $L \neq 0$.
- μ denotes a Borel probability measure on $\mathbb{S}^{n-1}$, and $r : \mathbb{S}^{n-1} \to \mathbb{R}$ a Borel measurable function that is *not* required to be nonnegative.
- $\mathbb{E}_{v\sim\mu}[\, r(v)vv^\top \,]$ denotes the $n \times n$ matrix whose (i,j) entry is $\mathbb{E}_{v\sim\mu}[\, r(v)v_iv_j \,]$.
- $\varepsilon > 0$ is the accuracy parameter (relative Frobenius distance to rank one) and $\delta > 0$ is the entropy exponent.
- Throughout, $0 \log 0 := 0$.

Parameters.

- n — ambient dimension; the sphere is the unit sphere of $\mathbb{R}^n$ and the second-moment matrix is n by n.
- ε — relative Frobenius accuracy: how close the reweighted second moment must be to a rank one matrix.
- δ — entropy exponent: the reweighting is allowed entropy cost $O(n^\delta)$. The content of the question is the regime $\delta < 1/2$.
- $C(\varepsilon, \delta)$ — the constant hidden by the paper's $O(n^\delta)$; it may depend on ε and δ but not on n or μ.

- r — the signed reweighting function; the point of the question is that r is not required to be nonnegative.
- L — the nonzero rank one n-by-n matrix that the reweighted second moment must approximate.

3 The conjecture

Definition 1 (Signed reweighting and its entropy cost). Let μ be a Borel probability measure on $\mathbb{S}^{n-1}$ and let $K \geq 0$. A Borel measurable function $r : \mathbb{S}^{n-1} \to \mathbb{R}$ is called a *signed reweighting of μ with entropy cost at most K* if r is μ-integrable and

$$\mathop{\mathbb{E}}_{v \sim \mu} \big|r(v)\big| = 1,$$

$$\mathop{\mathbb{E}}_{v \sim \mu} \Big[\big|r(v)\big| \log \big|r(v)\big| \Big] \leq K.$$

If in addition $r \geq 0$ holds μ-almost everywhere, then r is a *nonnegative* reweighting; in that case $d\mu' := r\, d\mu$ is a probability measure and the entropy cost is exactly the Kullback–Leibler divergence $\Delta_{\mathrm{KL}}(\mu' \,\|\, \mu) = \mathbb{E}_{v \sim \mu'} \log \big(d\mu' / d\mu\big)$.

Conjecture 1 (signed low-entropy reweighting to rank one). *For every $\varepsilon > 0$ and every $\delta > 0$ there exists a finite constant $C = C(\varepsilon, \delta)$ such that the following holds for every $n \in \mathbb{N}$ and every Borel probability measure μ on $\mathbb{S}^{n-1}$: there exist a signed reweighting r of μ with entropy cost at most $C \cdot n^{\delta}$ (in the sense of Definition 1) and a nonzero rank one matrix $L \in \mathbb{R}^{n \times n}$ such that*

$$\Big\| \mathop{\mathbb{E}}_{v \sim \mu} \big[r(v)\, v v^{\top} \big] - L \Big\|_F \leq \varepsilon \, \|L\|_F .$$

Remark 1 (three harvester-side readings a reviewer should check). First, quantifier placement on the constant: the paper writes the entropy bound as $O(n^{\delta})$, which by convention hides a constant that may depend on ε and δ but not on n or μ; that constant is made explicit above as $C(\varepsilon, \delta)$. Footnote 8 of the paper explicitly warns that for the log rank application one needs better control of the ε-dependence than the paper's own setting requires, so a solver should track C's dependence on ε rather than treat it as free. Second, regularity: the paper says only "distribution over $\mathbb{S}^{n-1}$" and does not specify measurability or integrability conditions on r; Definition 1 writes Borel measurable and μ-integrable. The essential case is finitely supported μ (this is what the sum-of-squares application produces), and the general case should follow by approximation, but this is the harvester's reading and not the paper's words. Third, currency: the paper's assertion that the nonnegative-r answer is No is stated without proof, resting on the remark that Theorem 2.3 is tight as stated; and the claims that Lovett's $\tilde{O}(\sqrt{n})$ remains the best known bound for the log rank conjecture [Lov14] and that Question 8.1 has not been answered for $\delta < 1/2$ are recalled from memory rather than from the paper. The log rank literature has moved since 2017 and a reviewer should verify both before publishing; no resolution is known to the harvester, but this is not certain.

4 Bibliography

[BKS16] B. Barak, P. Kothari, D. Steurer. *Proofs, beliefs, and algorithms through the lens of sum-of-squares.* Lecture notes in preparation, available on http://sumofsquares.org, 2016.

[GL14] D. Gavinsky, S. Lovett. *En route to the log-rank conjecture: New reductions and equivalent formulations.* ICALP (1), Lecture Notes in Computer Science, vol. 8572, Springer, pages 514–524, 2014.

[Lov14] S. Lovett. *Communication is bounded by root of rank.* STOC, ACM, pages 842–846, 2014.

[LS88] L. Lovász, M. E. Saks. *Lattices, möbius functions and communication complexity.* FOCS, IEEE Computer Society, pages 81–90, 1988.

[NW94] N. Nisan, A. Wigderson. *On rank vs. communication complexity.* FOCS, IEEE Computer Society, pages 831–836, 1994.

[Rot14] T. Rothvoß. *A direct proof for Lovett's bound on the communication complexity of low rank matrices.* CoRR abs/1409.6366, 2014.