Repairing the linearization radius in high-precision shadow estimation

Yuxuan Zhang
École Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland
Department of Physics, Princeton University, Princeton, NJ 08544, USA
zhangyx@comp.nus.edu.sg

17 June 2026*

Abstract

The linearization argument behind the high-precision shadow-estimation bounds of Chen, Li and Liu is valid only inside a radius $\|E\|_F \le (0.01/t)^4$. We record two things. First, the underlying tail lemma does not close as printed: the $j = 2$ term exceeds the stated bound by a factor of order 10^{11}, independently of the radius. Second, the constant in the Haar-moment estimate is slack enough that the radius can be replaced by $\|E\|_F \le 1/(4et)$ — fourth power to first — which moves the precision exponent from 12 to 5/2 for balanced and structured observables and to 11/2 in general, at unchanged copy complexity. We then place this honestly: the framework itself was subsequently bypassed.

1. Background

Shadow tomography asks for the expectation values of m observables on an unknown d-dimensional state to accuracy ε, using as few copies as possible [1]. The classical shadows protocol [2] made the problem practical, and the conjectured optimum is $\Theta(\log m/\varepsilon^2)$ copies with no poly(d) overhead.

Chen, Li and Liu [3] prove this in the high-precision regime. Their argument linearizes the relevant estimator around the maximally mixed state, and the linearization is controlled only for $\|E\|_F \le (0.01/t)^4$, with the supporting moment lemma proved in [4]. That fourth power is what confines the result to $\varepsilon \le d^{-12}$, and closing the gap down to $\varepsilon \sim d^{-1}$ was the stated open problem.

2. The problem

The step at issue. *The tail of the linearization error is bounded in [3] by*

$$\sum_{j\ge 2} 2^j \binom{t}{j}^2 (8j)^{8j} \|E\|^{2j} \;\le\; (100t)^4 \|E\|^4,$$

valid for $\|E\|_F \le (0.01/t)^4$. *How large can the radius be made, and is the stated inequality correct as printed?*

*Incoming: Department of Computer Science and Department of Physics, National University of Singapore, December 2026.

3. The lemma does not close as printed

The $j = 2$ term alone exceeds the claimed bound by a factor of order 1.8×10^{11}, and the excess is independent of $\|E\|$ — shrinking the radius does not rescue it. The replacement below therefore *repairs* the lemma rather than sharpening it, which matters for anyone citing the step.

4. The repaired radius

Using the exact identity

$$\big\|\mathbb{E}_U\big[(UEU^\dagger)^{\otimes t'}\big]\big\|_{\mathrm{op}} = \max_\lambda \frac{|s_\lambda(\alpha)|}{\dim V_\lambda^{t'}},$$

with $|s_\lambda| \le N_1(t')f^\lambda$ and $\max R = t'^{t'}/t'! \le e^{t'}$, the per-term constant is bounded by $t'^{t'}$. Resumming gives

$$\|E\|_F \;\le\; \frac{1}{4et}$$

in place of $(0.01/t)^4$, with a better constant besides. Downstream the precision exponent falls from 12 to 5/2 for balanced and structured observables and to 11/2 in general, with copy complexity unchanged at $O(\log(1/\delta)/\varepsilon^2)$.

5. Verification, and an exposition slip

The Murnaghan–Nakayama table was rebuilt from scratch and cross-checked against the power-sum and bialternant formulas; the logical link was confirmed on 40k signed traceless spectra with no violation; the tail ratio stays below 0.78 uniformly out to $t = 5000$. Two accounting gaps found adversarially were closed, both exponent-neutral. One is an exposition slip in [3], where a median bound is used as though it were a mean; the quantity is recoverable from an exact Schur–Weyl second moment, so nothing downstream breaks.

6. What remains open — and where this now sits

The honest placement matters more than the result. All of the above lives *inside* the linearization framework, and that framework has since been bypassed: Pelecanos, Spilecki and Wright [5] give an exactly unbiased estimator, which has no bias radius to respect and reaches $\varepsilon \lesssim d^{-1}$ directly, closing the moderate-precision gap and disproving the conjectured optimal scaling of [3]. The improvement recorded here is therefore of interest for the technique and for the record on the lemma, not for the state of the art.

What is genuinely open: matching *lower* bounds showing [5] optimal across the range, which those authors conjecture but do not prove; and the low-accuracy regime $\varepsilon \gtrsim d^{-1}$, where the $d^{2/3}/\varepsilon^{4/3}$ term dominates and no lower bound is known.

Disclosure

The mathematical content of this note — the construction, the case analysis, and the verification scripts — was produced by AI agents. My contribution was the choice of problem, the intuition for where a proof might come from, and the design of the adversarial procedure under which the claim was attacked and checked; I did not intervene in the derivation itself. Every load-bearing step was re-derived independently, in exact arithmetic where the problem allows it. This note is a record of a computer-assisted result. It has not been refereed.

References

[1] S. Aaronson, *Shadow tomography of quantum states*, arXiv:1711.01053.

[2] H.-Y. Huang, R. Kueng and J. Preskill, *Predicting many properties of a quantum system from very few measurements*, arXiv:2002.08953.

[3] S. Chen, J. Li and A. Liu, *Optimal high-precision shadow estimation*, arXiv:2407.13874.

[4] S. Chen, J. Li and A. Liu, *An optimal tradeoff between entanglement and copy complexity for state tomography*, arXiv:2402.16353.

[5] A. Pelecanos, J. Spilecki and J. Wright, *The debiased Keyl's algorithm: a new unbiased estimator for full state tomography*, arXiv:2510.07788.