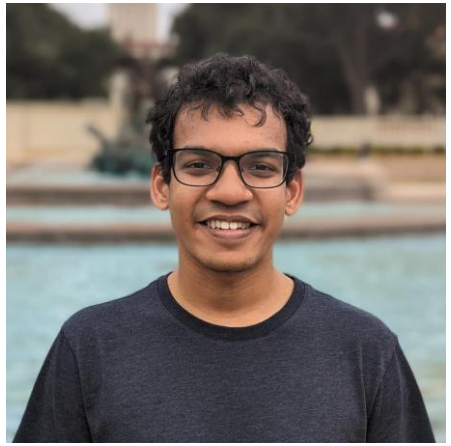
Computational Complexity of Testing in the Spiked Wigner Model

Alex Wein

UC Davis

Based on joint works with Ankur Moitra (MIT) and Ansh Nagda





High-Dimensional Testing: Spiked Wigner

[Johnstone '01; Peche '06 …]

- Given an $n \times n$ symmetric matrix Y drawn from one of two distributions…
  - "Null" $\mathbb{Q}$: $Y = W$ $\quad W_{ij} = W_{ji} \sim N(0,1/n),\ W_{ii} \sim N(0,2/n)$
  - "Planted" $\mathbb{P}$: $Y = \lambda \frac{xx^\intercal}{\|x\|^2} + W$ $\quad x \in \mathbb{R}^n$ i.i.d. from π, mean 0

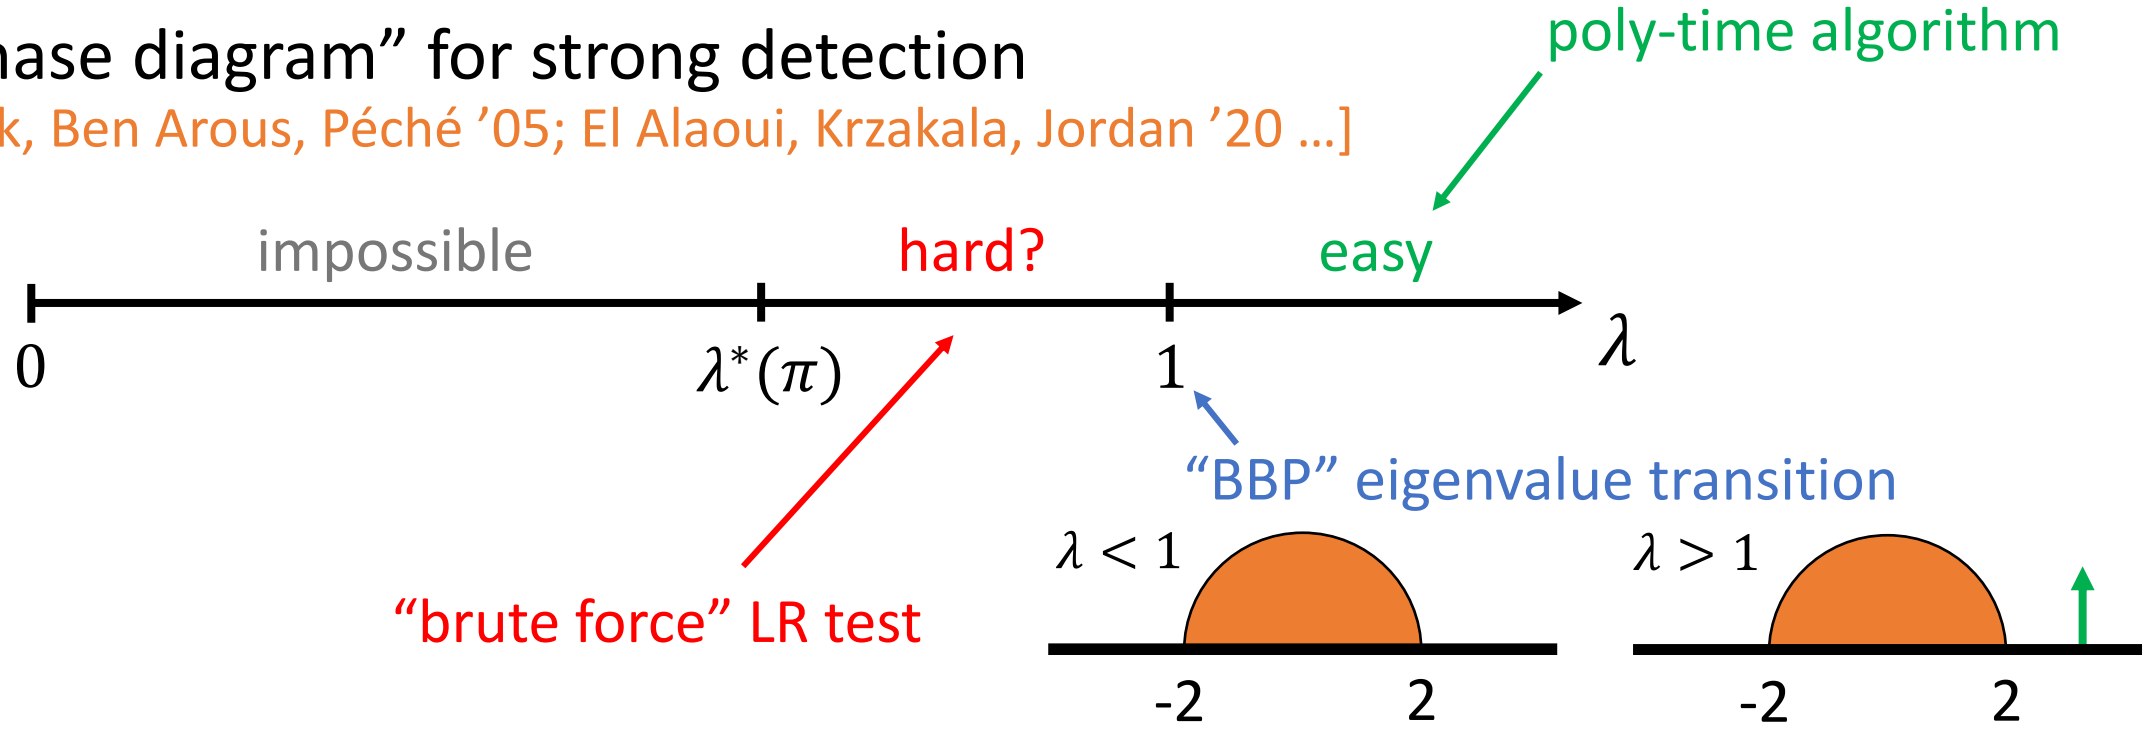
  … can you guess which distribution generated Y ($\mathbb{P}$ or $\mathbb{Q}$)?

- $n \to \infty$ with π and $\lambda > 0$ fixed, known
- Prior π: e.g., sparse Rademacher $\quad x_i = \begin{cases} +1 & \text{with probability } \rho/2 \\ -1 & \text{with probability } \rho/2 \\ 0 & \text{with probability } 1-\rho \end{cases}$

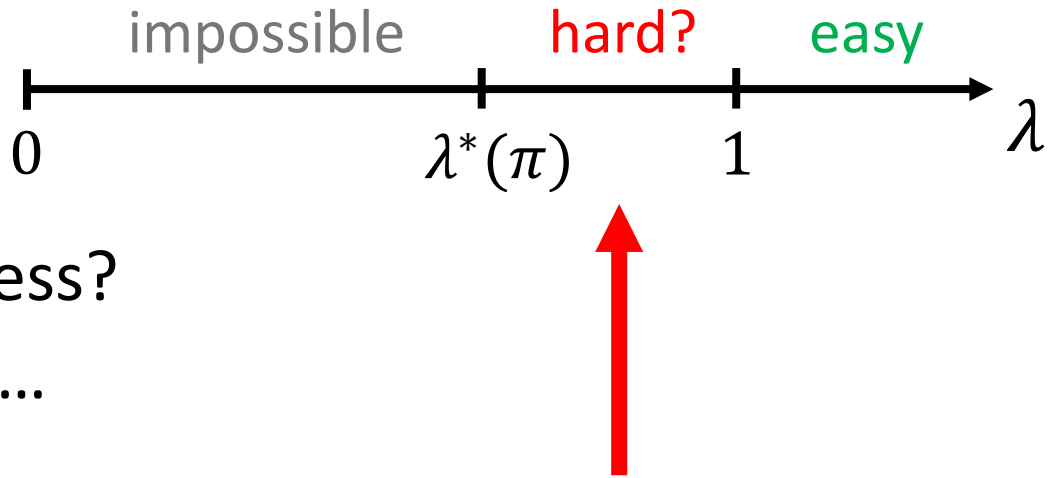
- Basic model for detecting planted "signal" buried in "noise"
- Testbed for developing mathematical tools…

Strong Detection

- Testing $\mathbb{P} = \mathbb{P}_n$ (planted) vs $\mathbb{Q} = \mathbb{Q}_n$ (null)
- Strong detection: type I + type II $\to 0$ as $n \to \infty$
  - Test t: $\mathbb{R}^{n \times n} \to \{0,1\}$
  - $\mathbb{Q}(t(Y) = 1) \to 0$ and $\mathbb{P}(t(Y) = 0) \to 0$
- "Phase diagram" for strong detection
  [Baik, Ben Arous, Péché '05; El Alaoui, Krzakala, Jordan '20 …]



“Hard” Regime



- How to show computational hardness?
- Average-case complexity is difficult…
- Instead:
  - Average-case reductions
    - E.g. from “planted clique” assumption
    - Current results don’t pinpoint “sharp” thresholds
  - Failure of restricted classes of algorithms
    - Statistical query (SQ) algorithms
    - Sum-of-squares (SoS) hierarchy
    - “Local search” algorithms
    - Approximate message passing (AMP)
    - …
    - Low-degree polynomials

Low-Degree Testing

[Hopkins, Steurer '17; Hopkins, Kothari, Potechin, Raghavendra, Schramm, Steurer '17; Hopkins '18 …]
Surveys: [Kunisky, W, Bandeira '19; W '25]

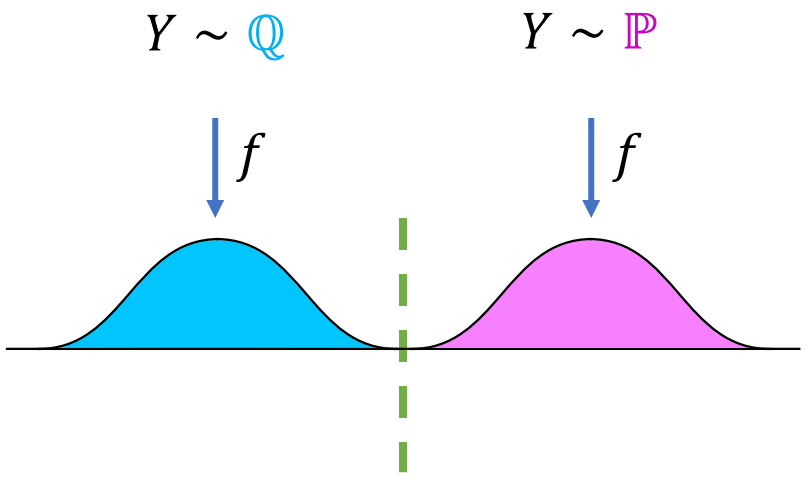
- **Degree-D test**: multivariate polynomial $f\colon \mathbb{R}^{n\times n} \to \mathbb{R}$ of degree $D = D_n$
- Examples:
  - $f(Y) = \mathrm{Tr}(Y) = \sum_i \mu_i$ (degree 1)
  - $f(Y) = \mathrm{Tr}(Y^p) = \sum_i \mu_i^p$ (degree p)
    - Approximates $\mu_{\max}$ when $p \approx \log n$
  - $f(Y) = \sum_i Y_{ii}^p$
    - Approximates largest diagonal entry

$\{\mu_i\}$ = eigenvalues of Y

- "Success": $f = f_n$ *strongly separates* $\mathbb{P}$ and $\mathbb{Q}$ if

$$\sqrt{\max\{\mathrm{Var}_{\mathbb{P}}(f), \mathrm{Var}_{\mathbb{Q}}(f)\}} = o(|\mathrm{E}_{\mathbb{P}}[f] - \mathrm{E}_{\mathbb{Q}}[f]|)$$

as $n \to \infty$

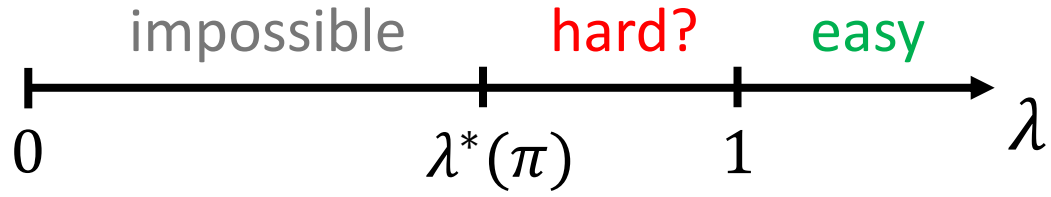


Degree vs Runtime

- Heuristic:
  - $\deg D \approx \text{time } n^{\widetilde{\Theta}(D)}$ (term-by-term evaluation)
  - $\deg O(1) < \text{poly time} < \deg O(\log n)$
- "Low-degree conjecture" [Hopkins '18 …]

  If low-degree polynomials fail, so do all algorithms of the corresponding runtime
  - Don't hope to prove it
  - Many success stories
  - But not true for all $\mathbb{P}$, $\mathbb{Q}$… [Buhai, Hsieh, Jain, Kothari '25; Mao '26 …]
  - Spiked Wigner is a "core" model in this area
- Or think of low-degree lower bounds as ruling out restricted algorithms

Back to Spiked Wigner



- Recall: poly-time algorithm for strong detection when $\lambda > 1$ (max eig)
- If $\lambda < 1$ and $D \ll n$, no degree-D polynomial strongly separates [Kunisky, W, Bandeira '19]

  $\Rightarrow$ Conjecture: strong detection requires time $\exp(n^{1-o(1)})$ when $\lambda < 1$

  Conjectured previously for *recovery* [Lesieur, Krzakala, Zdeborová '15]

- What about weak detection? (beat random guess)
  - $\mathbb{Q}(t(Y) = 1) + \mathbb{P}(t(Y) = 0) = 1 - \Omega(1)$
- Easy for any (constant) $\lambda > 0$ using $\mathrm{Tr}(Y)$
- This talk: what is the best possible error rate for weak detection?

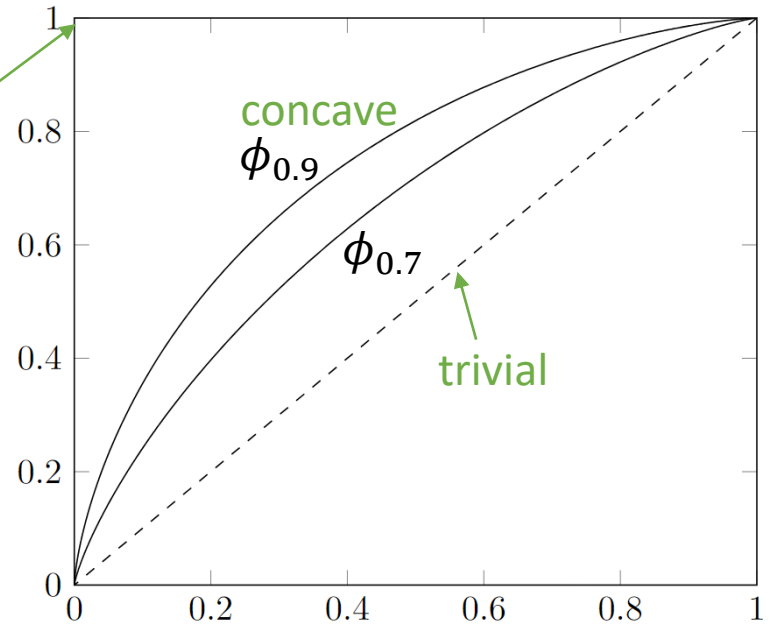
Linear Spectral Statistics (LSS)

[Bai, Silverstein '04; Bai, Yao '05; Baik, Lee '17; Chung, Lee '22 …]

- Best known poly-time algorithm for weak detection when $\lambda < 1$
- Threshold $\sum_i f_\lambda(\mu_i)$ where $f_\lambda(\mu) = -\log(1 - \sqrt{\lambda}\mu + \lambda)$
- Achieves a particular ROC (receiver operating characteristic) curve ϕ_λ

$$\phi_\lambda(\alpha) = 1 - \Phi[\Phi^{-1}(1-\alpha) - \sqrt{-\log(1-\lambda^2)/2}]$$

Φ = standard normal CDF



β = power
= true positive rate

α = type I error
= false positive rate

- IT optimal when $\lambda < \lambda^*(\pi)$
  [El Alaoui, Krzakala, Jordan '20]
- Poly-time optimal when $\lambda^*(\pi) < \lambda < 1$??
  - Strong detection is possible but hard

This Talk: LSS is Optimal

- Spiked Wigner model with $\lambda^*(\pi) < \lambda < 1$: strong detection possible-but-hard, but what is the best possible weak detection?
  - Previously no tools to address this type of question

- **Main result:** ROC curve of LSS cannot be improved by any poly-time test
  - Conditional on an appropriate complexity-theoretic conjecture…

- Consequence ("computational universality"): while IT threshold $\lambda^*(\pi)$ depends on prior π, the best computationally-efficient test only uses the spectrum and is thus oblivious to the prior

Part 1: LSS is Optimal, Conditional on the "Strong" Low-Degree Conjecture

"Precise Error Rates for Computationally Efficient Testing"
joint with Ankur Moitra (MIT)
arXiv:2311.00289 (2023)
Annals of Statistics, 2025



Low-Degree "Advantage" / Low-Degree Likelihood Ratio

[Hopkins, Steurer '17; Hopkins, Kothari, Potechin, Raghavendra, Schramm, Steurer '17; Hopkins '18 …]

- Standard tool for low-degree lower bounds: low-degree advantage

$$\mathrm{Adv}_{\leq D} := \sup_{f \text{ deg } D} \frac{\mathrm{E}_{Y\sim\mathbb{P}}[f(Y)]}{\sqrt{\mathrm{E}_{Y\sim\mathbb{Q}}[f(Y)^2]}}$$

- Goal: show $\mathrm{Adv}_{\leq D} = O(1)$, implying that no deg-D polynomial strongly separates $\mathbb{P}$ and $\mathbb{Q}$
  - Or $1 + o(1)$ to rule out "weak separation"
- Tractable to compute:

$$(\mathrm{Adv}_{\leq D})^2 = \sum_{\alpha} \mathrm{E}_{Y\sim\mathbb{P}}[h_\alpha(Y)]^2$$

  where $\{h_\alpha\}$ is a basis of orthogonal (Hermite) polynomials of deg $\leq D$ w.r.t. $\mathbb{Q}$
- Relation to χ^2-divergence: $(\mathrm{Adv}_{\leq\infty})^2 = \chi^2(\mathbb{P}\|\mathbb{Q}) + 1$

Strong Low-Degree Conjecture

- Wigner: for any $\lambda < 1$, any* π, and any $D = D_n = o(n/\log n)$,

$$\mathrm{Adv}_{\leq D} := \sup_{f \ \deg D} \frac{\mathrm{E}_{Y\sim\mathbb{P}}[f(Y)]}{\sqrt{\mathrm{E}_{Y\sim\mathbb{Q}}[f(Y)^2]}} \to (1-\lambda^2)^{-1/4} \text{ as } n \to \infty$$

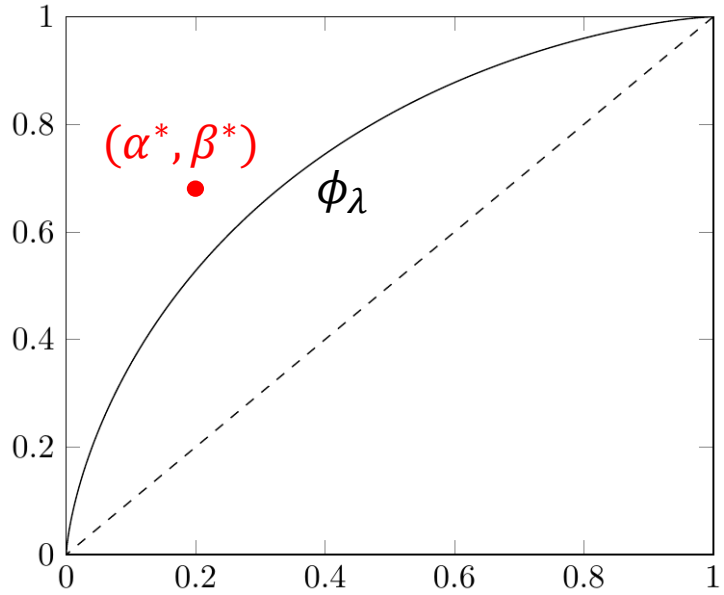
- This is $O(1)$, implying no strong separation by degree-D polynomials
- "Standard" LD conjecture: strong detection requires exponential runtime $\exp(n^{1-o(1)})$
- **Strong LD Conjecture (new)**: for spiked Wigner, any $f = f_n$ with

$$\limsup_{n\to\infty} \frac{\mathrm{E}_{Y\sim\mathbb{P}}[f(Y)]}{\sqrt{\mathrm{E}_{Y\sim\mathbb{Q}}[f(Y)^2]}} > (1-\lambda^2)^{-1/4}$$

requires runtime $\exp(n^{1-o(1)})$

Main Result

- Assuming strong LD conjecture, LSS has optimal ROC curve among efficient algorithms



- **Theorem**:
  - Fix $\lambda \in (0,1)$
  - Fix any* spike prior π
  - Assume the strong LD conjecture
  - Suppose (α^*, β^*) lies above the ROC curve of LSS
    - i.e., $\beta^* > \phi_\lambda(\alpha^*)$
  - Then any test with $(\text{type I}, \text{power}) \to (\alpha^*, \beta^*)$ requires runtime $\exp(n^{1-o(1)})$

Proof Idea

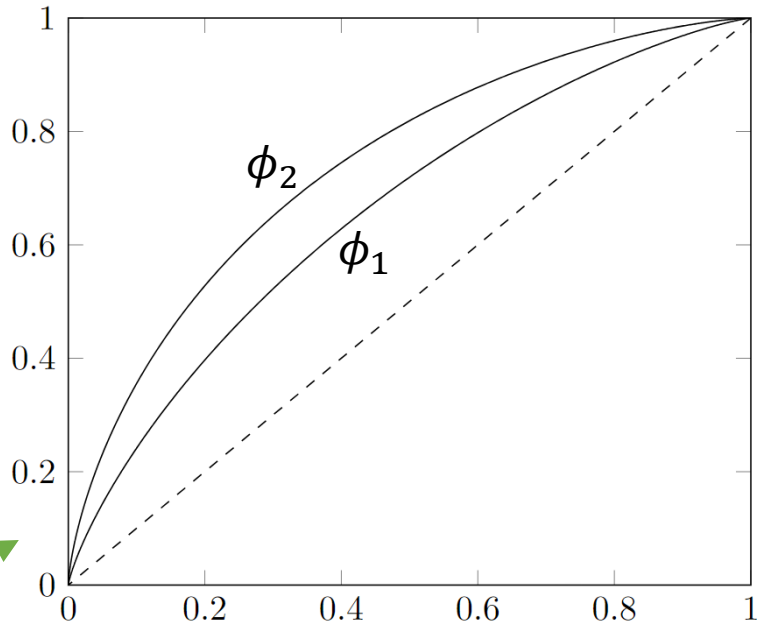
- Given achievable (concave) ROC curve ϕ, can construct f with

$$\mathrm{ratio}(f) = \mathrm{val}(\phi)$$

- $\mathrm{ratio}(f) := \dfrac{\mathrm{E}_{Y\sim\mathbb{P}}[f(Y)]}{\sqrt{\mathrm{E}_{Y\sim\mathbb{Q}}[f(Y)^2]}}$
- $\mathrm{val}(\phi) := \sqrt{\int_0^1 \left(\phi'(\alpha)\right)^2 d\alpha}$

- Better curve has better $\mathrm{val}(\phi)$

β = power
= true positive rate



α = type I error
= false positive rate

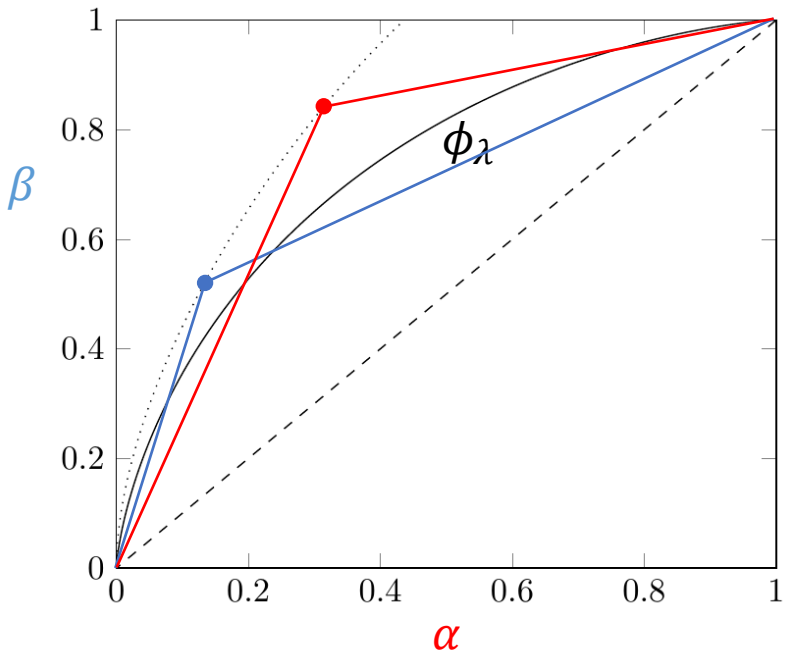
$\mathrm{val}(\phi_2) > \mathrm{val}(\phi_1)$

Proof Idea

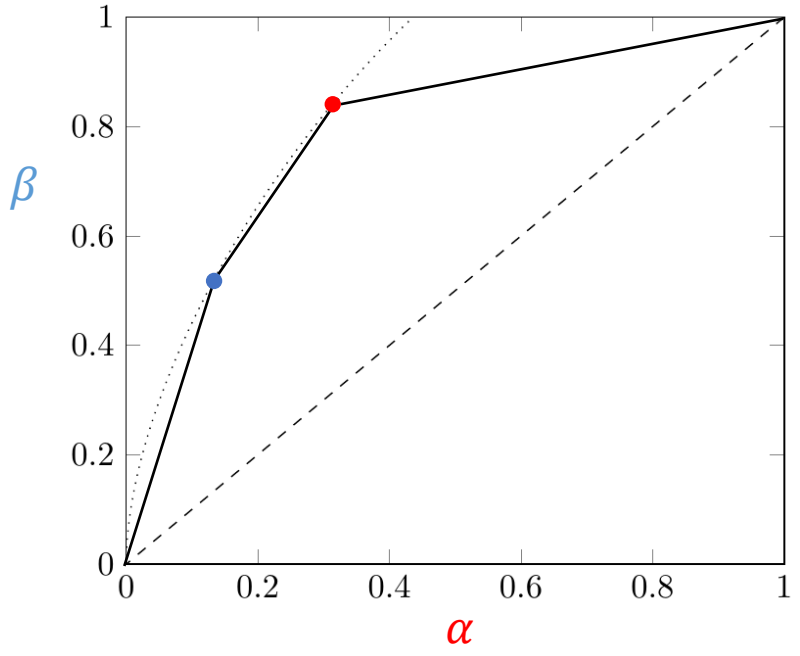
$$\mathrm{ratio}(f) \coloneqq \frac{\mathrm{E}_{Y\sim\mathbb{P}}[f(Y)]}{\sqrt{\mathrm{E}_{Y\sim\mathbb{Q}}[f(Y)^2]}}$$

$$\mathrm{val}(\phi) \coloneqq \sqrt{\int_0^1 \left(\phi'(\alpha)\right)^2 d\alpha}$$

- Recall: $\mathrm{ratio}(f) \leq (1-\lambda^2)^{-1/4}$
  - For low-degree f, and conjecturally for all efficiently-computable f
- Given this, what ROC curves are possible?
  - Must have $\mathrm{val}(\phi) \leq (1-\lambda^2)^{-1/4}$
  - Many possibilities…



but not

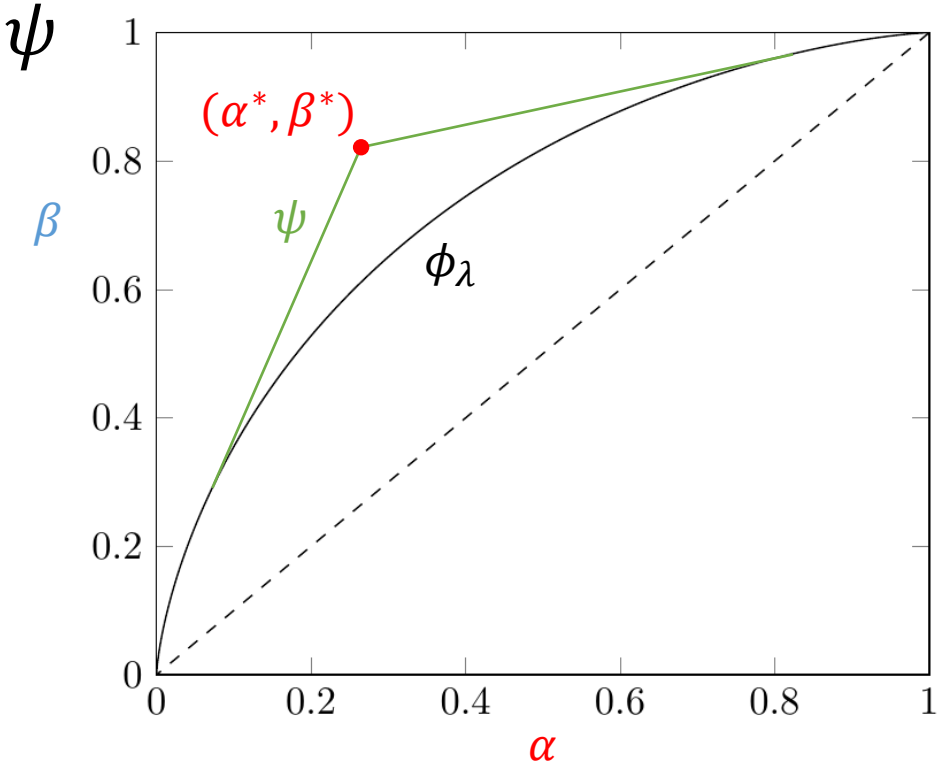


Proof

$$\text{ratio}(f) \coloneqq \frac{\mathrm{E}_{Y\sim\mathbb{P}}[f(Y)]}{\sqrt{\mathrm{E}_{Y\sim\mathbb{Q}}[f(Y)^2]}}$$

$$\text{val}(\phi) \coloneqq \sqrt{\int_0^1 \big(\phi'(\alpha)\big)^2 \, d\alpha}$$

- We know ϕ_λ is achievable in poly time, yielding ratio
$$\text{val}(\phi_\lambda) = (1-\lambda^2)^{-1/4}$$
- Assume for contradiction: some (α^*, β^*) above ϕ_λ is achievable
- Can then achieve an even better ROC curve ψ
- Thus achieving ratio
$\text{val}(\psi) > \text{val}(\phi_\lambda) = (1-\lambda^2)^{-1/4}$
- Contradicts strong LD conjecture
- Conclude: (α^*, β^*) not achievable (in sub-exponential time)



Recap

- To show optimality of an ROC curve ϕ, need two ingredients…
  1. A *positive result* achieving ϕ
  2. A precise bound on $\mathrm{Adv}_{\leq D}$, matching $\mathrm{val}(\phi)$
- Conclude: cannot beat ϕ in time $\approx n^D$, assuming strong LD conjecture

- Can prove (unconditional) *information-theoretic* results ($D = \infty$)

- But should we believe the strong LD conjecture…?

(Refined) Strong Low-Degree Conjecture

- **Strong LD Conjecture (v1)**: for spiked Wigner, any $f = f_n$ with

$$\limsup_{n\to\infty} \frac{\mathrm{E}_{Y\sim\mathbb{P}}[f(Y)]}{\sqrt{\mathrm{E}_{Y\sim\mathbb{Q}}[f(Y)^2]}} > (1-\lambda^2)^{-1/4}$$ requires runtime $\exp(n^{1-o(1)})$

- Refuted by Ansh Nagda (2025)
  - Issue: f can "cheat" by exploiting rare events (bet on one signal)
- **Strong LD Conjecture (v2)**: for spiked Wigner, any nice* $f = f_n$ with

$$\limsup_{n\to\infty} \frac{\mathrm{E}_{Y\sim\mathbb{P}}[f(Y)]}{\sqrt{\mathrm{E}_{Y\sim\mathbb{Q}}[f(Y)^2]}} > (1-\lambda^2)^{-1/4}$$ requires runtime $\exp(n^{1-o(1)})$

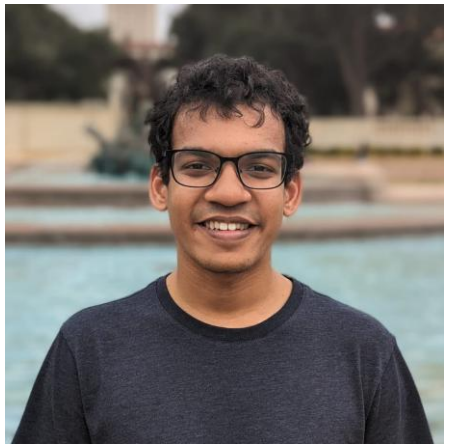
  - "Nice": outputs values in $[1/B, B]$ for a constant $B > 0$
  - Same proof works, assuming conjecture v2

Part 2: LSS is Optimal, Assuming Strong Detection is Hard

“From Weak to Strong Testing in Gaussian Models”

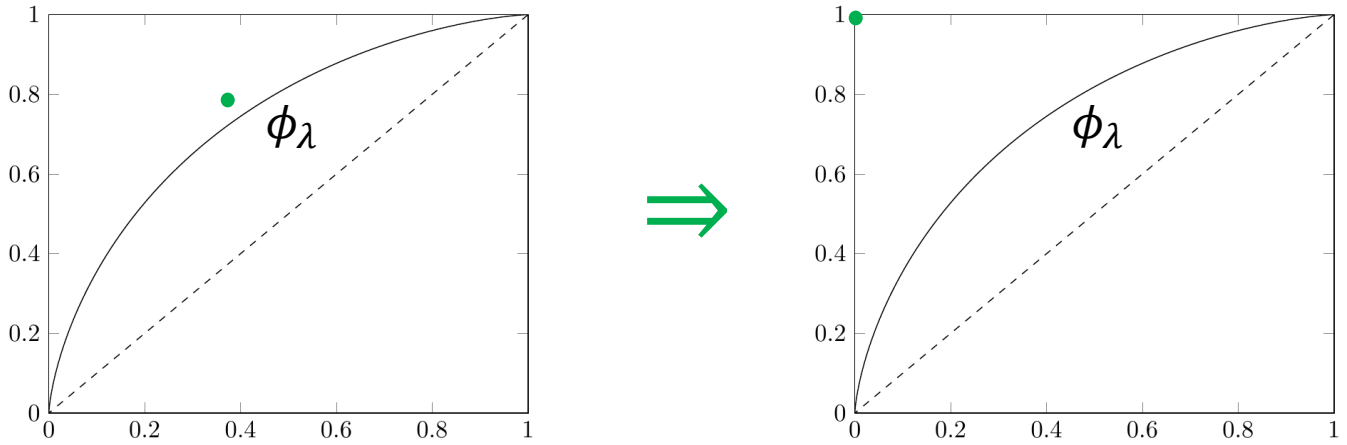
joint with Ansh Nagda

arXiv:2609.36050 (2026)



Main Result

- Strong LD Conjecture (v2) is true, assuming the "standard" LD Conjecture, i.e. *strong* detection is hard when $\lambda < 1$
- In other words, we prove *unconditionally* that for $\lambda < 1$:
  - *if* there is an algorithm beating the ROC curve of LSS (at a single point)
  - *then* there is an algorithm of similar runtime achieving *strong detection* for λ'



- Runtime of reduction is "almost polynomial": n^d for any $d = \omega(1)$
- Small loss in signal strength: λ' is any constant $\lambda < \lambda' < 1$

Self-Reducibility: Setup

- Fix a sequence $d = d_n \to \infty$ (slowly growing)
- Suppose we have an algorithm beating ROC curve of LSS
- From Part 1: can produce a nice $f: \mathbb{R}^{n \times n} \to \mathbb{R}$ such that

$$\frac{\mathrm{E}_{Y \sim \mathbb{P}(\lambda)}[f(Y)]}{\sqrt{\mathrm{E}_{Y \sim \mathbb{Q}}[f(Y)^2]}} > (1-\lambda^2)^{-1/4}$$ (beating degree-d polynomials)

- Goal: turn this into a strong distinguisher
- Inspired by "On optimal distinguishers for Planted Clique" [Nagda, Raghavendra '25]
  - Gives a related self-reduction for the planted clique problem in G(n,1/2)

Boosting the Ratio

[Nagda, Raghavendra '25]

- We have an algorithm computing a nice f such that

$$\frac{\mathrm{E}_{\mathbb{P}(\lambda)}[f]}{\sqrt{\mathrm{E}_{\mathbb{Q}}[f^2]}} \geq (1-\lambda^2)^{-1/4} + \Omega(1) \quad \text{(beating degree-}d\text{ polynomials)}$$

- Since f beats low-degree polynomials, its high-degree part is useful:

$$\frac{\mathrm{E}_{\mathbb{P}(\lambda)}[f_{>d}]}{\sqrt{\mathrm{E}_{\mathbb{Q}}[(f_{>d})^2]}} \geq \Omega(1)$$

- Consider the same ratio for "noisy" $Tf_{>d}$ (Ornstein–Uhlenbeck)
  - Numerator has small change in λ: $\mathrm{E}_{\mathbb{P}(\lambda')}[Tf_{>d}] = \mathrm{E}_{\mathbb{P}(\lambda)}[f_{>d}], \ \ \lambda' > \lambda$
  - Denominator shrinks by factor $\exp\big(\Omega(d)\big)$, since noise attenuates high degree
  - Ratio is now very large!

From Large Ratio to Strong Detection

- We now have an algorithm computing g such that
$$\frac{\mathrm{E}_{\mathbb{P}(\lambda')}[g]}{\sqrt{\mathrm{E}_{\mathbb{Q}}[g^2]}} \geq \exp\big(\Omega(d)\big)$$
- In general, large ratio need not imply strong detection
  - Why: numerator might be large due to rare event under $\mathbb{P}(\lambda')$
- In our setting, large ratio does imply a (1-sided) distinguisher
  - Test: threshold g
  - Analysis: bound first two moments of g under planted and null
- Ingredients:
  - f is nice
  - Every degree-d polynomial p has $\frac{\mathrm{E}_{\mathbb{P}(\lambda)}[p^2]}{\mathrm{E}_{\mathbb{Q}}[p^2]} \leq \mathrm{d}^{O(\sqrt{d})}$ for $\lambda < 1$

From 1-Sided to 2-Sided

- Previous argument gives a 1-sided distinguisher:
  - $\mathbb{Q}(t(Y)=1) \le \exp\left(-\Omega(d)\right)$ and $\mathbb{P}(t(Y)=1) \ge d^{-o(\sqrt{d})}$
- Additional step to boost this to (2-sided) strong detection
  - Idea: "clone" the input Y into many (correlated) copies
  - Each copy has slightly higher λ''
  - Run the 1-sided test on each copy, and output 1 if any copy outputs 1
  - The copies are highly correlated, but can still prove that the 1-sided test will accept at least one of them, with high probability under $\mathbb{P}$
- Technical difficulties that were not present in planted clique…
  - Here we care about sharp constants in λ
  - In our setting, weak detection is easy
  - Our signal is "inhomogeneous" (empirical distribution of x can vary)

Summary

- Self-reduction in spiked Wigner model: if you can beat ROC curve of LSS then you can boost this to strong detection
- The proof vindicates the Strong Low-Degree Conjecture (v2) and explains the importance of low-degree polynomials
- Proof ideas should (hopefully) transfer to other settings where precise error rates matter…

Thanks!