Average-Case Complexity and the Planted Clique Problem

Concentration, spectral algorithms, and random matrix norms

Part I

Worst-case and average-case complexity

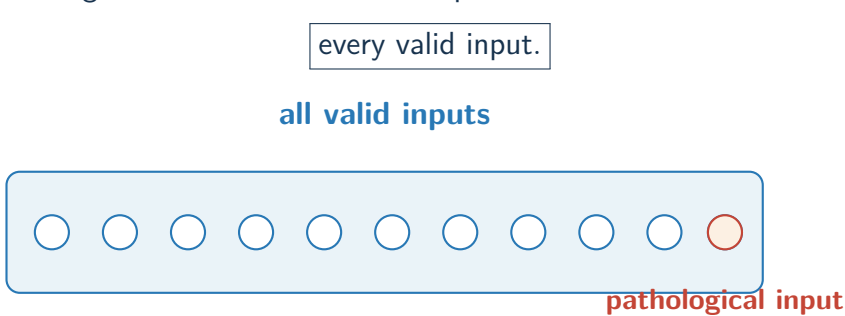
Worst-case algorithms handle every valid input

A classical computational problem specifies:

- a set of valid inputs;
- a set of acceptable outputs for each input.

A worst-case algorithm must return an acceptable answer for

every valid input.



A single difficult input can determine the worst-case complexity.

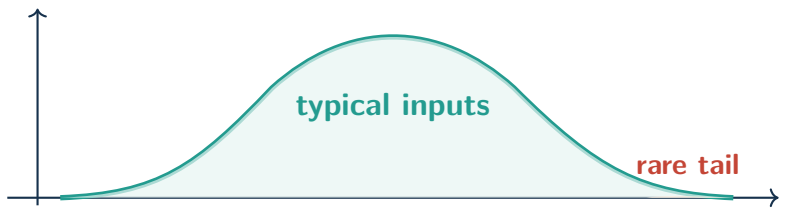
Average-case complexity includes an input distribution

A distributional problem specifies a probability distribution

$$\mathcal{D}_n$$

over inputs of size n.
The algorithm should succeed on a typical draw:

$$\Pr_{\substack{\text{algorithm randomness} \\ X \sim \mathcal{D}_n}} [A(X) \text{ is acceptable}] \geq \frac{2}{3}.$$



Random input models can reveal a different problem

Worst-case analysis remains essential, but it can hide useful structure.
Suppose:

- realistic inputs follow a specific random model;
- pathological instances have exponentially small probability;
- typical instances contain statistical patterns that an algorithm can exploit.

Average-case analysis asks two separate questions:

Can the distributions be distinguished at all?

Can they be distinguished in polynomial time?

The planted clique problem makes the gap between these questions especially clear.

Part II

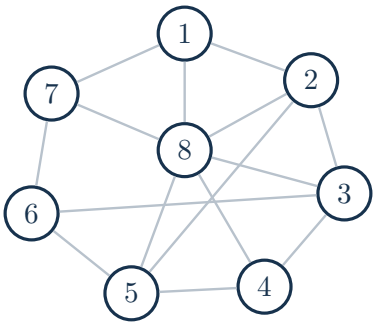
The planted clique problem

The null distribution is an Erdős–Rényi graph

Under the null hypothesis,

$$H_0 : \qquad G \sim G(n, 1/2).$$

Every unordered pair $\{i, j\}$ becomes an edge independently with probability $1/2$.



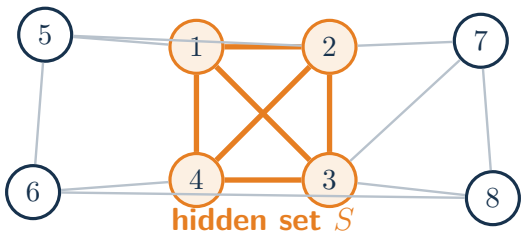
The planted distribution adds a hidden clique

Under the planted hypothesis,

$$H_1 : \qquad G \sim \mathcal{P}_1(n, k).$$

Generate G as follows:

1. sample $G_0 \sim G(n, 1/2)$;
2. choose $S \subseteq [n]$ uniformly among all k-element sets;
3. add every missing edge inside S.



Decision and recovery are different tasks

Decision problem

Construct a polynomial-time algorithm A such that

$$\Pr_{G \sim G(n,1/2)}[A(G) = 0] \geq \frac{2}{3}$$

and

$$\Pr_{G \sim \mathcal{P}_1(n,k)}[A(G) = 1] \geq \frac{2}{3}.$$

Recovery problem

Given

$$G \sim \mathcal{P}_1(n, k),$$

output the planted vertex set S with probability at least $2/3$.

Large cliques are extremely rare under the null

Let X_r be the number of r-vertex cliques in $G \sim G(n, 1/2)$.
For each fixed r-element set, all

$$\binom{r}{2}$$

internal edges must be present. Therefore

$$\mathbb{E}[X_r] = \binom{n}{r} 2^{-\binom{r}{2}}.$$

Markov's inequality gives

$$\Pr[X_r \geq 1] \leq \binom{n}{r} 2^{-\binom{r}{2}}.$$

For

$$r = \lceil 2 \log_2 n \rceil + 1,$$

this probability is at most $1/n$ for all sufficiently large n.

Planted clique is statistically detectable near $2 \log_2 n$

Under the planted distribution, the graph always contains a clique of size k.
Under the null distribution, with high probability,

$$\omega(G) \leq \lceil 2 \log_2 n \rceil ,$$

where $\omega(G)$ denotes the maximum clique size.
Therefore, for

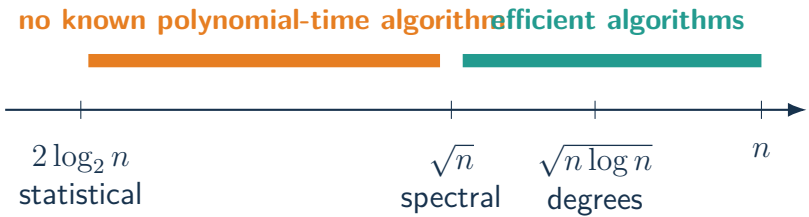
$$k > \lceil 2 \log_2 n \rceil ,$$

an exhaustive algorithm can distinguish the distributions:

1. compute a largest clique;
2. output H_1 if its size is at least k.

This establishes statistical distinguishability, but maximum clique search takes exponential time in the worst case.

The statistical and computational thresholds are far apart



The planted clique conjecture predicts polynomial-time hardness when

$$k = n^{1/2-\varepsilon}$$

for every fixed $\varepsilon > 0$.

Part III

A degree-based algorithm

Clique vertices have a larger expected degree

Under the null distribution,

$$\deg(v) \sim \mathrm{Bin}(n-1, 1/2),$$

so

$$\mu_0 := \mathbb{E}[\deg(v)] = \frac{n-1}{2}.$$

Under the planted distribution:

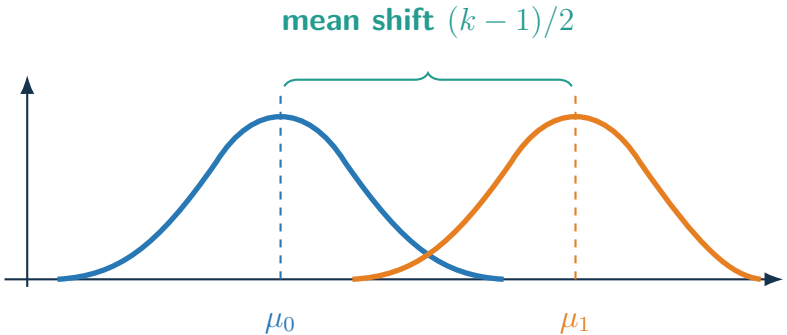
- if $v \notin S$, its degree distribution remains $\mathrm{Bin}(n-1, 1/2)$;
- if $v \in S$, then

$$\deg(v) = (k-1) + \mathrm{Bin}(n-k, 1/2),$$

and

$$\mu_1 = \frac{n+k-2}{2} = \mu_0 + \frac{k-1}{2}.$$

The degree distributions overlap unless the signal is large



To justify a degree threshold, we need a high-probability bound on the width of each distribution.

Part IV

Concentration inequalities

Markov's inequality converts expectation into a tail bound

Markov's inequality

If $X \geq 0$ and $a > 0$, then

$$\boxed{\Pr[X \geq a] \leq \frac{\mathbb{E}[X]}{a}.}$$

Define

$$Y := a\,\mathbf{1}_{\{X \geq a\}}.$$

Pointwise,

$$0 \leq Y \leq X.$$

Therefore

$$\mathbb{E}[X] \geq \mathbb{E}[Y] = a\,\mathbb{E}[\mathbf{1}_{\{X \geq a\}}] = a\,\Pr[X \geq a].$$

Dividing by a proves the inequality.

Chebyshev's inequality uses the variance

Chebyshev's inequality

For every random variable Y with finite variance and every $a > 0$,

$$\boxed{\Pr[|Y - \mathbb{E}[Y]| \geq a] \leq \frac{\mathrm{Var}(Y)}{a^2}.}$$

Apply Markov's inequality to the nonnegative random variable

$$Z := (Y - \mathbb{E}[Y])^2.$$

Then

$$\begin{aligned} \Pr[|Y - \mathbb{E}[Y]| \geq a] &= \Pr[Z \geq a^2] \\ &\leq \frac{\mathbb{E}[Z]}{a^2} \\ &= \frac{\mathrm{Var}(Y)}{a^2}. \end{aligned}$$

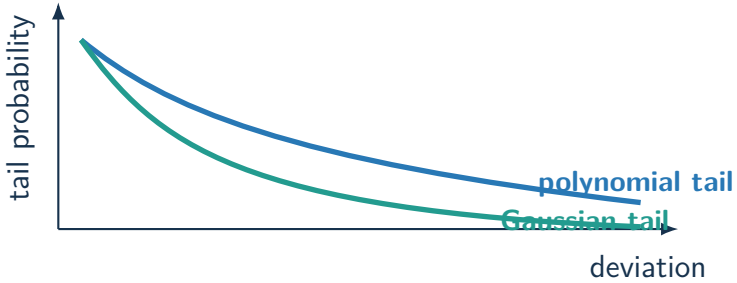
Polynomial tails are often too weak

Chebyshev gives a tail of order

$$\frac{1}{a^2}.$$

For sums of independent bounded variables, the true tail often behaves like

$$\exp(-ca^2).$$



The moment-generating function lets Markov's inequality see these exponentially small tails.

The moment-generating function

For a random variable Z, define

$$\boxed{\varphi_Z(\theta) := \mathbb{E}[e^{\theta Z}].}$$

If Z_1 and Z_2 are independent, then

$$\begin{aligned}\varphi_{Z_1+Z_2}(\theta) &= \mathbb{E}[e^{\theta(Z_1+Z_2)}] \\ &= \mathbb{E}[e^{\theta Z_1} e^{\theta Z_2}] \\ &= \mathbb{E}[e^{\theta Z_1}]\mathbb{E}[e^{\theta Z_2}] \\ &= \varphi_{Z_1}(\theta)\varphi_{Z_2}(\theta).\end{aligned}$$

Independence converts addition of variables into multiplication of their moment-generating functions.

Exponential Markov bounds one-sided tails

For every $\theta > 0$,

$$Z \geq t \quad \Longleftrightarrow \quad e^{\theta Z} \geq e^{\theta t}.$$

Markov's inequality gives

$$\begin{aligned} \Pr[Z \geq t] &= \Pr[e^{\theta Z} \geq e^{\theta t}] \\ &\leq \frac{\mathbb{E}[e^{\theta Z}]}{e^{\theta t}} \\ &= \boxed{\varphi_Z(\theta) e^{-\theta t}.} \end{aligned}$$

We are free to optimize over θ. This freedom produces the Gaussian-shaped exponent.

Subgaussian variables have Gaussian-size MGFs

Definition

A mean-zero random variable Z is σ-subgaussian if

$$\boxed{\mathbb{E}[e^{\theta Z}] \leq \exp\left(\frac{\sigma^2\theta^2}{2}\right)}$$

for every $\theta \in \mathbb{R}$.

If

$$Z \sim N(0, \sigma^2),$$

then equality holds:

$$\mathbb{E}[e^{\theta Z}] = \exp\left(\frac{\sigma^2\theta^2}{2}\right).$$

Independent subgaussian parameters add in squares

Suppose $Z_1, \ldots, Z_m$ are independent, mean-zero, and Z_i is σ_i-subgaussian. Then

$$\begin{aligned}\mathbb{E}\exp\left(\theta \sum_{i=1}^{m} Z_i\right) &= \prod_{i=1}^{m} \mathbb{E}[e^{\theta Z_i}] \\ &\leq \prod_{i=1}^{m} \exp\left(\frac{\sigma_i^2 \theta^2}{2}\right) \\ &= \exp\left(\frac{\theta^2}{2} \sum_{i=1}^{m} \sigma_i^2\right).\end{aligned}$$

Thus

$$\boxed{\sum_i Z_i \text{ is } \sqrt{\sum_i \sigma_i^2}\text{-subgaussian.}}$$

Bounded mean-zero variables are subgaussian

Hoeffding's lemma

If

$$\mathbb{E}[Z] = 0 \qquad \text{and} \qquad Z \in [a, b],$$

then

$$\boxed{\mathbb{E}[e^{\theta Z}] \leq \exp\left(\frac{\theta^2 (b-a)^2}{8}\right).}$$

In particular, if

$$|Z| \leq b,$$

then

$$\mathbb{E}[e^{\theta Z}] \leq e^{b^2 \theta^2 / 2}.$$

Therefore every mean-zero variable supported in $[-b, b]$ is b-subgaussian.

Convexity gives the key step in Hoeffding's lemma

For every $z \in [a, b]$, convexity of the exponential gives

$$e^{\theta z} \leq \frac{b - z}{b - a} e^{\theta a} + \frac{z - a}{b - a} e^{\theta b}.$$

Taking expectations and using $\mathbb{E}[Z] = 0$,

$$\mathbb{E}[e^{\theta Z}] \leq \frac{b}{b - a} e^{\theta a} - \frac{a}{b - a} e^{\theta b}.$$

The right-hand side is the MGF of a suitable two-point random variable. A one-variable calculus argument bounds it by

$$\exp\left(\frac{\theta^2 (b - a)^2}{8}\right).$$

Extremal bounded variables therefore behave no worse than a two-point distribution.

The MGF bound yields an exponential tail

If Z is σ-subgaussian, then for every $\theta > 0$,

$$\begin{aligned}\Pr[Z \geq t] &\leq \mathbb{E}[e^{\theta Z}]e^{-\theta t} \\ &\leq \exp\left(\frac{\sigma^2\theta^2}{2} - \theta t\right).\end{aligned}$$

The exponent

$$f(\theta) = \frac{\sigma^2\theta^2}{2} - \theta t$$

is minimized at

$$\theta^\star = \frac{t}{\sigma^2}.$$

Substitution gives

$$\Pr[Z \geq t] \leq \exp\left(-\frac{t^2}{2\sigma^2}\right).$$

Hoeffding's inequality for bounded sums

Applying the same argument to $-Z$ and using a union bound gives

$$\Pr[|Z| \geq t] \leq 2\exp\left(-\frac{t^2}{2\sigma^2}\right).$$

Therefore:

Hoeffding-type bound

Let $X_1, \ldots, X_m$ be independent random variables with

$$\mathbb{E}[X_i] = 0, \qquad |X_i| \leq b_i.$$

Set

$$B^2 := \sum_{i=1}^{m} b_i^2.$$

Then

Normalized deviations have Gaussian tails

Writing

$$t = \lambda B,$$

Hoeffding's inequality becomes

$$\boxed{\Pr\left[\left|\sum_i X_i\right| \geq \lambda B\right] \leq 2e^{-\lambda^2/2}.}$$

The scale

$$B = \sqrt{\sum_i b_i^2}$$

is deterministic.

Independence is essential. Without it, perfectly correlated bounded variables can add coherently and produce fluctuations of order m instead of order $\sqrt{m}$.

Part V

Concentration of graph degrees

A fixed null degree is a sum of Bernoulli variables

Fix $v \in [n]$. Under $G(n, 1/2)$,

$$\deg(v) = \sum_{u \neq v} Z_{uv},$$

where the Z_{uv} are independent $\mathrm{Bernoulli}(1/2)$ variables. Define

$$X_{uv} := Z_{uv} - \frac{1}{2}.$$

Then

$$\mathbb{E}[X_{uv}] = 0, \qquad |X_{uv}| \leq \frac{1}{2},$$

and

$$\deg(v) = \frac{n-1}{2} + \sum_{u \neq v} X_{uv}.$$

A fixed degree fluctuates by $O(\sqrt{n})$

Here

$$B^2 = \sum_{u \neq v} \left(\frac{1}{2}\right)^2 = \frac{n-1}{4}.$$

Hoeffding's inequality gives

$$\Pr\left[\left|\deg(v) - \frac{n-1}{2}\right| \geq t\right] \leq 2\exp\left(-\frac{2t^2}{n-1}\right).$$

Taking

$$t = C\sqrt{n \log n}$$

makes the failure probability

$$2n^{-\Omega(C^2)}.$$

A union bound controls every vertex simultaneously

Let

$$\mathcal{B}_v = \left\{ \left| \deg(v) - \frac{n-1}{2} \right| > C\sqrt{n \log n} \right\}.$$

Then

$$\Pr[\mathcal{B}_v] \leq 2n^{-\Omega(C^2)}.$$

By the union bound,

$$\Pr[\exists v : \mathcal{B}_v] \leq \sum_{v=1}^{n} \Pr[\mathcal{B}_v] \leq 2n^{1-\Omega(C^2)}.$$

For a sufficiently large constant C, with high probability every null degree lies in

$$\frac{n-1}{2} \pm C\sqrt{n \log n}.$$

Planted clique degrees concentrate around a shifted mean

For $v \in S$,

$$\deg(v) = (k-1) + \sum_{u \notin S} Z_{uv},$$

where

$$Z_{uv} \sim \text{Bernoulli}(1/2)$$

independently.

Hence

$$\mathbb{E}[\deg(v)] = k - 1 + \frac{n-k}{2} = \frac{n+k-2}{2}.$$

The random part contains at most n bounded independent variables, so with high probability every planted vertex satisfies

$$\deg(v) = \frac{n+k-2}{2} \pm C\sqrt{n \log n}.$$

A threshold separates the two degree ranges

Let

$$\tau := \frac{n-1}{2} + \frac{k-1}{4}.$$

Under the null hypothesis, with high probability,

$$\max_v \deg(v) \leq \frac{n-1}{2} + C\sqrt{n \log n}.$$

Under the planted hypothesis, every $v \in S$ has

$$\deg(v) \geq \frac{n-1}{2} + \frac{k-1}{2} - C\sqrt{n \log n}.$$

These intervals lie on opposite sides of τ when

$$k \geq C_0\sqrt{n \log n}$$

for a sufficiently large constant C_0.

The degree algorithm

Degree test

1. Compute every vertex degree.
2. Output H_1 if some vertex has degree at least

$$\tau = \frac{n-1}{2} + \frac{k-1}{4}.$$

3. Otherwise output H_0.

For

$$k \geq C_0 \sqrt{n \log n},$$

the algorithm distinguishes the two distributions with probability tending to one. Selecting all vertices above the threshold also recovers the planted clique in the same parameter range.

Part VI

The spectral planted clique algorithm

Center the adjacency matrix around zero

Define the signed adjacency matrix $M \in \mathbb{R}^{n \times n}$ by

$$M_{ii} = 0,$$

and, for $i \neq j$,

$$M_{ij} = \begin{cases} +1, & \{i,j\} \in E, \\ -1, & \{i,j\} \notin E. \end{cases}$$

Under $G(n, 1/2)$, the entries above the diagonal are independent Rademacher random variables:

$$\Pr[M_{ij} = 1] = \Pr[M_{ij} = -1] = \frac{1}{2}.$$

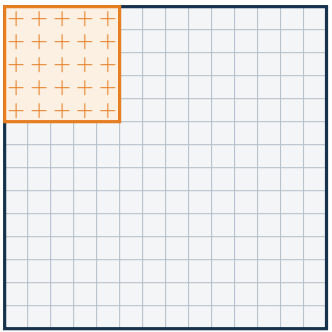
Thus

$$\mathbb{E}[M_{ij}] = 0.$$

Planting creates a coherent positive block

After permuting the vertices so that S comes first, the matrix has the schematic form

$$M = \begin{pmatrix} \text{all } +1 \text{ off diagonal} & \text{random signs} \\ \text{random signs} & \text{random signs} \end{pmatrix}.$$



planted $k \times k$ block

centered random noise

A k-clique forces a large eigenvalue

Let S be a k-clique and let $\mathbf{1}_S$ be its indicator vector.
Since every ordered pair of distinct vertices in S contributes $+1$,

$$\mathbf{1}_S^{\mathsf{T}} M \mathbf{1}_S = k(k-1).$$

Also,

$$\|\mathbf{1}_S\|_2^2 = k.$$

Since M is symmetric,

$$\|M\|_{\mathrm{op}} = \max_{\|x\|_2=1} |x^{\mathsf{T}} M x|.$$

Therefore

$$\boxed{\|M\|_{\mathrm{op}} \geq \frac{\mathbf{1}_S^{\mathsf{T}} M \mathbf{1}_S}{\|\mathbf{1}_S\|_2^2} = k-1.}$$

A null signed adjacency matrix has norm $O(\sqrt{n})$

Random matrix theorem

Let M be symmetric with $M_{ii} = 0$, and let the entries $M_{ij} = M_{ji}$, $i < j$, be independent uniform signs. Then

$$\boxed{\|M\|_{\mathrm{op}} \leq C\sqrt{n}}$$

with probability at least

$$1 - e^{-cn}$$

for universal constants $C, c > 0$.

The entries have size 1, but cancellations keep the largest eigenvalue at scale $\sqrt{n}$, rather than scale n.

We will prove the theorem using an ε-net.

The spectral decision algorithm

Spectral test

1. Construct the signed adjacency matrix M.
2. Compute $\|M\|_{\mathrm{op}}$.
3. Output H_1 if
$$\|M\|_{\mathrm{op}} > C\sqrt{n}.$$
4. Otherwise output H_0.

Under the null,
$$\|M\|_{\mathrm{op}} \leq C\sqrt{n}$$
with high probability.
Under the planted distribution,
$$\|M\|_{\mathrm{op}} \geq k - 1.$$
Thus the algorithm succeeds for

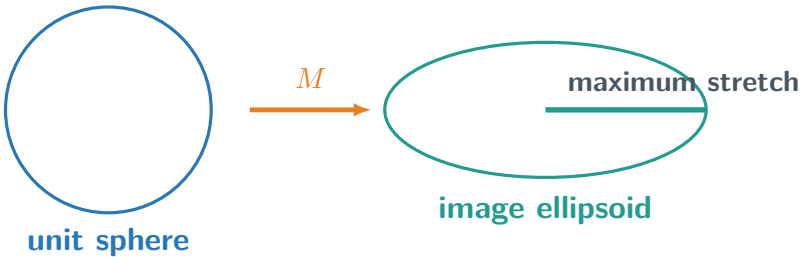
Part VII

Discretizing the unit sphere

The operator norm maximizes over infinitely many vectors

For symmetric M,

$$\|M\|_{\mathrm{op}} = \sup_{x \in \mathbb{S}^{n-1}} |x^{\mathsf{T}} M x|.$$



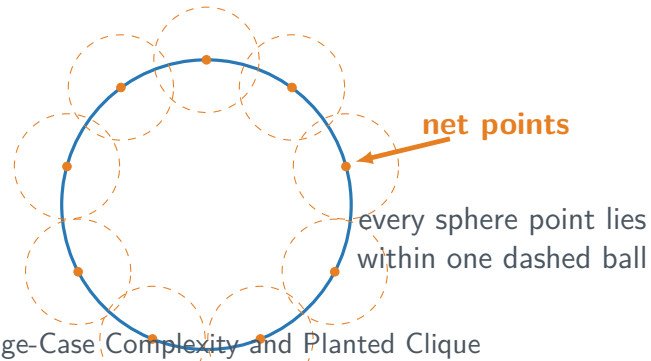
A concentration bound for one fixed vector does not directly control this uncountable supremum. We need a finite approximation to the sphere.

An ε-net covers the sphere

Definition

A set $\mathcal{N} \subseteq \mathbb{S}^{n-1}$ is an ε-net if every $x \in \mathbb{S}^{n-1}$ has some $y \in \mathcal{N}$ satisfying

$$\|x - y\|_2 \leq \varepsilon.$$



Covering numbers measure the size of a net

For $A \subseteq \mathbb{R}^n$, define

$$N(A, \varepsilon)$$

to be the minimum size of an ε-net contained in A.
Equivalently, $N(A, \varepsilon)$ is the smallest number of radius ε balls, centered in A, that cover A.
Our goal is an exponential-size net:

$$N(\mathbb{S}^{n-1}, 1/10) \leq 2^{O(n)}.$$

This may look enormous, but a fixed-vector tail of size

$$e^{-\Omega(n)}$$

can survive a union bound over $2^{O(n)}$ points.

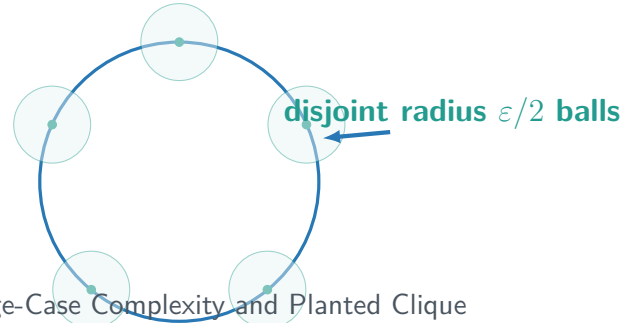
A packing keeps all selected points separated

Definition

A set $\mathcal{P} \subseteq A$ is an ε-packing if

$$\|x - y\|_2 > \varepsilon$$

for every distinct $x, y \in \mathcal{P}$.



A maximal packing automatically covers the set

Covering-packing lemma

Every inclusion-maximal ε-packing $\mathcal{P} \subseteq A$ is an ε-net for A.

Suppose instead that some $x \in A$ satisfies

$$\|x - y\|_2 > \varepsilon \qquad \text{for every } y \in \mathcal{P}.$$

Then

$$\mathcal{P} \cup \{x\}$$

remains an ε-packing, contradicting maximality.
Therefore

$$\boxed{N(A, \varepsilon) \leq P(A, \varepsilon).}$$

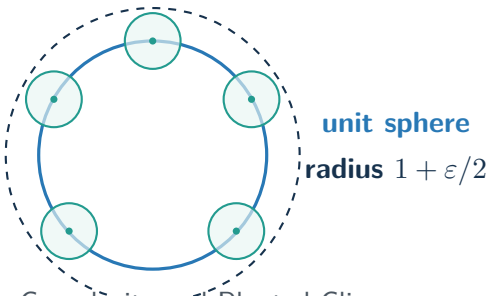
Disjoint small balls fit inside one larger ball

Let $\mathcal{P}$ be an ε-packing in $\mathbb{S}^{n-1}$.
The balls

$$B(p, \varepsilon/2), \qquad p \in \mathcal{P},$$

are pairwise disjoint.
Since $\|p\|_2 = 1$, each such ball lies inside

$$B(0, 1 + \varepsilon/2).$$



Volume bounds the packing number

Since the small balls are disjoint,

$$|\mathcal{P}| \ \mathrm{Vol}\Big(B(0,\varepsilon/2)\Big) \leq \mathrm{Vol}\Big(B(0,1+\varepsilon/2)\Big).$$

Volume in $\mathbb{R}^n$ scales as the n-th power of the radius. Therefore

$$\begin{aligned} |\mathcal{P}| &\leq \frac{\mathrm{Vol}(B(0,1+\varepsilon/2))}{\mathrm{Vol}(B(0,\varepsilon/2))} \\ &= \left(\frac{1+\varepsilon/2}{\varepsilon/2}\right)^n \\ &= \boxed{\left(1+\frac{2}{\varepsilon}\right)^n}. \end{aligned}$$

A constant-radius net has exponential size

Combining the packing and covering bounds,

$$\boxed{N(\mathbb{S}^{n-1}, \varepsilon) \leq \left(1 + \frac{2}{\varepsilon}\right)^n.}$$

For

$$\varepsilon = \frac{1}{10},$$

there exists a net $\mathcal{N} \subseteq \mathbb{S}^{n-1}$ satisfying

$$|\mathcal{N}| \leq 21^n = 2^{O(n)}.$$

This is the right scale: the sphere is continuous, but constant-accuracy discretization costs only exponentially many points.

Part VIII

The random matrix norm bound

A fixed quadratic form is a bounded independent sum

Fix $x \in \mathbb{S}^{n-1}$. Since M is symmetric with zero diagonal,

$$x^{\mathsf{T}} M x = 2 \sum_{i<j} M_{ij} x_i x_j.$$

The variables M_{ij}, $i < j$, are independent uniform signs.
Each summand

$$2 M_{ij} x_i x_j$$

has mean zero and magnitude at most

$$2|x_i x_j|.$$

Hoeffding's inequality applies once we bound the sum of the squared coefficients.

The coefficient squares sum to at most 2

We have

$$\begin{aligned}\sum_{i<j}(2x_ix_j)^2 &= 4\sum_{i<j}x_i^2x_j^2 \\ &= 2\left[\left(\sum_i x_i^2\right)^2 - \sum_i x_i^4\right] \\ &\leq 2\left(\sum_i x_i^2\right)^2 \\ &= 2.\end{aligned}$$

Therefore, for every fixed $x \in \mathbb{S}^{n-1}$,

$$\boxed{\Pr[|x^\mathsf{T}Mx| \geq t] \leq 2e^{-t^2/4}.}$$

The tail bound does not depend on the dimension.

The net union bound controls exponentially many vectors

Let $\mathcal{N}$ be a $1/10$-net with

$$|\mathcal{N}| \le 21^n.$$

For $t = C\sqrt{n}$,

$$\Pr[|x^{\mathsf{T}} M x| > C\sqrt{n}] \le 2e^{-C^2 n/4}$$

for each fixed $x \in \mathcal{N}$.
A union bound gives

$$\Pr\left[\exists x \in \mathcal{N} : |x^{\mathsf{T}} M x| > C\sqrt{n}\right] \le 2 \cdot 21^n e^{-C^2 n/4}.$$

For sufficiently large C, this is

$$e^{-\Omega(n)}.$$

A bound on the net controls the whole sphere

Define

$$T := \max_{x \in \mathcal{N}} |x^{\mathsf{T}} M x|.$$

Choose $u \in \mathbb{S}^{n-1}$ such that

$$\|M\|_{\mathrm{op}} = |u^{\mathsf{T}} M u|.$$

By the net property, there exists $x \in \mathcal{N}$ with

$$\|u - x\|_2 \le \varepsilon, \qquad \varepsilon = \frac{1}{10}.$$

Write

$$h := u - x.$$

Then

$$\|h\|_2 \le \varepsilon.$$

Approximation errors are proportional to the unknown norm

Using

$$u^\mathsf{T} M u - x^\mathsf{T} M x = h^\mathsf{T} M u + x^\mathsf{T} M h,$$

we get

$$\begin{aligned} \|M\|_{\mathrm{op}} &= |u^\mathsf{T} M u| \\ &\leq |x^\mathsf{T} M x| + |h^\mathsf{T} M u| + |x^\mathsf{T} M h| \\ &\leq T + \varepsilon \|M\|_{\mathrm{op}} + \varepsilon \|M\|_{\mathrm{op}}. \end{aligned}$$

Hence

$$(1 - 2\varepsilon) \|M\|_{\mathrm{op}} \leq T.$$

The random matrix theorem follows

With $\varepsilon = 1/10$,

$$\|M\|_{\mathrm{op}} \le \frac{T}{1-2\varepsilon} = \frac{5}{4}T.$$

With probability at least $1 - e^{-\Omega(n)}$, the net union bound gives

$$T \le C\sqrt{n}.$$

Therefore

$$\boxed{\|M\|_{\mathrm{op}} \le \frac{5}{4}C\sqrt{n} = O(\sqrt{n})}$$

with exponentially high probability.
This completes the proof of the random matrix norm bound.

The spectral method reaches the $\sqrt{n}$ scale

Under the null distribution,

$$\|M\|_{\mathrm{op}} = O(\sqrt{n})$$

with high probability.
Under the planted distribution,

$$\|M\|_{\mathrm{op}} \geq k - 1.$$

Hence a polynomial-time spectral algorithm distinguishes the two models whenever

$$k \geq C\sqrt{n}.$$

This improves the degree threshold

$$C\sqrt{n \log n}$$

by removing the factor $\sqrt{\log n}$.
No polynomial-time algorithm is known for planted cliques of size $n^{1/2-\varepsilon}$.

The proof combines geometry and probability



The probability of failure for one direction is exponentially small. Geometry reduces the continuous sphere to exponentially many directions. The two exponential scales balance.

Main ideas

1. Average-case problems include a distribution over inputs.
2. Planted clique is statistically detectable for $k \gtrsim 2 \log_2 n$, but efficient algorithms need a larger clique.
3. Moment-generating functions turn Markov's inequality into Gaussian tail bounds.
4. Degree concentration yields an algorithm for

$$k \gtrsim \sqrt{n \log n}.$$

5. A clique creates a signed-adjacency eigenvalue of size $k - 1$.
6. Nets and concentration show that null eigenvalues have size $O(\sqrt{n})$.
7. The spectral algorithm succeeds for

$$k \gtrsim \sqrt{n}.$$

Typical inputs can contain structure that worst-case analysis does not expose.

Concentration makes that structure algorithmically visible.