Understanding Graphs through Spectral Densities

David Bindel

17 July 2019

Department of Computer Science
Cornell University

Acknowledgements

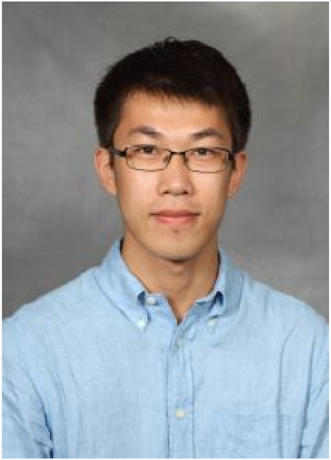

Thanks **Kun Dong** and Austin Benson, along with Anna Yesypenko, Moteleolu Onabajo, Jianqiu Wang.

Also: NSF DMS-1620038.

Eigenvalues Two Ways

What can we tell from *partial* spectral information (eigenvalues and/or vectors)

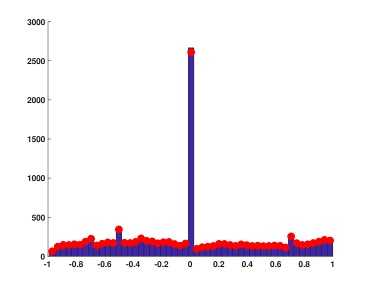

Claim: Most spectral analyses involve one of two perspectives:

- Approximate something via a *few* (extreme) eigenvalues.
- Look at *all* the eigenvalues (or all in a range).

Can One Hear the Shape of a Drum?

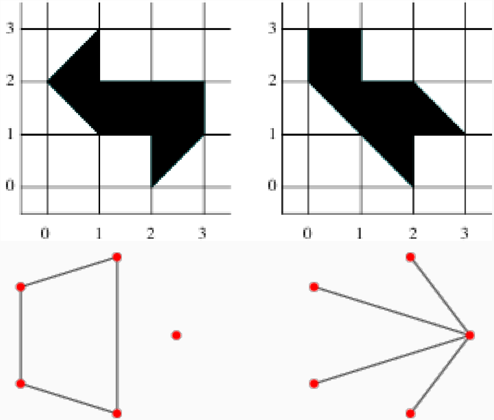

"You mean, if you had perfect pitch could you find the shape of a drum." — Mark Kac (quoting Lipmann Bers)
American Math Monthly, *1966*

What Do You Hear?

What information hides in the eigenvalue distribution?

1. Discretizations of Laplacian: something like Weyl's law
2. Sparse E-R random graphs: Wigner semicircular law
3. Some other random graphs: Wigner semicircle + a bit (Farkas *et al*, Phys Rev E (64), 2001)
4. "Real" networks: less well understood

Goal: Explore by estimating eigenvalue distributions (fast).

A Bestiary of Matrices

- Adjacency matrix: A
- Laplacian matrix: $L = D - A$
- Unsigned Laplacian: $L = D + A$
- **Random walk matrix**: $P = AD^{-1}$ (or $D^{-1}A$)
- **Normalized adjacency**: $\bar{A} = D^{-1/2}AD^{-1/2}$
- **Normalized Laplacian**: $\bar{L} = I - \bar{A} = D^{-1/2}LD^{-1/2}$
- Modularity matrix: $B = A - \frac{dd^T}{2n}$
- Motif adjacency: $W = A^2 \odot A$

All have examples of co-spectral graphs

… through spectrum uniquely identifies *quantum graphs*

Density of States

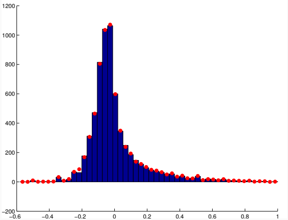

Spectra define a *generalized function* (a *density*):

$$\mathrm{tr}(f(H)) = \int f(\lambda)\mu(\lambda)\,dx = \sum_{k=1}^{N} f(\lambda_k)$$

where f is an analytic test function. Smooth to get a picture: a *spectral histogram* or *kernel density estimate*.

Exploring Spectral Densities

Kernel polynomial method (see Weisse, Rev. Modern Phys.)

- Spectral distribution on $[-1, 1]$ is a generalized function:

$$\int_{-1}^{1} \mu(x) f(x)\, dx = \frac{1}{N} \sum_{k=1}^{N} f(\lambda_k)$$

- Write $f(x) = \sum_{j=1}^{\infty} c_j T_j(x)$ and $\mu(x) = \sum_{j=1}^{\infty} d_j \phi_j(x)$, where $\int_{-1}^{1} \phi_j(x) T_k(x)\, dx = \delta_{jk}$
- Estimate $d_j = \operatorname{tr}(T_j(H))$ by stochastic methods
- Truncate series for $\mu(x)$ and filter (avoid Gibbs)

Much cheaper than computing all eigenvalues!

Alternatives: Lanczos (Golub-Meurant), maxent (Röder-Silver)

Stochastic Trace and Diagonal Estimation

$Z \in \mathbb{R}^n$ with independent entries, mean 0 and variance 1.

$$\begin{aligned} E[(Z \odot HZ)_i] &= \sum_j h_{ij} E[Z_i Z_j] = h_{ii} \\ \mathrm{Var}[(Z \odot HZ)_i] &= \sum_j h_{ij}^2. \end{aligned}$$

Serves as the basis for stochastic estimation of

- Trace (Hutchinson, others; review by Toledo and Avron)
- Diagonal (Bekas, Kokiopoulou, and Saad)

Independent probes $\implies$ $1/\sqrt{N}$ convergence (usual MC).
(Can go beyond independent probes.)

Example: PGP Network

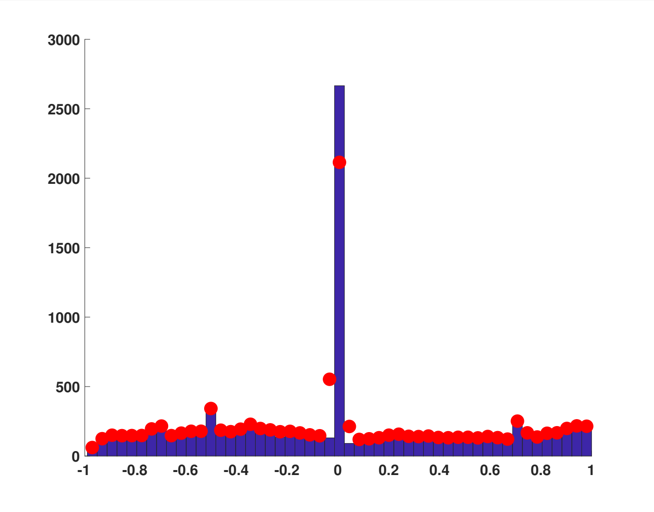

Spike (non-smoothness) at eigenvalues of 0 leads to inaccurate approximation.

Suppose $PH = HP$. Then

$$\mathcal{V} \text{ a max invariant subspace for } P \implies$$
$$\mathcal{V} \text{ a max invariant subspace for } H$$

So *local symmetry* $\implies$ *localized eigenvectors*.

Simplest example: P swaps (i, j)

- $e_i - e_j$ an eigenvector of P with eigenvalue -1
- $e_i - e_j$ an eigenvector of $\bar{A}$ with eigenvalue

$$\lambda = \rho_{\bar{A}}(e_i - e_j) = \begin{cases} d^{-1}, & (i, j) \in \mathcal{E} \\ 0, & \text{otherwise.} \end{cases}$$

- All other eigenvectors (eigenvalue -1) satisfy $v_i = v_j$

Motifs in Spectrum

- $\lambda = 0$

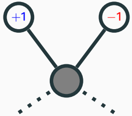

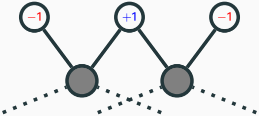

- $\lambda = \pm 1/2$

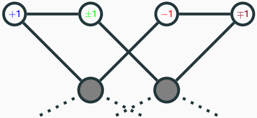

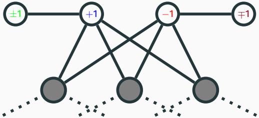

- $\lambda = -1/2$

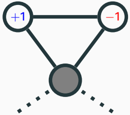

$\lambda = \pm 1/\sqrt{2}$

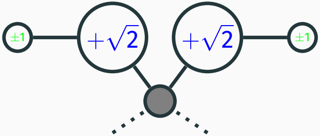

Motif Filtering

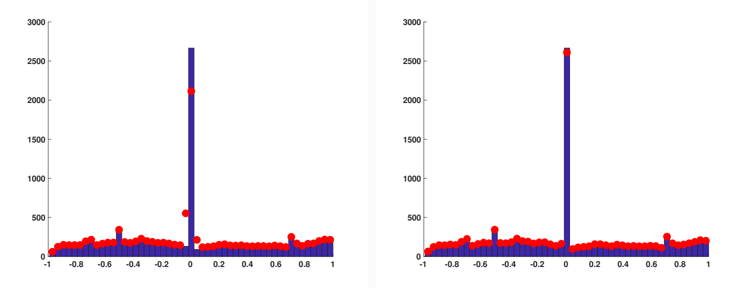

Motif "spikes" slow convergence – deflate motif eigenvectors!

If $P \in \mathbb{R}^{n \times m}$ an orthonormal basis for the quotient space,

- Apply estimator to $P^T \overline{A} P$ to reduce size for $m \ll n$.
- or use $Proj_P(Z)$ to probe the desired subspace.

Diagonal Estimation and LDoS

Diagonal estimation also useful for *local* DoS $\nu_k(x)$;
in the symmetric case with $H = Q\Lambda Q^T$, have

$$\int f(x)\nu_k(x)\,dx = f(H)_{kk} = e_k^T Q f(\Lambda) Q^T e_k$$

$$\nu_k(x) = \sum_{j=1}^{n} q_{kj}^2\,\delta(x - \lambda_j)$$

DoS is sum of local densities of states:

$$\mu(x) = \sum_{k=1}^{n} \nu_k(x)$$

KPM for LDoS

Same game, different moments:

- Estimate $d_j = [T_j(H)]_{kk}$ by diag estimation
- Truncate series for $\mu(x)$ and filter (avoid Gibbs)

Diagonal estimator gives moments *for all k simultaneously*!

Alternatives: Lanczos (Golub-Meurant), maxent (Röder-Silver)

LDoS Information

Can compute common *centrality measures* with LDoS

- Estrada centrality: $\exp(\gamma A)_{kk}$
- Resolvent centrality: $\left[(I - \gamma \bar{A})^{-1}\right]_{kk}$

Some motifs associated with localized eigenvectors:

- Chief example: Null vectors of $\bar{A}$ supported on leaves.
- Use LDoS + topology to find motifs?

Other uses: clustering and role discovery. What else?

Exploring Spectral Densities (with David Gleich)

- Compute spectrum of normalized Laplacian / RW matrix
- Compare KPM to full eigencomputation

Things we know

- Eigenvalues in $[-1, 1]$; nonsymmetric in general
- Stability: change d edges, have

$$\lambda_{j-d} \leq \hat{\lambda}_j \leq \lambda_{j+d}$$

- kth moment = P(return after k-step random walk)
- Eigenvalue cluster near 1 $\sim$ well-separated clusters
- Eigenvalue cluster near -1 $\sim$ bipartite structure
- Eigenvalue cluster near 0 $\sim$ leaf clusters

What else can we “hear”?

Erdos

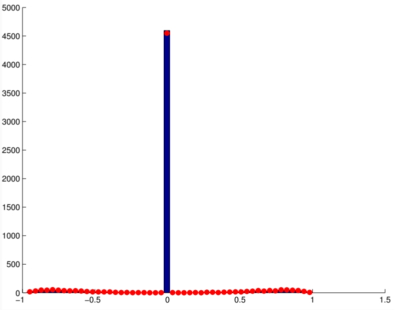
5000
4500
4000
3500
3000
2500
2000
1500
1000
500
0
−1
−0.5
0
0.5
1
1.5

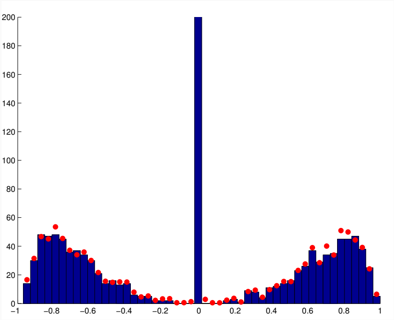
200
180
160
140
120
100
80
60
40
20
0
−1
−0.8
−0.6
−0.4
−0.2
0
0.2
0.4
0.6
0.8
1

Erdos (local)

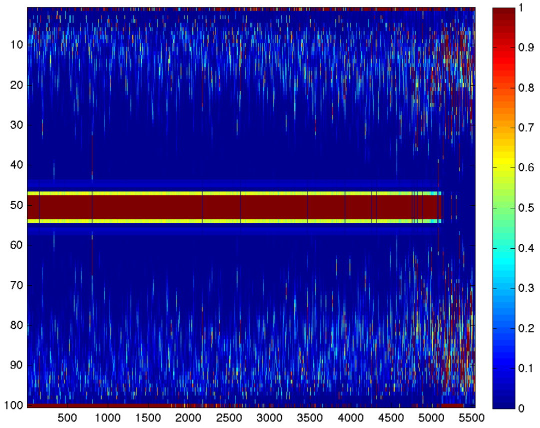
10
20
30
40
50
60
70
80
90
100
500
1000
1500
2000
2500
3000
3500
4000
4500
5000
5500
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0

Internet topology

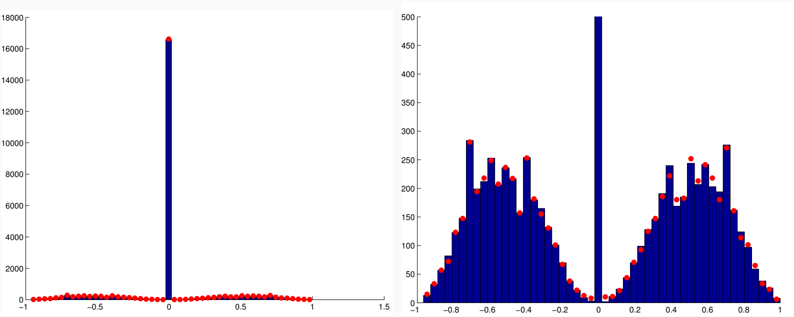

Internet topology (local)

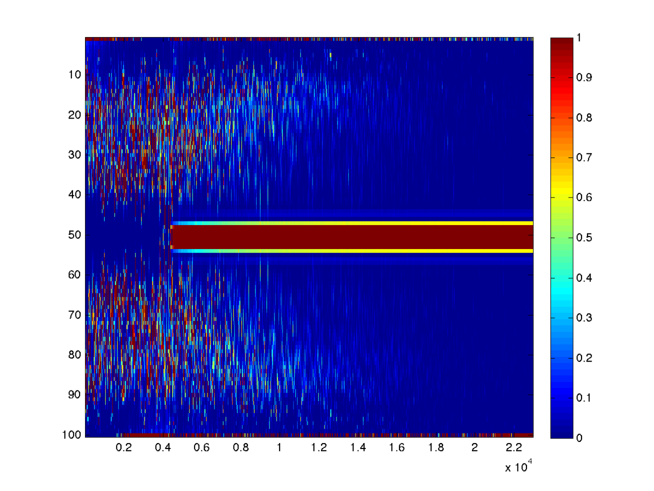

PGP

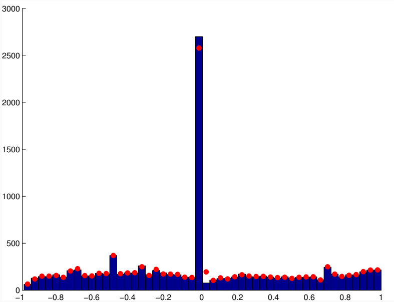
3000
2500
2000
1500
1000
500
0
−1
−0.8
−0.6
−0.4
−0.2
0
0.2
0.4
0.6
0.8
1

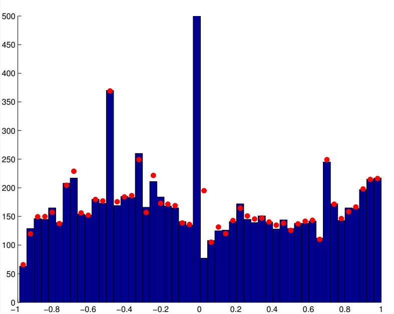
500
450
400
350
300
250
200
150
100
50
0
−1
−0.8
−0.6
−0.4
−0.2
0
0.2
0.4
0.6
0.8
1

PGP (local)

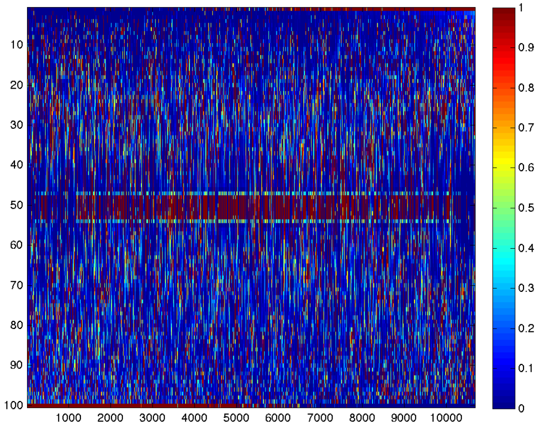
10
20
30
40
50
60
70
80
90
100
1000
2000
3000
4000
5000
6000
7000
8000
9000
10000
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0

Yeast

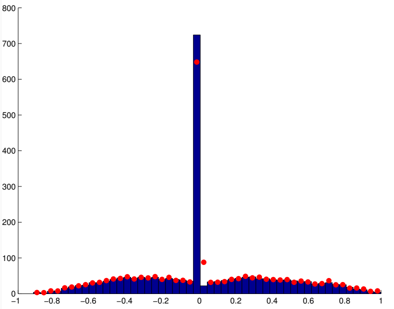

800
700
600
500
400
300
200
100
0
−1
−0.8
−0.6
−0.4
−0.2
0
0.2
0.4
0.6
0.8
1

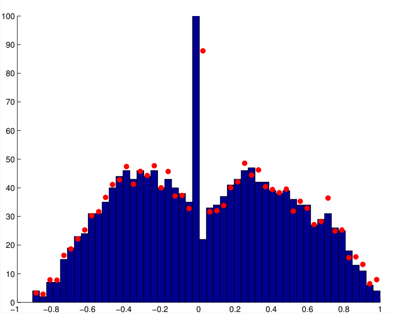

100
90
80
70
60
50
40
30
20
10
0
−1
−0.8
−0.6
−0.4
−0.2
0
0.2
0.4
0.6
0.8
1

Yeast (local)

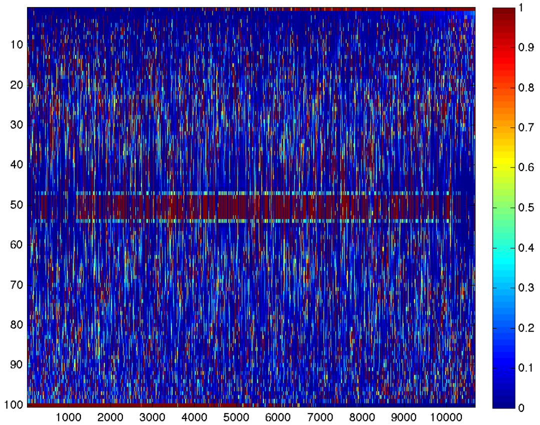
10
20
30
40
50
60
70
80
90
100
1000
2000
3000
4000
5000
6000
7000
8000
9000
10000
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0

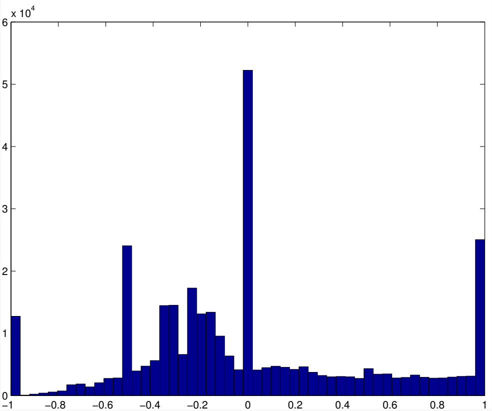

$N = 326186$, $nnz = 1615400$, 80 s (1000 moments, 10 probes)

Hollywood 2009 (LAW)

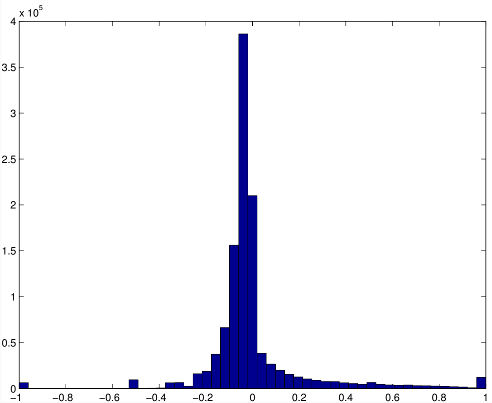

$N = 1139905$, $nnz = 113891327$, 2093 s (1000 moments, 10 probes)

Barabási–Albert model

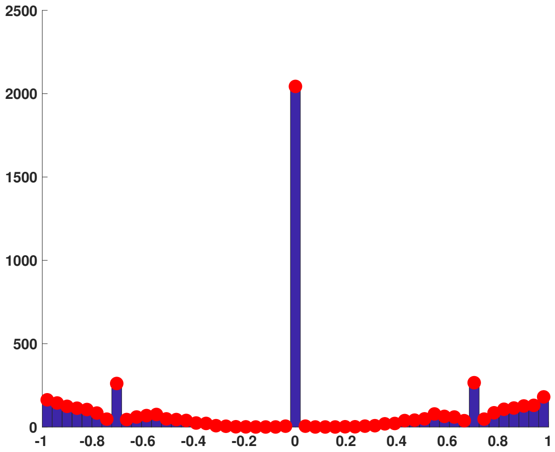

Scale-free network (5000 nodes, 4999 edges)

Watts–Strogatz

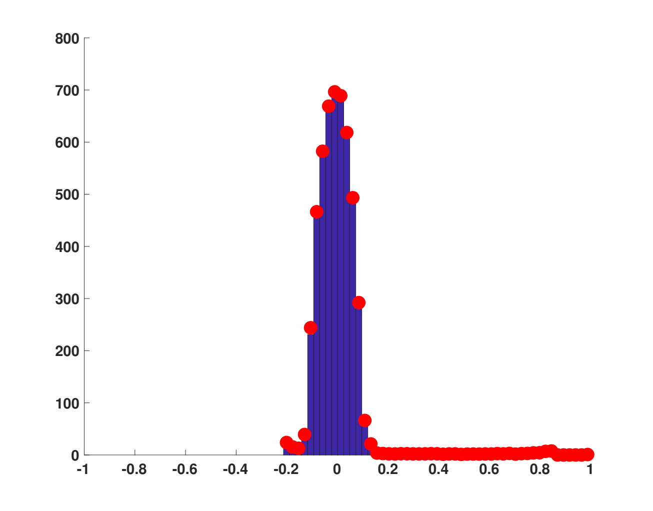

Small world network (5000 nodes, 260000 edges)

Model Verification: BTER

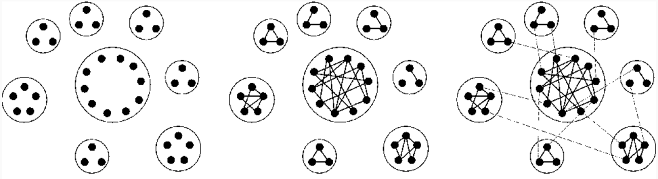

(a) Preprocessing: Distribution of nodes into affinity blocks

(b) Phase 1: Local links within each affinity block

(c) Phase 2: Global links across affinity blocks

Kolda et al, SISC (36), 2014

Block Two-Level Erdős-Rényi model (BTER)

- First Phase: Erdős-Rényi Blocks
- Second Phase: Using Chung-Lu Model to connect blocks with $p_{ij} = p(d_i, d_j)$

Model Verification: BTER

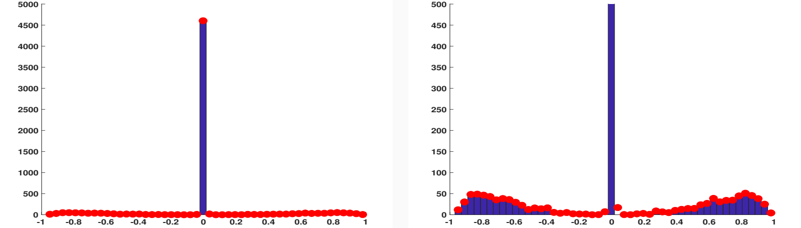

Figure 1: Erdos collaboration network.

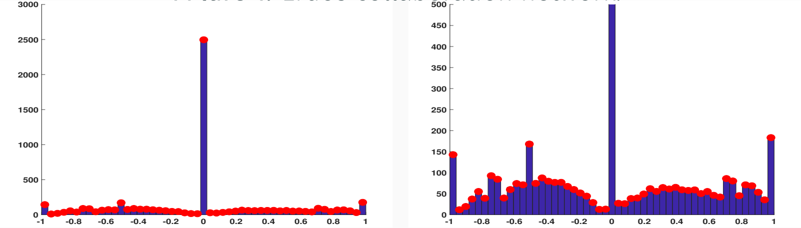

Figure 2: BTER model for Erdos collaboration network.

What Do You Hear?

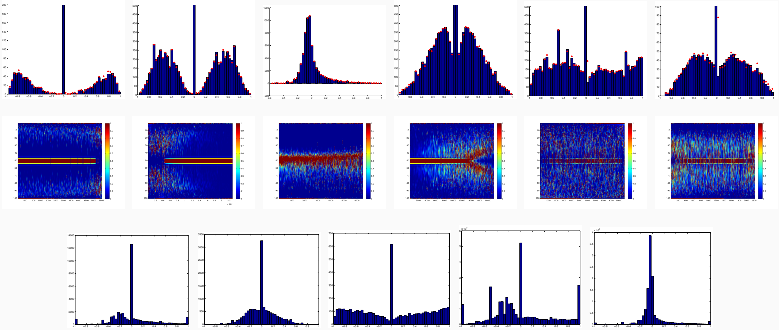

Latest:

- Dong, Benson, Bindel (KDD 2019).
- Longer talk at ILAS 2019 (slides online)