Spectral Graph Theory

*D. J. Kelleher[1]

[1]Department of Mathematics
University of Connecticut

UConn— SIGMA Seminar — Fall 2011

Basic graph theory stuff

Formally, a graph is a pair $G = (V, E)$, where

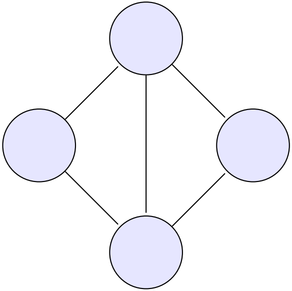

Basic graph theory stuff

Formally, a graph is a pair $G = (V, E)$, where

V is the vertex set.

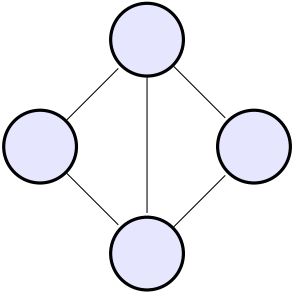

Formally, a graph is a pair $G = (V, E)$, where

V is the vertex set.

$E \subset V \times V$ is the edge set.

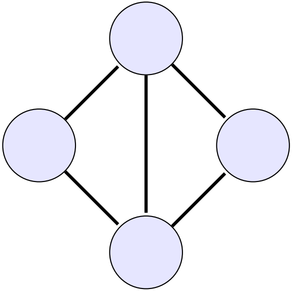

Basic graph theory stuff

Formally, a graph is a pair $G = (V, E)$, where

V is the vertex set.

$E \subset V \times V$ is the edge set.

We say that $x \sim y$ if $(x, y) \in E$.

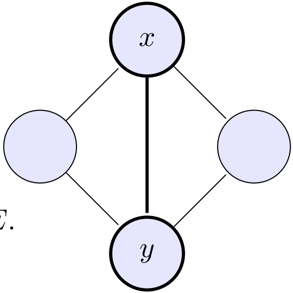

Basic graph theory stuff

Formally, a graph is a pair $G = (V, E)$, where

V is the vertex set.

$E \subset V \times V$ is the edge set.

We say that $x \sim y$ if $(x, y) \in E$.

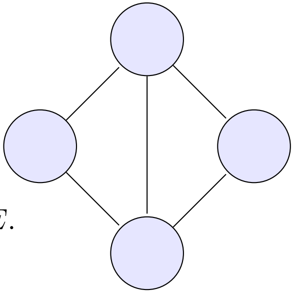

We could also add edge weights, directions to the edges, and there are generalizations of most of what follows. However, we will assume that all graphs are simple, i.e. E is symmetric.

Encoding a graph as a matrix

We want to start using matrices, so we take functions $f : V \to \mathbb{R}$, if we agree to an enumeration of the vertex set, this allows us to write these functions as vectors

$$\begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \end{array} \begin{pmatrix} f_1 \\ f_2 \\ f_3 \\ f_4 \end{pmatrix}$$

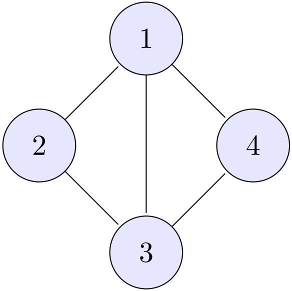

This will allow us to think of operators on these functions as matrices.

Encoding a graph as a matrix

The degree matrix D, where $D_{ij} = 0$ if $i \neq j$ and $D_{jj} = d_j$ is the degree of the jth vertex — the number of edges that j is on.

$$\begin{array}{c c} & \begin{array}{cccc} 1 & 2 & 3 & 4 \end{array} \\ \begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \end{array} & \begin{pmatrix} 3 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix} \end{array}$$

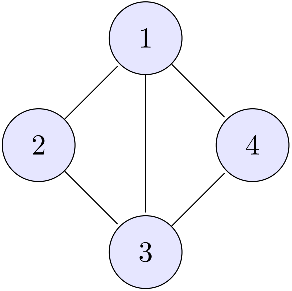

The degree matrix is a diagonal matrix.

Encoding a graph as a matrix

The adjacency matrix A where $A_{ij} = 1$ ith and jth verteces are connected.

$$\begin{array}{c c} & \begin{array}{cccc} 1 & 2 & 3 & 4 \end{array} \\ \begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \end{array} & \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{pmatrix} \end{array}$$

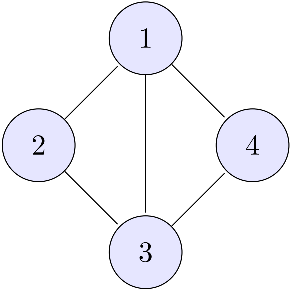

The adjacency matrix is symmetric and often sparse in practice.

Encoding a graph as a matrix

The (non-normalized) Laplacian matrix is $\Delta = D - A$.

$$\begin{array}{c c} & \begin{array}{cccc} 1 & 2 & 3 & 4 \end{array} \\ \begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \end{array} & \begin{pmatrix} 3 & -1 & -1 & -1 \\ -1 & 2 & -1 & 0 \\ -1 & -1 & 3 & -1 \\ -1 & 0 & -1 & 2 \end{pmatrix} \end{array}$$

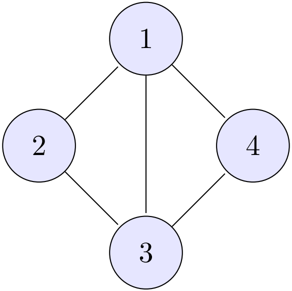

Some literature refers to this as the negative of the Laplacian, in our case, we take the negative because it is non-negative definite, In particular all eigenvalues are between 0 and ∞ (with zero included).

Encoding a graph as a matrix

To normalized Laplacian $\mathcal{L} = D^{-1/2}\Delta D^{-1/2}$, so the previous matrix we were using becomes

$$\begin{array}{c c} & \begin{array}{cccc} 1 & 2 & 3 & 4 \end{array} \\ \begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \end{array} & \begin{pmatrix} 1 & -1/\sqrt{6} & -1/3 & -1/\sqrt{6} \\ -1/\sqrt{6} & 1 & -1/\sqrt{6} & 0 \\ -1/3 & -1/\sqrt{6} & 1 & -1/\sqrt{6} \\ -1/\sqrt{6} & 0 & -1/\sqrt{6} & 1 \end{pmatrix} \end{array}$$

$$\text{i.e. } \mathcal{L}_{xy} = \begin{cases} 1 & \text{if } x = y, \\ -\frac{1}{\sqrt{d_x d_y}} & x \sim y \\ 0 & \text{otherwise.} \end{cases}$$

Encoding a graph as a matrix

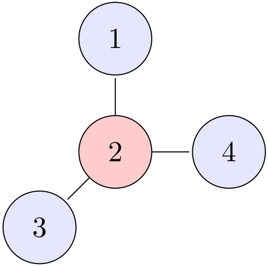

We can also look at delta as an linear operator acting on functions $f : V \to \mathbb{R}$, given by

$$\Delta f(x) = \sum_{x \sim y} (f(x) - f(y))$$

or

$$\mathcal{L} f(x) = \frac{1}{\sqrt{d_x}} \sum_{x \sim y} \left(\frac{f(x)}{\sqrt{d_x}} - \frac{f(y)}{\sqrt{d_y}} \right)$$

Encoding a graph as a matrix

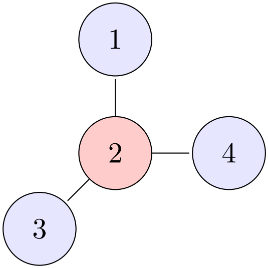

The matrix $D^{-1/2}\mathcal{L}D^{1/2} = D^{-1}\Delta$ is conjugate to $\mathcal{L}$, and thinking of it as an operator is slightly nicer

$$d_x^{-1}\sum_{x\sim y}(f(x) - f(y))$$

Spectral Theorem

The spectrum of a matrix is the set of eigenvalues, for the this talk I will refer to the spectrum of a graph as the spectrum of the Laplacian

$$\mathcal{L}f = \lambda f$$

λ is an eigenvalue, f is an eigenfunction. The eigenspace of λ is the set of eigenfunctions which satisfy the above equations. The λ-eigenspace is a linear space. Note that because Δ and $\mathcal{L}$ are non-negative definite, we have a full set of non-negative real eigenvalues.

The 0-eigenspace is the set of (globally) harmonic function.

Spectral Theorem

The spectrum of a matrix is the set of eigenvalues, for the this talk I will refer to the spectrum of a graph as the spectrum of the Laplacian

$$\mathcal{L}f = \lambda f$$

λ is an eigenvalue, f is an eigenfunction. The eigenspace of λ is the set of eigenfunctions which satisfy the above equations. The λ-eigenspace is a linear space. Note that because Δ and $\mathcal{L}$ are non-negative definite, we have a full set of non-negative real eigenvalues.

The 0-eigenspace is the set of (globally) harmonic function.

Spectral Theorem

The spectrum of a matrix is the set of eigenvalues, for the this talk I will refer to the spectrum of a graph as the spectrum of the Laplacian

$$\mathcal{L}f = \lambda f$$

λ is an eigenvalue, f is an eigenfunction. The eigenspace of λ is the set of eigenfunctions which satisfy the above equations. The λ-eigenspace is a linear space. Note that because Δ and $\mathcal{L}$ are non-negative definite, we have a full set of non-negative real eigenvalues.

The 0-eigenspace is the set of (globally) harmonic function.

Spectral Theorem

Spectral Theorem *If A is a real symmetric $n \times n$-matrix, then each eigenvalue is real, and there is an orthonormal basis of $\mathbb{R}^n$ of eigenfunctions (eigenvectors) of A.*

$\{e_j\}_{j=1}^n$ is orthonormal if $e_j \cdot e_k = \delta_{jk} = \begin{cases} 0 & \text{if } j \neq k \\ 1 & \text{if } j = k. \end{cases}$

Further, *If A is non-negative definite, that is*

$$(Af) \cdot f = f^T A f \geq 0, \quad \forall f \in \mathbb{R}^n$$

then all of the eigenvalues are non-negative.

Δ and $\mathcal{L}$ are both symmetric and non-negative definite.

Spectral Theorem *If A is a real symmetric $n \times n$-matrix, then each eigenvalue is real, and there is an orthonormal basis of $\mathbb{R}^n$ of eigenfunctions (eigenvectors) of A.*

$\{e_j\}_{j=1}^n$ is orthonormal if $e_j \cdot e_k = \delta_{jk} = \begin{cases} 0 & \text{if } j \neq k \\ 1 & \text{if } j = k. \end{cases}$

Further, *If A is non-negative definite, that is*

$$(Af) \cdot f = f^T A f \geq 0, \quad \forall f \in \mathbb{R}^n$$

then all of the eigenvalues are non-negative.

Δ and $\mathcal{L}$ are both symmetric and non-negative definite.

What the spectrum tells us

Question What kind of functions are harmonic?

Looking at

$$d_x^{-1} \sum_{x \sim y} (f(x) - f(y)) = 0$$

we quickly realize constants are harmonic (as in $\mathbb{R}^n$).

... well, locally constant functions at least. i.e. $f(x) = f(y)$ for all $x \sim y$.

Question What kind of functions are harmonic?
Looking at

$$d_x^{-1} \sum_{x \sim y} (f(x) - f(y)) = 0$$

we quickly realize constants are harmonic (as in $\mathbb{R}^n$).

... well, locally constant functions at least. i.e. $f(x) = f(y)$ for all $x \sim y$.

Question What kind of functions are harmonic?
Looking at

$$d_x^{-1} \sum_{x \sim y} (f(x) - f(y)) = 0$$

we quickly realize constants are harmonic (as in $\mathbb{R}^n$).

... well, locally constant functions at least. i.e. $f(x) = f(y)$ for all $x \sim y$.

What the spectrum tells us: Connected components

Locally constant functions are functions which are constant on the connected components of our graph G.

If we re-enumerate, then the Laplacian is a block diagonal matrix:

$$\mathcal{L} = \begin{pmatrix} \mathcal{L}_1 & 0 & \dots & 0 \\ 0 & \mathcal{L}_2 & \dots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \dots & \mathcal{L}_k \end{pmatrix}$$

where $\mathcal{L}_i$ is the Laplacian matrix of the connected components G_i of G.

In particular, the number of connected components of a graph dimension of the 0-eigenvalue (multiplicity of 0 as an eigenvalue).

What the spectrum tells us: Bipartite graphs

We say that $G = (V, E)$ is bipartite if $V = V_1 \cup V_2$ (disjoint) such that $x \sim y$ only if $x \in V_1$ and $y \in V_2$ or visa versa. Assume G is bipartite.

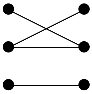

Then 2 is an eigenvalue of the normalized Laplacian. The converse holds true as well.

A graph is a tree if it contains no cycles. A subgraph of G is called a spanning tree if it is a tree that contains all the verteces of G.

These are the 8 spanning trees of the graphs we showed earlier.

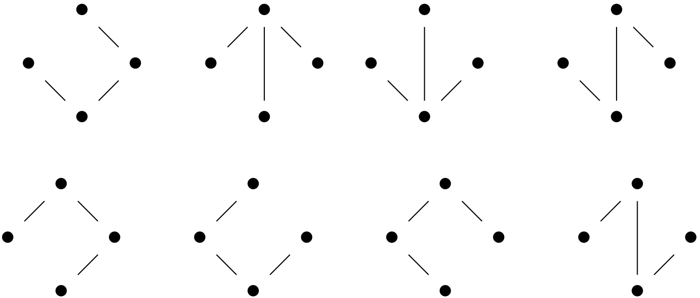

What the spectrum tells us: Spanning trees

Kirchhoff's theorem If G is a connected simple graph, and Δ is the non-normalized Laplacian G, and $\lambda_1, \ldots, \lambda_{n-1}$ are the non-zero eigenvalues of Δ, then the number of spanning trees of G is given by

$$\frac{1}{n}\lambda_1 \cdots \lambda_{n-1}.$$

Fractals

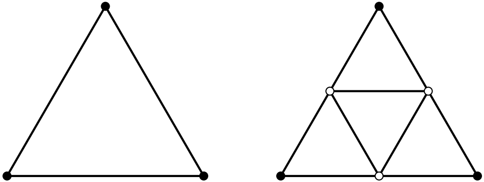

To get a Laplacian on a finitely ramified fractal, we take a set of approximating graphs G_m, take their non-normalized Laplacians Δ_m, then we take a scaled limit, so in the case of Siepriński gasket

$$\Delta f(x) = \frac{3}{2} \lim_{m\to\infty} 5^m \Delta_m f(x).$$

Note that points in the graph are being identified with points in the fractal.

Fractals

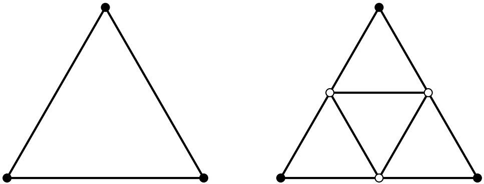

So graph approximations allow us to approximate the Laplacian on the limiting space, this gives us valuable insights into things like

- Spectral dimension
- Random walks
- Heat kernels

And in special cases like with the Sierpinski gasket, this formulation allows us to say things about the spectrum and eigenfunctions of the limit Laplacian.

The hexacarpet

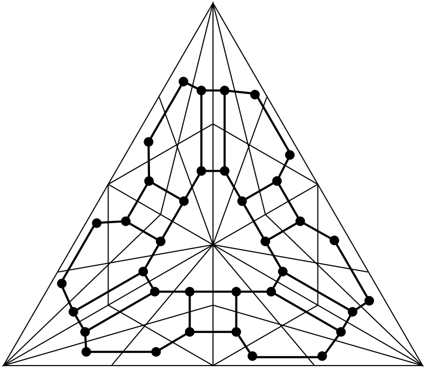

Last semester, Sasha, Hugo, Ryan, Aaron, Matt and D started looking at a fractal that comes from barycentric subdivisions of

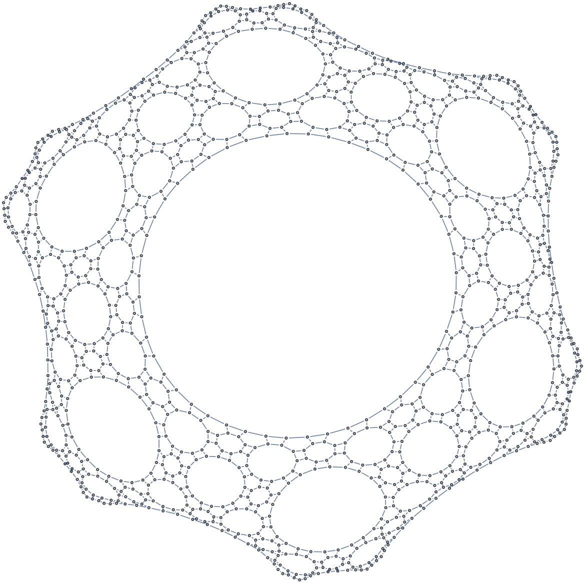

This is what the 4th level approximation looks like. (thanks to Hugo Panzo)

Eigenfunction coordinates

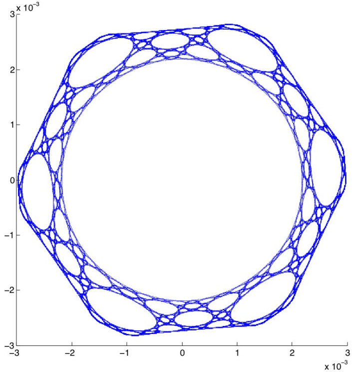

If we take the two (preferably linearly independent) eigenfunctions ϕ_1 and ϕ_2 of your Laplacian, and graph the set of points $(\phi_1(a), \phi_2(a))$ for $a \in V$ in euclidean space, you get something like the above picture. (thanks to Matt Begue)

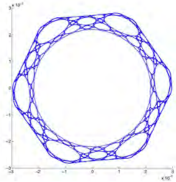

(A) (φ_2, φ_3)

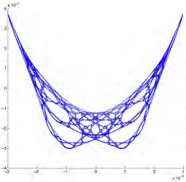

(B) (φ_2, φ_4)

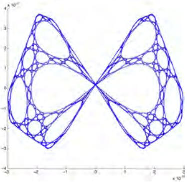

(C) (φ_2, φ_5)

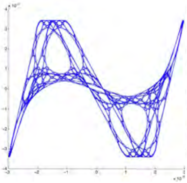

(D) (φ_2, φ_6)

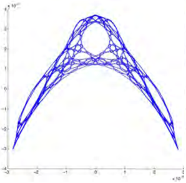

(E) (φ_3, φ_4)

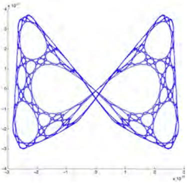

(F) (φ_3, φ_5)

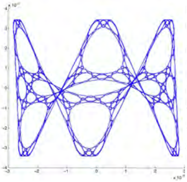

(G) (φ_3, φ_6)

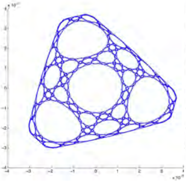

(H) (φ_4, φ_5)

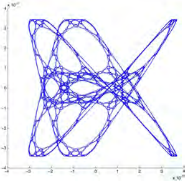

(I) (φ_4, φ_6)

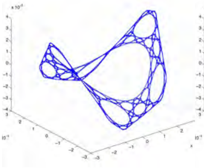

(A) $(\varphi_2, \varphi_3, \varphi_4)$

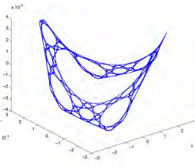

(B) $(\varphi_2, \varphi_3, \varphi_5)$

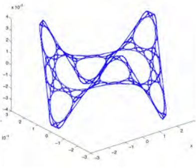

(C) $(\varphi_2, \varphi_3, \varphi_6)$

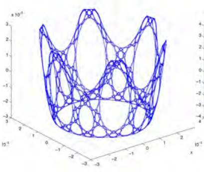

(D) $(\varphi_2, \varphi_3, \varphi_7)$

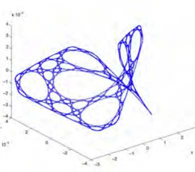

(E) $(\varphi_2, \varphi_4, \varphi_5)$

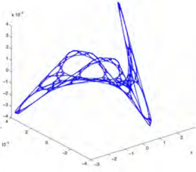

(F) $(\varphi_2, \varphi_4, \varphi_6)$

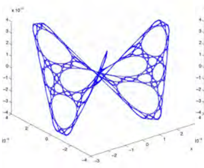

(G) $(\varphi_2, \varphi_5, \varphi_6)$

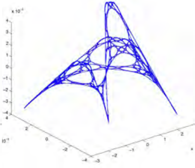

(H) $(\varphi_3, \varphi_5, \varphi_6)$

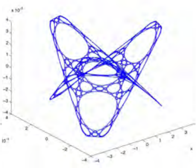

(I) $(\varphi_4, \varphi_5, \varphi_6)$

Eigenfunction coordinates

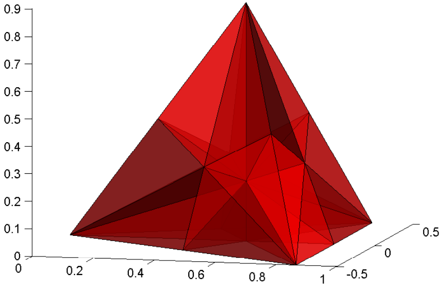

Over the summer the REU students Diwakar Raisingh, Gabriel Khan, as well as Matt Beque and DK started to consider the 3-simplex version of this fractal.

Eigenfunction coordinates

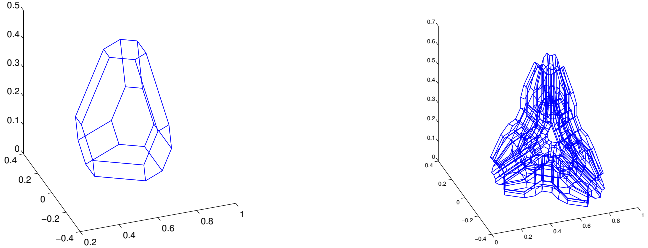

Figure: First and second level graph approximations to the 3-barycentric sponge.

Eigenfunction Video

Eigenfunction Pictures

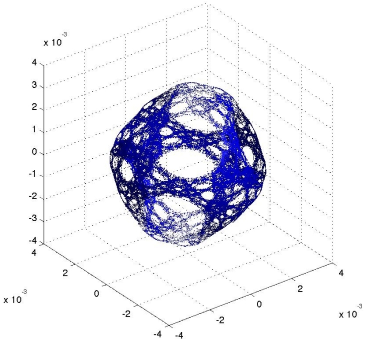

Harmonic Extension

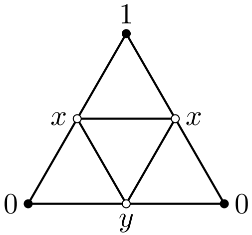

Harmonic extension / Dirichlet Problem Say we know a function f on $V_0 \subset V$ (Filled dots above), we want to extend to a function $\tilde{f}$ on V with

$$\Delta f(p) = 0, \quad \forall \ \ p \notin V_0.$$

We call V_0 the boundary, and $\tilde{f}$ a harmonic Extension of f.

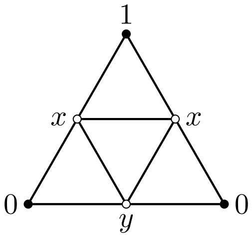

Let's try!

$$0 = 2(y - 0) + 2(y - x) = 4y - 2x$$

So

$$y = x/2$$

Harmonic Extension

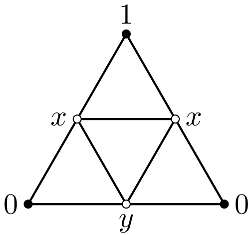

Let's try!

$$0 = 2(y-0) + 2(y-x) = 4y - 2x$$

So

$$y = x/2$$

Harmonic Extension

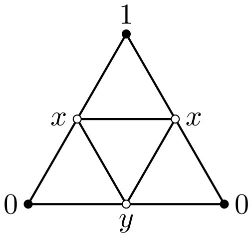

Let's try!

$$0 = 2(y - 0) + 2(y - x) = 4y - 2x$$

So

$$y = x/2$$

Harmonic Extension

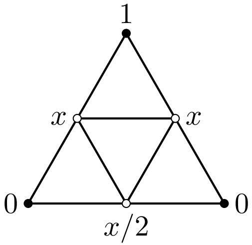

Let's try!

$$0 = (x - 0) + (x - x/2) + (x - x) + (x - 1) = \frac{5}{2}x - 1$$

So

$$x = \frac{2}{5}, \quad \text{and} \quad y = \frac{1}{5}$$

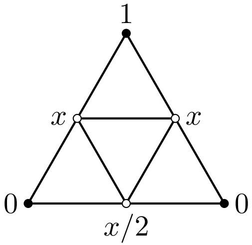

Let's try!

$$0 = (x - 0) + (x - x/2) + (x - x) + (x - 1) = \frac{5}{2}x - 1$$

So

$$x = \frac{2}{5}, \quad \text{and} \quad y = \frac{1}{5}$$

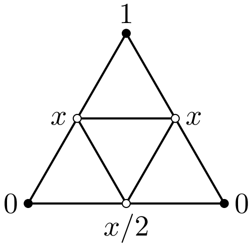

Let's try!

$$0 = (x - 0) + (x - x/2) + (x - x) + (x - 1) = \frac{5}{2}x - 1$$

So

$$x = \frac{2}{5}, \quad \text{and} \quad y = \frac{1}{5}$$

Harmonic Extension

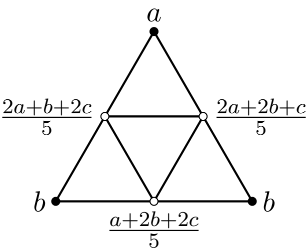

We've done more Harmonic extension is linear, and applying symmetries of the graph, we can compute the extension of any function.

Even more than that, if we iterate this process, we can compute harmonic extensions of all approximating graphs to the Sierpinski Gasket, and then extend to the entire space.

Harmonic Extension

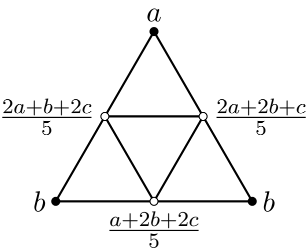

We've done more Harmonic extension is linear, and applying symmetries of the graph, we can compute the extension of any function.

Even more than that, if we iterate this process, we can compute harmonic extensions of all approximating graphs to the Sierpinski Gasket, and then extend to the entire space.

Harmonic Coordinates

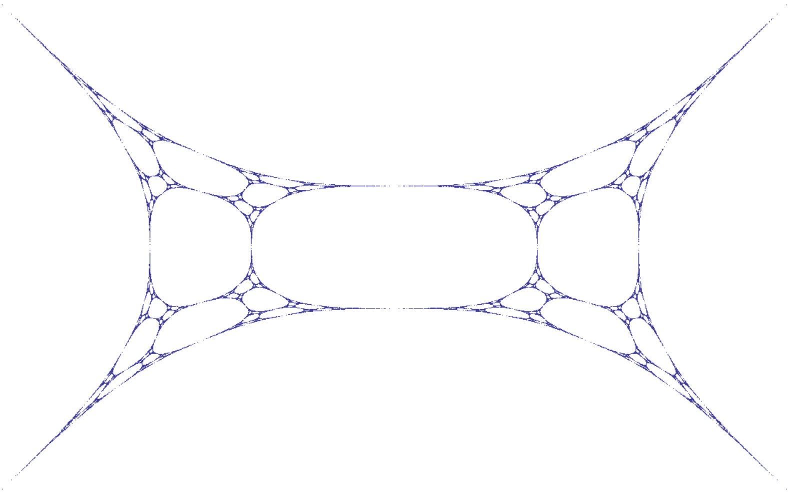

Similarly to the eigenvalue coordinates, take two harmonic functions with specific boundary values, f_1 and f_2 and plot the points $\{(\tilde{f}_1(p), \tilde{f}_2(p)) \mid p \in V_n\}$ in $\mathbb{R}^2$. (Picture by Jason Marsh)

What about science?

A worm-brain idea

Following the cue of a scientist Dmitri Chklovskii, part of the REU summer 2011 was also to examine the brain of *Caenorhabditis elegans* (C. Elegans).

Chklovskii formed a Laplacian matrix of the graph representing 279-neuron brain. Naturally, Tyler Reese, Dylan Yott, Antoni Brzoska, and DK asked the question... "is it a fractal."

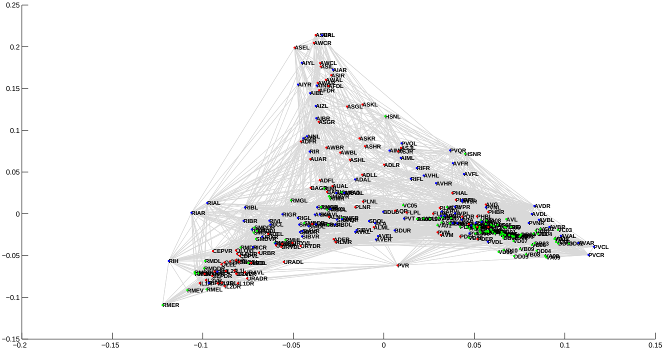

C. Elegans eigenfunction coordinate representation, thanks to Dylan, Tyler, Toni

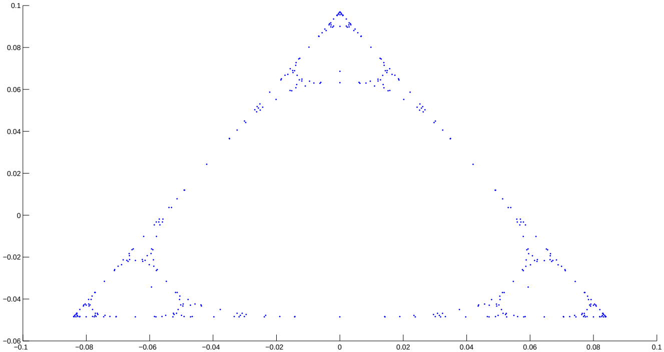

It's doesn't look very fractal.

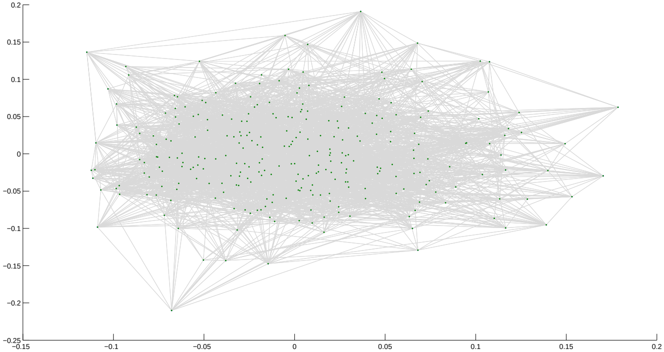

But it doesn't look random either.

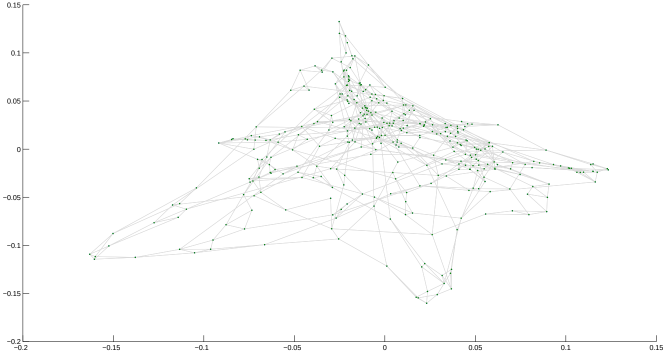

Maybe more like this guy?

We also compared the spectrum of this graph looking at

- Spectral Gaps
- Eigenvalue counting functions (Weyl ratios)
- Small world properties
- Localization of eigenfunctions

Comparing this to

- random graphs
- random trees
- fractals
- rewired fractals

None of which proved to be that satisfying of a model.

Places to find out more:

- Fan Chung's website: math.ucsd.edu/∼fan
- Chklovskii: neurop-tikon.org/projects/display/chklovskiilab/Publications
- DK's Page: math.uconn.edu/∼kelleher

Thanks!

- Hexacarpet: Matt Begue, Hugo Panzo, Aaron Nelson, Ryan Pellico
- UConn Fractals REU 2011
- Barycentric sponge: Gabriel Khan and Diwakar Raisingh
- Worm-Brain Stuff: Dylan Yott, Tyler Reese and Antoni Brzoska
- Pillow harmonic coordinates: Jason Marsh
- Data and code: Dmitri Chklovskii
- Research supported in part by NSF grant DMS-0505622.
- Special thanks: Sasha Teplyaev