

Graph-Theoretic Methods

Notes are derived from GTM1819; main addition is a section on persistent homologies.

Motivation and Introduction

One is often faced with analyzing large spatial or spatiotemporal datasets – say involving N nodes, or N time series. If one is only interested in the individual behavior of each node or time series, analysis remains tractable (or at least, as order N) – a “massively univariate” approach (for spatial datasets). However, this approach (a) ignores any possible interactions between the nodes or time series, and, implicitly, makes the assumption that the N nodes or time series directly represent the underlying variables of interest -- the more likely alternative being that the observables each represent a mixture of the underlying variables.

A natural approach is therefore to look at covariances (in the case of spatial datasets) or, in the case of time series, cross-spectra – the quantities $P_{x_i, x_j}(\omega)$ defined earlier. While this is comprehensive (at least for pairwise interactions), there’s a scaling problem – the number of covariances (or cross-spectra) is $N(N - 1) / 2$. So even if N is relatively small (say, 25 to 50), N^2 can be sufficiently large so as to be difficult to visualize, and when N is large (e.g., imaging pixels), a comprehensive approach may be computationally impractical. So this motivates the development of methods that summarize the collective properties of a set of covariances, or a matrix of cross-spectra.

One kind of approach is to recognize that these covariances do form a matrix, and, as such, the standard ways of describing a matrix are applicable. A good example is the “global coherence” – this is the ratio of the largest eigenvalue of the matrix of cross-spectra to its trace, i.e., a measure of the extent to which the cross-spectral matrix can be approximated by a matrix of rank 1 (see Homework 3 for LSBB, 2012-2013). A rank-1 cross-spectral matrix implies that the observed covariances can all be accounted for by a single common source of noise, so the global coherence has an immediate interpretation.

Graph-theoretic methods represent another kind of approach. Here, the basic idea is also one of dimensional reduction, but the dimensional reduction occurs at an earlier stage: reducing the original pairwise measurements, which may be real or complex numbers, to something simpler – typically a matrix A whose elements a_{ij} are just 0’s and 1’s. Then, one looks at the properties of this array.

Overall, the advantage of this kind of approach is that it focuses on the qualitative nature of the interrelationships of the observed variables, and their network-like (or “topological”) attributes. The main cautionary note is that the reduction from covariances or cross-spectra to binary quantities involves thresholding (and, in the case of cross-spectra, the choice of analysis frequency.) So to check on the robustness of the results, one might want to investigate how they depend on this choice – even if the choice is “objective” (as in, a level of statistical significance).

One should also note that graph-theoretic methods have in general been developed for application to domains in which this thresholding is not an issue – i.e., domains in which the fundamental measurements are binary. A good example is the case of social networks: a matrix element $a_{ij} = 0$ means that individuals i and j are strangers, while $a_{ij} = 1$ means that they know each other.

Note also that in this case, the matrix A is necessarily symmetric. This symmetry typically is taken to be part of the definition of a graph. But many graph-theoretic measures are applicable far beyond the restricted domain of symmetric, binary graphs. Asymmetric but binary arrays correspond to “directed” graphs (for example, i has contacted j , or website i links to website j). Arrays that are not binary correspond to “weighted” graphs (some links are stronger than others). Depending on the data type, these extensions may be particularly natural. For example, diffusion tractography data provides a measure of the strength of connectivity, but not its directionality. In all of these cases (and in everything we consider), $a_{ij} \geq 0$. Relaxing this condition (i.e., allowing negative “strengths”) changes the character of the problem greatly – for example, the graph Laplacian is no longer positive definite, so that “diffusion” can diverge. Taking this even further – allowing a_{ij} to be complex, with $a_{ji} = \overline{a_{ij}}$, essentially yields the picture of general linear dynamics at a single frequency, as the connectivities now can be considered to represent values of a transfer function.

There is also a connection to Markov chains. If every column of a graph matrix sums to 1 (the “Markov” condition), we can view a_{ij} as the probability that a particle at node j will, at the next time step, move to node i . Note also that we can take any graph matrix and column-normalize it (i.e., replace a_{ij} by $a_{ij} / \sum_k a_{kj}$) to ensure that its columns sum to 1, so that we can always turn a graph matrix into a Markov matrix.

Elements

A web-based source for this material may be found in Radu Horaud, <http://csustan.csustan.edu/~tom/Clustering/GraphLaplacian-tutorial.pdf>.

General definitions

A *graph* G consists of a set of vertices V (a.k.a. nodes) and an adjacency matrix A whose rows and columns correspond to the vertices.

The vertices are typically labeled by the integers $\{1, \dots, N\}$ -- but typically, these are abstract labels, and their numerical values are irrelevant. All elements of the adjacency matrix A are assumed to be non-negative.

A *subgraph* of G is a subset of the vertices, along with the matrix formed from the corresponding rows and columns of the adjacency matrix A .

If $a_{ij} > 0$, we say that there is an edge connecting vertex i to vertex j , and write $i \sim j$.

The *degree* of a vertex is the number of vertices that it connects with. In the case of a directed graph, one needs to distinguish between the outgoing degree and the incoming degree.

The *distance* between two vertices is the minimum number of edges that must be traversed to pass between them. Each example of a minimum-length path is called a *geodesic*. The *diameter* of a graph is the length of its largest geodesic. It is the lowest power p for which A^p has no zero off-diagonal elements.

A *clique* is a set of vertices for which every pair is connected by an edge. Two vertices connected by an edge trivially form a clique. A three-clique is a set of vertices whose subgraph is a triangle.

On this general setup, we typically add one or more kinds of restrictions:

If we allow nonzero a_{ij} 's to be quantities other than 1, the graph is a “weighted” graph.


If we allow the matrix A to be asymmetric, the graph is a “directed” graph; a_{ij} is the strength of the connection from j to i (this convention makes the connection with Markov matrices most straightforward).

Typically, we will deal with “simple” graphs; three conditions apply: all a_{ij} are 0 or 1, they are symmetric, and the on-diagonal elements are 0. The last condition means that there are no self-loops.

Going from a simple graph to a weighted graph is often a straightforward generalization, in that a graph with integer weights can be thought of as a graph in which there are multiple connections between a single pair of nodes. So one can expect that this generalization will not affect approaches that depend on the algebraic structure of A (though it may affect approaches that depend on counting the number of paths).

In contrast, going from a simple graph to a directed graph means allowing A to be asymmetric. This will change its algebraic properties markedly, since A goes from being self-adjoint to non-self-adjoint. However, combinatorial approaches (e.g, approaches that rely on counting paths between nodes) may generalize readily. For example, each element of A^k counts the number of k -step paths from node j to node i .

A simple graph is just a set of points and lines that connect them:



For this graph, the adjacency matrix is given by $A = \begin{pmatrix} 0 & 1 & 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$.

The cliques consist of the seven pairs of vertices connected by edges, and also the 3-clique $\{1,2,3\}$.

Some special kinds of graphs

The following are fairly common, and often useful:

A *connected* graph is a graph for which, for every pair of vertices, there is a sequence of edges that connects them. Equivalently, it is a graph for which a sufficiently high power of the adjacency matrix A contains no zero elements. Since the elements of A and its powers are non-negative, this is equivalent to the condition that $1 + zA + z^2A^2 + \dots = \frac{1}{1 - zA}$ (series converges for sufficiently small z) has no zero elements. Equivalently, a connected graph is a graph whose diameter is finite.

A *cyclic* graph is a connected graph for which every vertex has degree 2.

A *forest* is a graph that has no subgraphs that are cycles.

A *tree* is a forest that is connected.

A *triangle-free* graph is a graph that has no subgraphs that are 3-cliques. All forests are triangle-free.

A *regular* graph is a graph for which all vertices have the same degree.

A *complete* graph is a graph for which all edges are present. Equivalently, the entire graph is a clique.) Equivalently, it is a graph of diameter 1.

A *star* graph is a graph in which one vertex is connected to all of the others, and there are no cycles. It has diameter 2.

A *bipartite* graph is a graph whose vertices can be partitioned into two disjoint subsets, and the vertices within each subset are not connected with each other. (Similarly for a *k-partite* graph.)

Some special kinds of graph architectures

The above terms apply to specific graphs, and they are fragile – the properties can be gained or lost by insertion or deletion of a single vertex or edge. The terms below are different – they are intended to apply to graph architectures (strategies of building graphs with an arbitrarily large number of nodes, as the number of nodes is large).

A *scale-free* graph is a graph for which the frequency of nodes with a given degree is a power-law function of the degree. That is, the probability that a node has degree between d and $d + \Delta d$ is given by $p(d) = d^{-\gamma} \Delta d$; typically γ is in the range of 2 to 3. Obviously this can only hold over a range of degrees, and must be approximate. But the basic idea is that there are a small number of vertices that have a large number of connections (“hubs”), a larger number of nodes that have somewhat fewer connections, etc.

A *small-world* graph (or network) is a graph in which the average distance between pairs of vertices is small in comparison to the total number of vertices. A complete graph is small-world (with N vertices, the ratio is $1/N$); so is a star graph (ratio is $1/N$). A cyclic graph is not small-world (with N vertices, the ratio is $N/2$). Typically, the term is reserved for situation that the average distance grows no faster than the logarithm of the number of vertices.

Scale-free graphs, with exponents γ in the range of 2 to 3, are small-world. But not all small-world graphs are scale-free. A classic example of this is the small-world architecture of Watts and Strogatz (Nature 1998). These graphs are made by (i) starting with a cyclic graph, and (ii) adding a small number of long-range cross-connections.

An Erdos-Renyi graph is a graph in which the probability of a connection between any two nodes is set to some value p . For $p = 1$, this yields a complete graph. For $p < 1$, the graph may not be connected. There are numerous results on the behavior of these graphs as the number of vertices grows – for example, the probability that the graph is connected, the probability that it has a single large component, etc. For large N , $\frac{\ln N}{N}$ determines connectedness: for $p > \frac{\ln N}{N}$, the graph is almost always connected, for $p < \frac{\ln N}{N}$, the graph is almost always not connected.

The graph Laplacian

The graph Laplacian is a matrix that captures many aspects of a graph, including how dynamical processes on the graph evolve, and how the graph can best be represented by a one-dimensional structure. It is also central to many algorithms (“spectral partitioning”, “spectral clustering”) that use graphs for many other purposes (image segmentation, VLSI layout).

We will define the Laplacian for a simple (unweighted, un-directed) graph. The extension to a weighted graph is straightforward: instead of a 1 or a 0 to indicate whether an edge is present, one can insert a positive real number to indicate its strength. (One can consider the intermediate case – the possibility of a multiplicity of equally-weighted edges between each pair of points – and this corresponds exactly to the case of integer weights.) However, the extension of the graph Laplacian to a directed graph is not, as the symmetry of the in-connections and out-connections from a node is an essential feature.

The starting point is to consider a function on the graph, i.e., an assignment x_i of values to each vertex on the graph. We can think of each value x_i to represent an amount of material (or heat) at the vertex i . Now allow this material to diffuse, according to the graph's connectivity. The amount at each vertex at time t is given by $x_i(t)$. At each time step, some of the material goes to each of the d_i vertices j that is connected to i , but also, some of the material at these neighboring vertices flows back to i . So,

$$\frac{d}{dt}x_i = K \left(-d_i x_i + \sum_{j \sim i} x_j \right),$$

where the first term on the right represents the loss to neighboring vertices, and the second term represents the influx. The constant K is irrelevant (it provides an overall temporal scale, and can be set to 1), so we can write

$$\frac{d}{dt}\vec{x} = -L\vec{x},$$

where \vec{x} is the column matrix of the x_i 's and L is the *graph Laplacian*. The graph Laplacian is given by $L = D - A$, the difference of the diagonal matrix of the degrees (D) and the adjacency matrix. Note that L is self-adjoint, as D and A are both real and symmetric.

A justification for calling L a Laplacian is that the above resembles the standard heat (or diffusion) equation: for $x(s,t)$ is the temperature of a conducting bar at position s and time t , the evolution is $\frac{\partial}{\partial t}x = K \frac{\partial^2}{\partial s^2}x$. The graph Laplacian, like the ordinary Laplacian, is proportional to the difference between the value at a point, and its local average. There is a sign difference – merely a matter of convention -- the graph Laplacian approximates the *negative* of the second derivative, while the continuous-system Laplacian, $\frac{\partial^2}{\partial s^2}$ or ∇^2 , is the (positive) second derivative.

We now would like to solve the equation $\frac{d}{dt}\vec{x} = -L\vec{x}$, given some set of initial conditions $\vec{x}(0)$.

Since L is self-adjoint, it has a complete set of eigenvectors $\vec{\varphi}_m$ (i.e, a set that spans the vector space of functions on the graph). Each of these has an associated eigenvalue λ_m , with $L\vec{\varphi}_m = \lambda_m\vec{\varphi}_m$. Each eigenvector $\vec{\varphi}_m$ corresponds to a solution that dies down exponentially, with rate λ_m :

$$\frac{d}{dt} \vec{\varphi}_m e^{-\lambda_m t} = -\lambda_m \vec{\varphi}_m e^{-\lambda_m t} = -L \vec{\varphi}_m e^{-\lambda_m t}.$$

So to find a solution that begins with $\vec{x}(0)$ at time zero, we write $\vec{x}(0)$ in the eigenvector basis

set, $\vec{x}(0) = \sum_{m=1}^N c_m \vec{\varphi}_m$. This yields a solution:

$$\vec{x}(t) = \sum_{m=1}^N c_m \vec{\varphi}_m e^{-\lambda_m t}.$$

The bottom line is that the eigenfunctions of the Laplacian indicate the patterns that die out exponentially, and the eigenfunctions with lowest nonzero eigenvalues die out the slowest – and thus, characterize the dynamics of functions that evolve on G .

There are some useful variants of the Laplacian that differ in row- or column-normalization.

The “normalized” Laplacian is defined as $L_{norm} = D^{-1/2} L D^{-1/2}$. Since

$L_{norm} = D^{-1/2} L D^{-1/2} = D^{-1/2} (D - A) D^{-1/2} = I - D^{-1/2} A D^{-1/2}$, it is equal to the identity on its diagonal. For a weighted graph with “weak” connections, $L_{norm} = I - \varepsilon D^{-1/2} A D^{-1/2}$. Since this is close to the identity, one can approximate repeated applications of L_{norm} by

$(L_{norm})^r = (I - \varepsilon D^{-1/2} A D^{-1/2})^r \approx \left(\exp^{-\varepsilon D^{-1/2} A D^{-1/2}} \right)^r = \exp^{-\varepsilon r D^{-1/2} A D^{-1/2}}$. L_{norm} is symmetric, and its eigenvalues can also be shown to be non-negative (see below).

The “transition matrix” Laplacian allows for a connection with Markov chains:

$L_{transition} = D^{-1} L = I - D^{-1} A$. $L_{transition}$ differs from the standard graph Laplacian in that its rows are divided by the degree. This means that the matrix now represents particle diffusion in the following sense: a particle can choose to leave a node, and if it does so, it chooses any of the edges randomly to go to the next node. So it’s diffusion of probabilities, not heat (as was the case for the Laplacian as initially defined).

To see the difference between this and the graph Laplacian considered above, consider a “star” graph with $N + 1$ vertices. For the standard graph Laplacian, the function that is constant on all nodes is an eigenvector of eigenvalue 0. For the transition-matrix Laplacian, the function that has value N on the central vertex and value 1 on the peripheral ones is an eigenvector: on every step, a particle at the central vertex moves to any one of the peripheral ones, but a particle at any peripheral node must move centrally.

Eigenstructure of the Laplacian: null space, positive-definiteness

We first determine the eigenvectors corresponding to a zero eigenvalue via elementary means, then show that the other eigenvectors are positive, and then use some group theory to determine them in some interesting special cases. We focus on the standard graph Laplacian, but all the arguments work for the other variants.

The null space

A function on the graph is mapped to zero by the Laplacian if its value at every vertex is equal to the average of its neighbors. So a function that is constant on the graph is mapped to zero. Conversely, if a function is not equal to the average value of its nearest neighbors, then it is not mapped to zero.

If a graph is not connected, the Laplacian consists of blocks, one for each connected component. This is because D is diagonal, and A is in blocks. A function that is constant on each component separately (but perhaps has a different value on each component) is therefore mapped to zero by the Laplacian.

So the null space of the Laplacian is the set of functions that are constant on each component separately.

Positive-definiteness, the incidence matrix, and representing the graph on a line

The above comments determine the eigenvectors of eigenvalue 0 for the graph Laplacian. With a bit more work, we can show that all of the eigenvalues of the graph Laplacian are non-negative. We do this by representing it as a sum of outer products, and this leads to some other applications and interpretations of the graph Laplacian.

Positive-definiteness is a property that is expected from the diffusion interpretation. Using the negative of the Laplacian as the right side of a differential equation moves the value at a node towards the local average. Since this cannot lead to divergence, no time-dependence can grow without bound, so the negative of the Laplacian cannot have any positive eigenvalues.

To represent the Laplacian as a sum of outer products, we define the incidence matrix. The *incidence matrix* (or *vertex incidence matrix*) of a graph is a matrix Q whose rows correspond to the edges of the graph, and whose columns correspond to its vertices. It is defined as follows: For each row (edge) r , which consists of a connection between vertex b and vertex c , set $q_{rb} = 1$ and $q_{rc} = -1$, and $q_{rj} = 0$ for $j \notin \{b, c\}$. Note that the assignment of +1 or -1 to b or c is arbitrary, so the adjacency matrix is defined only up to sign-flips of each row. However, we will only be concerned with properties of Q that are independent of such sign-flips, especially properties of $Q^T Q$.

In the above case, and ordering the edges in lexicographic order along the rows, the incidence matrix is

$$Q = \begin{pmatrix} +1 & -1 & 0 & 0 & 0 & 0 \\ +1 & 0 & -1 & 0 & 0 & 0 \\ +1 & 0 & 0 & -1 & 0 & 0 \\ +1 & 0 & 0 & 0 & 0 & -1 \\ 0 & +1 & -1 & 0 & 0 & 0 \\ 0 & 0 & +1 & 0 & -1 & 0 \\ 0 & 0 & 0 & +1 & -1 & 0 \end{pmatrix} \quad \text{with (arbitrary) row order} \quad \begin{pmatrix} \{1,2\} \\ \{1,3\} \\ \{1,4\} \\ \{1,6\} \\ \{2,3\} \\ \{3,5\} \\ \{4,5\} \end{pmatrix}.$$

The basic observation that we need to demonstrate is that $L = Q^T Q$. To see this, i.e., to see that $Q^T Q = A - D$: On the diagonal (say in position i), $Q^T Q$ has a contribution of $(\pm 1)^2$ for every edge that begins or ends at vertex i , which is the degree of that vertex. Off the diagonal, $Q^T Q$ is nonzero only if column i and column j have a nonzero entry in the same position (e.g., row r). This means that there is an edge (the edge corresponding to row r) that starts at i and ends at j , or vice-versa. So these entries in Q must be opposite in sign, and $(Q^T Q)_{ij} = -1$, as required by $Q^T Q = D - A$. Conversely, $(Q^T Q)_{ij} = 0$ means that no edge connects i to j , and that the adjacency matrix is zero at that location, also as required by $Q^T Q = D - A$.

Since each row of Q sums to 0, the vector $\vec{1}$ consisting of a column of 1's is in the null space of Q , and hence, an eigenvector of $Q^T Q$ with eigenvalue 0. (We already knew this, from the above comment about what is in the null space of the Laplacian.)

The relationship of L to Q now shows that all eigenvalues of L are non-negative. For say λ is an eigenvalue. Then $\vec{x}^T L \vec{x} = \vec{x}^T Q^T Q \vec{x} = (Q \vec{x})^T Q \vec{x} = |Q \vec{x}|^2$ but also $\vec{x}^T L \vec{x} = \vec{x}^T \lambda \vec{x} = \lambda \vec{x}^T \vec{x} = \lambda |\vec{x}|^2$, so $\lambda \geq 0$.

The same argument shows that the normalized Laplacian is also positive-definite:

$L_{norm} = D^{-1/2} L D^{-1/2} = D^{-1/2} Q^T Q D^{-1/2} = (Q D^{-1/2})^T Q D^{-1/2}$. From this, it follows that the eigenvalues of the transition-matrix Laplacian (an asymmetric matrix), are also non-negative, since $L_{transition} = D^{-1} L = D^{-1/2} D^{-1/2} L D^{-1/2} D^{1/2} = D^{-1/2} L_{norm} D^{1/2}$. So the transition-matrix Laplacian and the normalized Laplacian are similar matrices, and their eigenvalues are the same.

The relationship $L = Q^T Q$ provides another application of the Laplacian. Writing out $\vec{x}^T L \vec{x} = \vec{x}^T Q^T Q \vec{x}$ in coordinates yields $\vec{x}^T L \vec{x} = \vec{x}^T Q^T Q \vec{x} = \sum_{i \sim j} (x_i - x_j)^2$. The eigenvector of lowest positive eigenvalue minimizes this quantity, subject to the constraint $|\vec{x}|^2 = 1$.

The interpretation is as follows: let's say that you wanted to arrange the graph vertices at points x_i along a line, so that vertices that connected map to points that are close together. Then you would want to minimize the above expression, $\vec{x}^T L \vec{x} = \sum_{i \sim j} (x_i - x_j)^2$. Obviously you could do this by setting all $x_i = 0$, but this wouldn't preserve the structure of the graph. So it makes sense to add an overall size constraint, $|\vec{x}|^2 = 1$. Another way to see this connection is that the distribution of heat on the graph that dies out the slowest (other than the trivial distribution of equal amounts of heat on every vertex) is one in which there is the least difference in heat on connected nodes – since the heat difference is what drives the equalization. So we want to minimize $\sum_{i \sim j} (x_i - x_j)^2$.

Viewing $\vec{x}^T L \vec{x}$ as an ellipsoid, our task is to find its smallest axis, avoiding the one dimension in which it is flat. This is equivalent to finding the smallest eigenvector of L . We will see this below in deriving PCA. Also, it is part of a general way of approaching constrained-minimization problems: the eigenvectors of L are the solutions of the Lagrange multiplier problem $\frac{d}{dx}(\vec{x}^T L \vec{x} - \lambda \vec{x}^T \vec{x}) = 0$, which works out to $L \vec{x} = \lambda \vec{x}^T$.

In sum, the map from vertex i to x_i is the map from the graph to the line that best approximates the graph structure: it minimizes the total squared distance between the vertices that are connected, given a constraint on the total extent of the map.

The idea extends to further dimensions: Consider a map M from each vertex i to the row vector

$v_i = \left(\frac{1}{\sqrt{\lambda_p}} (\vec{\varphi}_p)_i, \dots, \frac{1}{\sqrt{\lambda_q}} (\vec{\varphi}_q)_i \right)$ (where $\vec{\varphi}_p, \dots, \vec{\varphi}_q$ are the eigenvectors of nonzero eigenvalue) is

an embedding of the graph in a space of dimension $q - p + 1$. The first coordinate is proportional to the above mapping to the line; each subsequent coordinate is the best possible map to a line that is orthogonal to all previous mappings. A very hand-wavy argument as to why this set of coordinates is a good way to represent the graph: Let C be the matrix of coordinates

of these vectors, i.e., $C_{i,r} = (v_i)_r = \frac{1}{\sqrt{\lambda_r}} (\vec{\varphi}_r)_i$, which consists of the eigenvectors $\vec{\varphi}_r$ in columns,

with each column normalized by $\frac{1}{\sqrt{\lambda_r}}$. Then $C^T L C = C^T Q^T Q C = C^T \left(\sum_m \lambda_m \varphi_m \varphi_m^T \right) C = I$. So


in some sense, $C^T C$ is like the inverse of the Laplacian, which in turn indicates how easy it is to diffuse between two points. So when two vertices have coordinates (columns of C) with high covariance, they are close in the sense of diffusion on the graph.

The nonzero eigenvalues


A simple calculation

One might think that for a graph whose topology is that of a line, the dominant eigenvector of the graph Laplacian is just a linear function on the line. But this can't actually be the case: where the value assigned to a node is equal to the average of its neighbors, the Laplacian goes to zero. So the Laplacian maps a linear function on the graph to a function that is zero everywhere, except at the endpoints. So unless $N = 2$ or $N = 3$, this cannot be an eigenvector.

So let's calculate the lowest eigenvector and eigenvalue of the Laplacian for a graph consisting of N points in a line. For $N=2$, the eigenvector (trivially) is $[-1 \ 1]$; for $N=3$, almost as trivially, it is $[-1 \ 0 \ 1]$. But for $N \geq 4$, it is not a linear function of position on the graph, and, as the calculations show, it approaches a half cycle of a sinusoid, with a peak at one end of the graph, a zero-crossing in the middle, and a trough at the other end.



We can understand this asymptotic behavior as follows, via the “reflection” trick. For this graph, at every point except the endpoints, the graph Laplacian is a discrete approximation to the (negative of the) second derivative: the value at position x_j is replaced by $2x_j - (x_{j-1} + x_{j+1})$. At the endpoints, this is not the case – BUT – if we imagine adjoining reflections of the graph at each endpoint, the second-derivative behavior is recovered. More formally, we “guess” that the solution $\{x_1, \dots, x_N\}$ is one section of an infinite line (black segment in figure below), where $x_{2N-j} = x_j$ (the mirror) and $x_{2N-2+j} = x_j$ (periodicity). So now a solution to the graph Laplacian heat equation is a segment of a solution to the standard Laplacian heat equation on the standard line, provided that this solution is mirror-symmetric at N , and periodic with period $2(N-1)$.



Since the heat equation on the line is translation-invariant, its solutions are the eigenfunctions of the translation operator – the sinusoids. For $x(s) = \exp(i\omega s)$, $Lx(s) = -\frac{\partial^2}{\partial s^2} \exp(i\omega s) = \omega^2 x(s)$. So the lowest nonzero eigenfunction is the one with the lowest frequency consistent with the mirror-symmetry condition, i.e., a sinusoid with period $2(N-1)$.

This analysis properly suggests that group-theoretic considerations will help understand the solution of the heat equation on other graphs. If there is a permutation of the vertices that leaves the graph’s connectivity unchanged, then it necessarily leaves the Laplacian unchanged as well.

So we can immediately calculate the eigenfunctions of the Laplacian on a cyclic graph with N elements – they are $\vec{x}_j^{[r]} = \exp(\frac{2\pi i j r}{N})$. To find the corresponding eigenvalues:

$$(L\vec{x}^{[r]})_j = 2(\vec{x}^{[r]})_j - (\vec{x}^{[r]})_{j-1} - (\vec{x}^{[r]})_{j+1} = e^{\frac{2\pi i j r}{N}} (2 - e^{-2\pi i r/N} - e^{+2\pi i r/N}) = (2 - e^{-2\pi i r/N} - e^{+2\pi i r/N})(\vec{x}^{[r]})_j$$

so $\lambda_r = (2 - e^{-2\pi i r/N} - e^{+2\pi i r/N}) = 2(1 - \cos(\frac{2\pi r}{N})) = \sin^2(\frac{\pi r}{N})$. The smallest nonzero eigenvalue corresponds to $r = 1$.

Further implications of symmetry

Since we're dealing with graphs, there's the possibility for more "interesting" (i.e., larger, non-commutative) symmetry groups to be relevant. For the complete graph with N vertices, the full permutation group leaves the graph invariant. So now our task is to find out how the full permutation group acts in the space of function on N objects.

We recall that every finite group G has a specific, finite set of "irreducible representations", i.e., the distinct ways in which it can operate in a vector space V . These actions are defined by a set of linear operators U_g in $Hom(V, V)$, which respect the group operation: $U_{gh} = U_g U_h$. For any nonzero vector v in that space, the set of images $U_g v$ must span the space. (If they did not span the space, then they would span a subspace of V preserved by G , and U would not be irreducible.)

A main difference between commutative groups (for example, the cyclic group) and non-commutative groups (for example, the full permutation group with 3 or more elements) is that non-commutative groups contains irreducible representations of dimension 2 or more.

What are the consequences of this in the current set-up, when a group acts on a graph in a structure-preserving way? First, the action of the group on the graph induces an action of the group on functions of the graph. This means that there is a representation of the group (say, W) in the vector space X of functions on the graph. But also, since the group operation preserves the graph structure, it necessarily preserves the Laplacian. So for any group element g , $W_g L = L W_g$.

Let's say we have an eigenvector $\vec{\varphi}$ of L , with $L\vec{\varphi} = \lambda\vec{\varphi}$. Then for any group element g , $L(W_g\vec{\varphi}) = W_g L\vec{\varphi} = W_g \lambda\vec{\varphi} = \lambda(W_g\vec{\varphi})$, so $W_g\vec{\varphi}$ is also an eigenvector of L , and has the same eigenvalue. Since there may be some irreducible subspaces of X that have dimension 2 or more, this means that two or more eigenvectors of L will be related by the group symmetry, and have the same eigenvalue. So this puts a premium on finding the irreducible components of how G acts in the space X of functions on the graph – for every irreducible component of dimension k , we will find k eigenvectors sharing the same eigenvalue.

Now consider the specific case of the complete graph of N points. The space of functions on the graph also has dimension N . The elements of the permutation group act on these functions via $N \times N$ permutation matrices. We can tell that this representation is *not* irreducible by a simple character calculation. Specifically, the number of copies of an irreducible representation M in a

given representation W is $\frac{1}{\#|G|} \sum_g \overline{\chi_M(g)} \chi_W(g)$, where $\chi_W(g) = \text{tr}(W_g)$. For permutation matrices, the trace is the sum of the number of items that are not relabeled. So for the representation W , $\chi_W(g) \geq 0$ everywhere, and, for some group elements, the inequality is strict. Now taking M to be the trivial representation E (that maps every group element to the number 1), it follows that there is at least one copy of E inside of W . So $W = E \oplus X$.

For representations of the full permutation group, there's a classic and general (but very involved) way to show that X is irreducible. Here we will use an elementary argument, by characterizing the matrices that commute with all elements of W . First, consider W_σ for a particular pair-swap permutation $\sigma = (pq)$. The corresponding W_σ is a matrix for which $(W_\sigma)_{p,q} = 1$, $(W_\sigma)_{q,p} = 1$, $(W_\sigma)_{j,j} = 1$ for $j \notin \{p,q\}$, and $(W_\sigma)_{j,k} = 0$ otherwise. If a matrix B commutes with W_σ , then it follows from $BW_\sigma = W_\sigma B$ that $B_{q,p} = B_{p,q}$. Therefore, a matrix that commutes with W_τ all pair-swap permutations must have all of its off-diagonal elements equal to the same value:

$$B = \begin{pmatrix} b & c & \cdots & c \\ c & b & \ddots & \vdots \\ \vdots & \ddots & \ddots & c \\ c & \cdots & c & b \end{pmatrix}. \text{ Finally, since the number of irreducible components of } W \text{ is the number of}$$

linearly independent operators that commute with all W_g , the number of such components is $2 -$ one of dimension 1 (that corresponds to the trivial representation that maps all group elements to the identity) and one of dimension $N - 1$, which consists of the “natural” representation in permutation matrices, but with the subspace corresponding to the identity removed. Since the Laplacian commutes with all of the matrices in this irreducible representation, it must act as a multiple of the identity – so the Laplacian has $N - 1$ eigenvectors in this space, and they all have the same eigenvalue.

More explicitly: the eigenvector corresponding to the trivial representation is the vector $\vec{1}$ of all 1's, and it has eigenvalue 0. The other eigenvectors consist of any vector whose mean is zero, as this guarantees that they are orthogonal to $\vec{1}$. In this subspace, the Laplacian acts by replacing each element x_i by $(N - 1)x_i - \sum_{j \neq i} x_j = Nx_i - \sum_j x_j = Nx_i - (\vec{x} \bullet \vec{1} = \sum_j x_j = 0)$. So the eigenvalues in this space are N .

We can use a similar approach for other many other simple graphs. For example, in the star-topology, all of the peripheral vertices are equivalent under the action of the full symmetric group; the wagonwheel topology, the cube topology, and many others can be treated in a similar way.

Symmetry arguments are also useful when there is an approximate symmetry, as in this case, the exact eigenvalues and eigenvectors can be considered to be perturbations of the corresponding

quantities of the fully symmetric graph. So one can group the eigenvalues and eigenvectors in a meaningful way.

Embedding in a simpler context

It is useful to compare the calculation of the eigenvectors of the Laplacian with basic procedures for exploratory analysis of multidimensional data that arise in a simpler context: principal components analysis and multidimensional scaling. In principal components analysis (PCA), the data already have coordinates, and we look for a dimensionally-reduced set of coordinates that still represents the same data well. Multidimensional scaling is a bridge from PCA to Laplacian embedding – in MDS, one knows the distances between the datapoints but not their coordinates, and seeks a dimensionally-reduced representation. For the graph Laplacian, one doesn't know distances, one just knows which vertices are connected.

See material in MVAR2021.pdf for more details on PCA and MDS.

Combinatorial measures

We now return to graphs, and consider some examples of “combinatorial” measures – measures that emphasize counting connections, rather than the algebraic structure of A .

Global clustering coefficient

The *global clustering coefficient* C is a way of measuring whether the connectivity is mostly local (highly “clustered”) or not. It is defined as three times the ratio of the number of triangles to the number of connected triples, where a “connected triple” is a subset of nodes i, j , and k for which $i \sim j$ and $j \sim k$. The factor of 3 is because every triangle necessarily has three connected triples (each cyclic order). A complete graph has $C = 1$, a triangle-free graph has $C = 0$.

For a random graph (an Erdos-Renyi graph) in which nodes are connected with probability p , we can calculate that the expected number of connected triples is $\frac{N(N-1)(N-2)}{2} p^2$ and the expected number of triangles is $\frac{N(N-1)(N-2)}{6} p^3$, so the global clustering coefficient is p .

More generally, the global clustering coefficient is not “normalized” for the overall connection density of the graph. It is also not straightforward to extend the idea of a global clustering coefficient to a graph with weighted edges.

Community structure

Community structure is an approach advanced over the last decade to analyze graphs that arise in many different contexts. The idea is to attempt to partition the vertices into “communities”, in a way that most adjacency relationships are within communities, rather than between. See

Newman (PNAS 2006, <http://www.pnas.org/content/103/23/8577.full.pdf+html>) for an application to social networks. For an application to fMRI, see Bassett... Porter... Grafton, <http://www.pnas.org/content/early/2011/04/15/1018985108.full.pdf+html>). We use Porter's normalization conventions but the algebraic approach of Newman.

A community structure consists of an assignment of each vertex i to a community c_i . The extent to which this assignment captures the graph structure is determined by a “quality function”, typically

$$Q = \sum_{i,j} (a_{ij} - p_{ij}) \delta(c_i, c_j).$$

The quantities a_{ij} are the entries in the adjacency matrix. The quantities p_{ij} correspond to the expected probability of a connection between i and j , subject to a null hypothesis. The natural null hypothesis is $p_{ij} = \frac{d_i d_j}{2m}$, where m is the total number of edges. This is the expected

probability of a connection, given the total number of edges in the graph, and the number of edges present at i and at j , and no further organization. Put another way, if you create an adjacency matrix that is as random as possible given a specified set of degrees at each vertex,

then its off-diagonal entries would be $p_{ij} = \frac{d_i d_j}{2m}$. Note that $m = \frac{1}{2} \sum_{i,j} a_{ij} = \frac{1}{2} \sum_q d_q$, since

counting the degree of each vertex counts each connection twice.

Q is large if edges occur within a community in a way that is more than chance ($a_{ij} > p_{ij}$ if $\delta(c_i, c_j) = 1$). In the trivial case that all vertices are assigned to the same community,

$$Q = \sum_{i,j} (a_{ij} - p_{ij}) = \sum_{i,j} \left(a_{ij} - \frac{d_i d_j}{2m} \right) = \sum_{i,j} a_{ij} - 2m \sum_i \frac{d_i}{2m} = \sum_{i,j} a_{ij} - \sum_i d_i = \sum_{i,j} a_{ij} - 2m = 0.$$

Note that any permutation that respects the symmetry of the graph results in a reassignment of communities that leaves Q invariant.

One difficulty in applying this approach is that the general problem of determining the community structure that maximizes Q is intractable – specifically, it is an NP-hard problem, equivalent in difficulty to the traveling salesman problem. However, finding approximately optimal community structures (an informal concept) is much easier. Here we outline the Newman (2006) approach.

The approach is “greedy” – the first step is to find the best partitioning into two communities; at each subsequent step, the existing community structure is the starting point and the algorithm seeks to divide one of the existing communities into two. There is no guarantee that the optimal community structure can be identified this way, so this is the first sense in which the algorithm is inexact. (“Greedy” algorithms typically have this property.) At each stage, the algorithm can terminate – by indicating that no further subdivision improves Q . Newman's algorithm consists of an explicit eigenvector computation at each stage, but he also suggested ways to refine each stage by a local search, in which individual vertices are reassigned to the alternative community.

The key observation is that, for the special case of dividing a graph into two communities, we can rewrite Q as a quadratic form. Let the community structure be determined by $s_i = \pm 1$, where the two signs correspond to the two communities. Then, $\delta(c_i, c_j) = \frac{1}{2}(s_i s_j + 1)$. As noted above, $\sum_{i,j} (a_{ij} - p_{ij}) = 0$ (i.e., for the assignment of all vertices to the same community, $Q = 0$), so

$$Q = \frac{1}{2} \sum_{i,j} (a_{ij} - p_{ij}) s_i s_j .$$

Now using the s_i 's to constitute a column vector \vec{s} , and B to be the matrix whose entries $b_{ij} = a_{ij} - p_{ij}$ express the difference between the true adjacency matrix and the “null hypothesis” matrix, we have

$$Q = \frac{1}{2} \vec{s}^T B \vec{s} .$$

The problem of finding the optimal community structure is now the problem of maximizing this quadratic form, subject to the constraint that all of the components of \vec{s} are ± 1 . If, instead, the constraint was that $|\vec{s}| = 1$, we would simply find the eigenvector of B that corresponds to its largest eigenvalue. So here is the second place at which the algorithm is heuristic rather than exact: the Newman procedure identifies this eigenvector \vec{s}_{max} , and then sets s_i to be ± 1 according to the sign of $(\vec{s}_{max})_i$. This yields the vector \vec{s} that is most closely aligned with \vec{s}_{max} (in the dot-product sense), but not necessarily the one that yields the largest Q . Of course it makes sense that these are similar – and in specific examples, as checked by an exhaustive search, this is the case – but it is not guaranteed. The procedure to be followed when $s_i = 0$ is unspecified, but there are some obvious choices: choose randomly, or do an exhaustive search.

The numerical values of $(\vec{s}_{max})_i$ have a meaning – they are the extent to which each vertex contributes to the community structure.

As is the case with the graph Laplacian, the matrix B is left invariant by any relabeling of the graph vertices that leaves the connectivity matrix invariant. So we can use group-theoretic tools to find B 's eigenvalues in many simple cases (complete graph, cyclic graph, star graph, etc.).

Importantly, the matrix B could be negative-definite -- there may be no eigenvectors with positive eigenvalue. In this case, the Newman procedure simply terminates – as it can find no community structure that increases Q above its value for a single-community assignment. As an example, for the complete graph of size N , $a_{ij} = 1$ and $p_{ij} = \frac{(N-1)(N-1)}{N(N-1)/2} = \frac{2(N-1)}{N}$, so for $N \geq 3$, all off-diagonal $b_{ij} < 0$. In this case, the algorithm stops, as no eigenvector will increase the quality Q of the community structure.

But the situation is more complex: because the a true eigenvector of B is replaced by one whose entries are coordinates are ± 1 , it is possible that B has a positive eigenvector, but none of the community assignments increase Q . In this case, the algorithm again terminates. An example of this is a graph with a “wagonwheel” connectivity and 6 vertices (5 peripheral): B has an eigenvalue of $(\sqrt{5-1})/2 \approx 0.618$; this corresponds to an assignment of 3 vertices to each community and a decrement of Q by 0.2.

What may be even more of a concern – but this relates more to the definition of Q than to the algorithm -- is that Q may be large even for graphs in which the community structure does not seem relevant. For example, partitioning a cyclic graph into two connected halves increases Q by $N - 4$. But there is no particular reason to choose any of the many ways that one can hemisect the graph; all lead to an equal increase in Q . are equally good. Put another way, the matrix B has several eigenvectors that are identical; they are guaranteed to be identical because of the symmetry of the graph. So this suggests a measure of robustness of the community assignment: if the next-largest eigenvalue of B is close to the largest one, the partitioning is not robust.

This also suggests an algorithmic tweak: when there is no gap, or only a small gap, between the largest eigenvalue of B and the next-largest eigenvalue, these top eigenvectors should be considered simultaneously. There are at least two ways of doing this: first, a binary the community assignment could be carried out by finding the vector containing ± 1 's that is closest to the hyperplane that they span; second, one could look for a non-binary partition in that subspace (as in the way that the Fisher discriminant is used for more than binary partitioning).

It is interesting to note that the “community” approach becomes tractable (i.e., reduces to a set of eigenvalue problems) only if one ignores the combinatorial aspects of the problem – or at least, uses the heuristics that they can be approximated by an algebraic approach. It is also interesting to note that the above eigenvalue problem is very similar to the one that arises in the Fisher discriminant, where the problem is to find an axis that maximally separates two clusters. The difference is that in the Fisher problem, one is given coordinates of the points (and from this, one can compute their distances). Here, one is only given the distances between the points (effectively 1 if connected, large if not connected), and one needs to put them into a space first, so that one can find directions.

Analogous to the Fisher problem, one wonders whether a less-greedy approach, in which one allows for a multipart subdivision at a single stage if B has more than one positive eigenvalue – would be a useful generalization.

See material in MVAR1819 for further details on linear discriminant analysis a.k.a. the Fisher discriminant.

Comparing community structures

Often one has several observations of data that one represents as a graph – e.g., measures of functional connectivity on several occasions, or on different tasks, or in different subjects. Given

such data, one might want to compare the graphs, or, one might want to carry out a joint analysis to determine which set(s) of connections are robust across conditions.

See for example Bassett et al, Dynamic reconfiguration of human brain networks during learning. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3088578/>, for algorithms and an interesting application.

Topological Data Analysis: Persistent homology

Motivation

Topological data analysis is an approach to characterize the global properties of data. If we think of the data as a collection of a large number of points, we could wonder whether these points formed a single cloud, several clouds, a ring, a donut, a spherical shell, a distorted pretzel, etc – i.e., how can we characterize the “shape” of the data, in terms of properties that are invariant with respect to continuous deformation of the space.

The starting point is the sample of points in the space that we observe. We then construct a graph by connecting the pairs of points that are sufficiently close. We then think of this graph as a skeleton of the full dataset, and attempt to characterize it: how many disjoint pieces are there? Are there portions of the graph that have “holes,” i.e., four or more points connected in a ring, but no connections bridging across the interior of the ring? Are there portions of the graph that have “voids” (3-d analogs of holes), i.e., a geodesic-dome-like structure, with no connections bridging across the interior? What about higher-dimensional analogs?

We need tools to count up these structures. We can then apply them to a graph, constructed by assigning connections to pairs of nodes whose distance is less than some criterion. The distance can be defined in many ways – for example, the actual distance in the space of the data, or the geodesic distance within the data domain, or from a correlation coefficient. We then want to apply the analysis parametric in the distance criterion, i.e., apply the analysis to a sequence of graphs that have more and more connections as we relax the distance criterion – what is known as a “filtration.”

Below we describe this counting-up process, which requires what may be a surprising amount of mathematical structure. The basic reason that a sophisticated approach is required is that one wants to have a way to only count loops that are independent – for example, if the graph formed a pretzel, one would only want to count three loops (the loops around the holes), and not a fourth loop that went around the entire pretzel. But how should this be done in general? And how to generalize this to higher-dimensional “holes”?

Simplexes

These entities generalize the notion of a graph. As we will see, one can always create simplexes from a graph, but there are also other ways of making simplexes.

An “ n -simplex” is merely a labeled set of n points, along with an ordering of those points.

Critically, two orderings are considered the same if they differ by an even permutation. We also define the negative of a simplex is the same set of points, but associated with an ordering that differs by an odd permutation.

A 0-simplex is merely a point. Ordering a single point is vacuous.

A 1-simplex is a pair of points, say X and Y , and the ordering is the same as a choice of polarity (i.e., (X, Y) vs. (Y, X)). So a 1-simplex consists of two linked nodes of a graph and an assignment of a direction to the link between them.

A 2-simplex is a triplet of points, and the ordering is a choice of a cyclic order (i.e., (X, Y, Z) vs. (X, Z, Y)). Note that since permutations that differ by an even permutation are considered to be the same “order”, (X, Y, Z) , (Y, Z, X) , and (Z, X, Y) are equivalent orders; while (X, Z, Y) , (Z, Y, X) , (Y, X, Z) are all the opposite order. Thus, a 2-simplex can be considered to be a three-element clique of graph, and an assignment of “clockwise” vs. “counterclockwise” direction to the triangle that they form.

A 3-simplex is a quadruplet of points, an assignment of an order and all its even-permutation equivalents. So a 3-simplex is a four-element clique of a graph, and an assignment of handedness to the tetrahedron that they form.

Given a set of points and distances between them, we can construct a sequence of graphs, parameterized by a threshold D , by connecting all points that have a distance D or less. For sufficiently small D , the graph will be completely disconnected. For sufficiently large D , all points will be connected. For intermediate values, we can anticipate that the structure of the graph will describe the topology of the space that the points are drawn from. For each of these graphs, we can consider the collection of the n -simplexes of every dimension n .

Note that there are other ways of building a sequence of n -simplexes besides using distances. For example, the points could represent individuals, and we could group the individuals into simplexes if the individuals were all members of the same club. Then not every triplet of points which formed 1-simplexes would necessarily form a 2-simplex – only if all three were in the same club. In an early application of this approach (Dabaghian et al., PLOS Computational Biology 2012), simplexes were formed from hippocampal place cells, according to whether their place fields had an overlap.

A “simplicial complex” is merely a collection of simplexes, which may be of different orders.

A linear structure on simplexes

Next, we will create a vector spaces associate with n -simplexes. We will then create a crucial linear operator from the vector space of n -simplexes to the vector space of $(n - 1)$ -simplexes, known as the “boundary.”

The vector space corresponding to a set of simplexes σ_i of dimension n consists is the “free vector space” on the simplexes, -- that is, the set of all formal sums $\sum_i c_i \sigma_i$, where the coefficients c_i are drawn from a field. Addition and scalar multiplication is term-by-term. Two formal sums are identical only if each of their terms are identical. As above, we consider $-\sigma_i$ to

be the same simplex as σ_i , but with the alternative ordering. We denote the vector space for the n -simplexes by V_n .

This construction will allow us to define a the “boundary” operator, a mapping from V_n to V_{n-1} , that will solve the “counting of holes” problem. The “boundary” operator is traditionally denoted (unfortunately) by ∂ . The boundary operator is defined on the simplexes themselves, and then extends to the entire vector space by linearity.

The boundary of every 0-simplex (a point) is zero.

The boundary of a 1-simplex $\sigma = (X, Y)$ is defined as a difference of two 0-simplexes, namely, $\partial(X, Y) = (Y) - (X)$. That is, the boundary of a directed segment is the difference between its endpoints.

The boundary of a 2-simplex $\sigma = (X, Y, Z)$ is a linear combination of three 1-simplexes, namely, $\partial(X, Y, Z) = (X, Y) - (X, Z) + (Y, Z) = (X, Y) + (Y, Z) + (Z, X)$. That is, the boundary of a 2-simplex is the triplet of 1-simplexes that enclose it.

In general, for an n -simplex

$$\partial(X_0, X_1, X_2, \dots, X_n) = (X_1, X_2, \dots, X_n) - (X_0, X_2, \dots, X_n) + (X_0, X_1, X_3, \dots, X_n) \dots + (-1)^n (X_0, X_1, X_2, \dots, X_{n-1})$$

The boundary operator, which we have defined on simplexes, clearly extends to sums of simplexes, i.e., to all of V_n . But we only have to check that $\partial(-\sigma) = -\partial(\sigma)$. That this holds can be seen by considering a single pair-swap of σ and how it propagates from V_n to V_{n-1} .

To see how the boundary operator this might help us, we first consider a graph with several disconnected components and use this machinery to count the number of such components. The dimension of the vector space V_0 is the number of points. We want to show that if two points are in the same component of the graph, then the difference between them – which is a formal sum in V_0 -- is in ∂V_1 . If two points (say A_0 and A_1) are connected, there is a 1-simplex


$\sigma = (A_0, A_1)$, whose boundary yields this difference: $\partial(A_0, A_1) = (A_1) - (A_0)$. If, further, a point A_2 is not directly connected to A_1 but only via A_0 , then this still holds: $\partial(A_0, A_2) = (A_2) - (A_0)$ and $\partial(A_1, A_2) = (A_2) - (A_1)$ so $(A_2) - (A_0) = \partial(A_1, A_2) + \partial(A_0, A_1)$ is in ∂V_1 . In a similar manner, for all points A_j that are in the same component as A_0 , it follows that $(A_j) - (A_0)$ is in ∂V_1 , namely, it is the boundary of the sum of 1-simplexes that link A_j to A_0 .

What this means is that within each connected component, we can reach any element of V_0 by specifying the coefficient at a single node in that component (say A_0), and then adding an appropriate element of ∂V_1 to yield the desired coefficients at all the connected nodes.

Conversely, if a point B is in a different component of the graph than A (i.e., there is no sequence of links that connects them), then $(B) - (A_0)$, which is in V_0 , is not in ∂V_1 . So we cannot specify the coefficient of (B) by adding an element of ∂V_1 to (A_0) .


That is, the dimensionality of V_0 exceeds the dimensionality of ∂V_1 by the number of disconnected components of the graph.

Now let's extend the idea to the relationship between V_1 and ∂V_2 . V_1 , essentially, consists of all the formal sums of directed edges. But since we want to count "holes", let's only consider the sums of directed edges in V_1 that form a closed circuit – for example, this graph, which has links from A_0 to A_1 , A_1 to A_2 , A_2 to A_3 , and A_3 to A_0 .



So we consider an element in V_1 corresponding to this loop, given by $v = (A_0, A_1) + (A_1, A_2) + (A_2, A_3) + (A_3, A_0)$. Note that $\partial v = 0$. And conversely, if, for some $w \in V_1$, $\partial w = 0$, then w corresponds to a sum of such closed circuits. Put another way, the elements in V_1 that are in the kernel of ∂ are the elements that can be decomposed into a sum of closed circuits, where the weights within each closed circuit is a constant.

Now, when can $v = (A_0, A_1) + (A_1, A_2) + (A_2, A_3) + (A_3, A_0)$ be written as the boundary of some $x \in V_2$? In the above example, it cannot – as there are no 2-simplexes.



But in the example below – in which the graph also has a connection from A_0 to A_2 -- there are now two 2-simplexes, (A_0, A_1, A_2) , and (A_0, A_2, A_3) . We compute $\partial(A_0, A_1, A_2) = (A_1, A_2) - (A_0, A_2) + (A_0, A_1)$ and $\partial(A_0, A_2, A_3) = (A_2, A_3) - (A_0, A_3) + (A_0, A_2)$, so

$$\begin{aligned} \partial((A_0, A_1, A_2) + (A_0, A_2, A_3)) &= \\ (A_1, A_2) - (A_0, A_2) + (A_0, A_1) + (A_2, A_3) - (A_0, A_3) + (A_0, A_2) &= \\ (A_1, A_2) + (A_0, A_1) + (A_2, A_3) + (A_0, A_3) &= v \end{aligned}$$

That is we "filled in" the hole that v encircled, and now we can find v in the linear span of the boundaries of 2-simplexes.

The general picture

Just as the way that the 1-simplexes whose boundaries were zero corresponded to closed paths, the 2-simplexes whose boundaries are zero correspond to closed surfaces. And just as a closed path that is NOT the boundary of a 2-simplex is a "hole", a closed surface that is NOT the boundary of a 3-simplex is a "void" (bubble). The free-vector space construction allows us to count now many of these are independent – this corresponds to the dimension of the space.

So, to count the number of k -dimensional holes/voids/etc., we want to determine how many objects whose boundaries are zero are NOT boundaries of a $(k+1)$ -dimensional simplex. To count the number of k -dimensional objects whose boundaries are zero, we use the dimension of the subspace of V_k for which $\partial V_k = 0$, which we call C_k -- the "cycles." To count the number of cycles that ARE the boundaries of a $(k+1)$ -dimensional simplex, we use the dimension

of the subspace of V_k formed by ∂V_{k+1} , which we call U_k -- the “boundaries.” Then the count of the k -dimensional holes/voids/etc. is given by $B_k = \dim(C_k) - \dim(U_k)$, which is known as the k th Betti number.

An important fact that we haven’t shown is that U_k is a subspace of C_k . That is, $\partial(\partial V_{k+1}) = 0$. This is also something that can be seen directly from the definition of the boundary operator. As a consequence, $B_k = \dim(C_k / U_k)$, where C_k / U_k is a “quotient space,” namely, a vector space whose elements are the sets of elements of C_k that differ by an element of U_k .

A characterization via persistent homology then consists of carrying out the above calculations across a filtration, i.e., a sequence of simplicial complexes. This sequence is often defined by the cliques in a graph, where the nodes of the graph become progressively more filled-in as the criterion for connection becomes more permissive. But (as mentioned above) simplexes can be defined in other ways, and it need not be the case that every k -clique is a k -simplex. **There is an option for further generality, as the field used for the vector-space calculations need not be real or complex numbers. Choosing a finite field can lead to a richer characterization of the topology of the dataset.**

The dependence of each B_k on the filtration threshold characterizes the topological properties of the space from which the data points are sampled. In practice, to understand these characterizations, it is necessary to simulate the Betti-number calculations in a model, as analytic expressions for the Betti-number curves are not available.

Below is an example of this from Zhou, Smith, and Sharpee, *Sci. Adv.* 2018: the dashed curves are the empirical Betti numbers derived from a graph based on similarities of odor samples; the solid curves are the Betti numbers from alternative models – hyperbolic space (left) and Euclidean space (right).

