Rowan Shigeno May 2, 2025

Cycles on Metric Graphs of given Entropy.

Abstract

In this paper, we examine connected metric graphs G (graphs for which it is possible to travel via edges to all vertices in the graph, and with lengths of all edges being specified). Loops (edges that connect vertices to themselves) and multiple edges between the same pair of vertices are allowed, and there are no bounds on edge length. Our goal is to analyze the nature of the growth of the number of possible cycles in such graphs less than or equal to a specified length as the number of vertices and edges (and therefore paths) increase. Specifically, we will introduce a probability measure μ for analyzing the nature of such cycles, and find that there are multiple, seemingly unrelated ways of computing μ. Subsequently, we run a series of experiments on different graphs and their subgraphs to obtain a greater understanding of the number of cycles and μ as a consequence of this.

Contents

1 Introduction

We begin by considering the definition of a connected metric graph, with examples, and proceed to think about what it means to count cycles on such graphs (under certain restrictions regarding what we can count as valid cycles). We then examine a related quantity called entropy; after deriving some properties of entropy, we then consider what it means to contract a graph and how this is important with respect to our understanding of

entropy. Subsequently, we introduce the Gibbs measure, essentially a probability measure pertaining to the distribution of cycles across edges of the graph, and explain how it can be derived using the Perron-Frobenius theorem.

After some proofs and technical computations, we can then derive a much more precise relationship between entropy and edge lengths on a graph, using the Gibbs measure to our aid. Finally, we examine subgraphs on roses of given entropy.

To start, we clarify what we mean by connected metric graphs.

Definition 1. *A **graph** G is defined to be a set of **vertices**, V (points in $\mathbb{R}^2$), along with a set of pairs of vertices called **edges**, E (whose elements are not necessarily distinct), which we can represent visually as lines linking pair of points. That is, $G = (V, E)$ [**83**, 4].*

*We say that a graph is **connected** if any vertex in the graph is reachable from any other vertex via some sequence of edges. A **metric graph** is a graph whose edges are weighted, i.e. assigned positive numerical values.*

For the sake of what we are interested in, we will not restrict ourselves to "straight" edges, or for that matter even edges with lengths that allow us to realistically construct them in Euclidean space. Additionally, we will allow for self-loops, or edges that connect vertices to themselves.

For instance, the set of points representing the corners of this page, along with 7 edges corresponding to the 4 edges of the page, a diagonal edge, and 2 self-loops (one on each non-diagonal-adjacent vertex), could be considered such a graph; alternatively, a single vertex with 4 self-loops of lengths 1/10, π, 7, and a billion is also perfectly valid.

Definition 2. *A **cycle** on graph G is a path, or a sequence of oriented edges $e_1, e_2, \ldots, e_n$, that "close up": that is, for each i, the initial vertex of e_{i+1} is the terminal vertex of e_i, and the initial vertex of e_1 is the terminal vertex of e_n. A **reduced cycle** is a cycle that doesn't back-track, meaning we cannot have the sequence of edges e_i then e_i' (e_i with the opposite orientation), for any i at any point in the reduced cycle.*

Definition 3. *The **length** of a cycle is the sum of the weightings of all the edges it contains. We let $\mathbf{N_G(L)}$ denote the number of possible reduced cycles in the graph G that are of total length $\leq L$.*

For example, considering the "lamppost" graph shown below, let $\ell(e)$ denote the length or weighting of edge e, and set $\ell(a) = 1$, $\ell(b) = 2$, $\ell(c) = 4$,

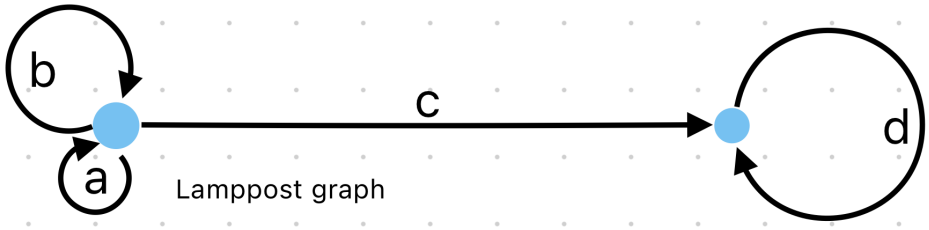

Figure 1: example of a connected metric graph.

and $\ell(d) = 3$. Then examples of reduced cycles with length ≤ 20 include $d'd'd'd'd'$ (d in reverse 5 times), $acdc'bb$, $baab'b'ab$, and $c'acd'c'a'$; in contrast, some invalid examples are a^{21} (total length > 20), $ddd'd$ (not a valid reduced cycle, due to back-tracking), and $abcdcd$ (not a valid cycle due to teleportation from the right vertex to the left).

At this point, we should additionally note that vertices that have fewer than 3 neighboring edges (i.e. a valency less than 3) are not of our concern, as their contributions to cycle-counting are insignificant: this is because vertices with only 1 neighboring edge do not contribute any cycles (as a cycle cannot pass through that edge without back-tracking), and vertices with only 2 neighboring edges are essentially just subdivisions of a single edge into two (and thus have no contribution to the overall complexity or nature of cycles on the graph).

In general, for a vertex being adjacent to k edges (not necessarily distinct, as the beginning of one may be the end of another in the case of a self-loop), a reduced cycle arriving at that vertex can only leave along one of the other $k-1$ edges. In the simplest case, a graph has only 1 vertex and n self-loops (more on this later): then the initial loop has $2n$ possible paths to begin from (counting both the "forward" and "reverse" directions of each edge), and each subsequent edge has $2n-1$ options in order to avoid backtracking. In general, we have the following theorem.

Theorem 1. *$N_G(L)$ grows exponentially in L, i.e. $N_G(L) \sim e^{K \cdot L}$, where K depends on the topology of G and the weights. We call K the **entropy** of the graph, which can also be written as*

$$K = \lim_{L \to \infty} \frac{\log N_G(L)}{L} \text{ [2, 5].} \tag{1}$$

We also introduce an important quantity called the rank of a graph, where the rank of a single-edge, single-vertex graph is 1. In general, we

take the the rank of graph G to be *the number of petals (self-loops) in the contraction of G*. This is explained on the following page.

Theorem 2. *For connected metric graph G, $K > 0$ if and only if rank$(G) > 1$.*

An exception to this positive-entropy rule is the graph consisting of a single vertex with a single loop (edge) to itself— this graph only has rank 1, with only 2 neighboring edges to the vertex and therefore no leeway in cycle choices with each pass. As the number of total reduced cycles of this particular graph grows linearly with L (as there is only one possible direction to continue one each loop), $K = 0$ in this case.

In regards to one of our ultimate goals of investigating entropic properties of subgraphs, an important question will be: given a fixed graph of rank ≥ 2 with an entropy of $K = 1$, does there exist some $\lambda > 0$ (depending *only* on the number of vertices) s.t. there is a **proper** subgraph with entropy $\geq \lambda$?

Now instead of "graph", it should be noted that we are in reality referring to a "1-complex" (which allows for loops and multiple paths between the same pair of vertices) in the literature that follows [**7**, 4]. So, our object is really a {1 complex, weightings of edges}— but for simplicity, we will continue to refer to these objects as "graphs".

An important property of entropy is described as follows.

Lemma 1. *$h(\lambda G) = \frac{1}{\lambda} h(G)$, where $h(G)$ is the entropy of G [**15**, 1], and "λG" is simply a scaling of all edge weights of G by a factor of λ. Notice in particular that when $\lambda = h(G)$, $h(h(G) \cdot G) = 1$.*

Observation 1. *A reduced cycle of length L on G has length λL on λG. This follows directly from the definition of λG.*

Proof of lemma. Use the observation and the definition of entropy. Letting $T = L/\lambda$,

$$\begin{aligned} \lambda h(\lambda G) &= \lambda \lim_{L\to\infty} \frac{1}{L} \log N_{\lambda G}(L) = \lambda \lim_{L\to\infty} \frac{1}{L} \log N_G(L/\lambda) \\ &= \lambda \lim_{T\to\infty} \frac{1}{\lambda T} \log N_G(T) = \lim_{T\to\infty} \frac{1}{T} \log N_G(T) = h(G). \end{aligned}$$

□

Now given a graph G, we can **contract** the edges of the graph by each edge that connects two distinct vertices down to length 0, thereby combining the vertices. Notice that every time we do this, we lose 1 edge and 1 vertex, as illustrated by figure 2; and by default, we continue contracting until we are left with only 1 vertex and no non-loop edges.

Lemma 2. *Let V be the number of vertices and E the number of edges of G. Then the number of resulting petals (i.e. the number of remaining edges, or self-loops) after contraction is $E-(V-1)$.*

Corollary 1. *The number of petals in the resulting "rose" of the contracted graph does not depend on choices we made regarding the order in which we contract the edges [**2**, 4].*

This allows us to make the following definition:

Definition 4. *The **rank** of a graph $G=(V,E)$ is equal to $E-V+1$.*

This is clear from the fact that each contraction removes an edge and a vertex, and the number of contractions we make is equivalent to the number of vertices minus 1 (since we remove all but 1 vertex). Thus, the number of edges or petals we are left with is $E-(V-1)$.

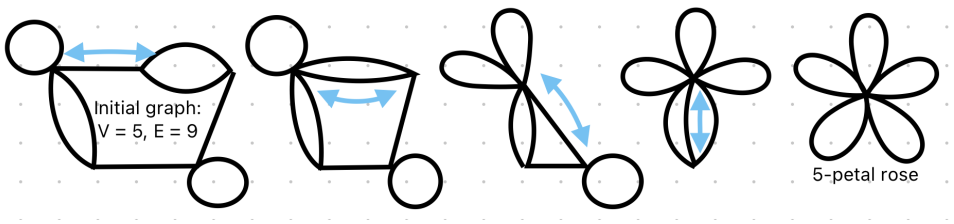

Figure 2: contraction of a graph to a 5-petal rose.

2 Some Computations

As roses are the easiest types of graphs to analyze in terms of cycle-counting, entropy computation, etc., they will hold a crucial role in this paper. To begin, we compute the entropy of a general rose.

Theorem 3. *The entropy of an n-petal rose whose petal lengths are all 1 is equal to* $\log(2n-1)$.

Proof. For a 2-petal rose with each petal length $=1$, there are $4\cdot 3^{L-1}$ paths of length L, since there are 3 paths to choose from each time the path goes through the vertex, and 4 paths for the initial choice. Thus, there are $4\cdot 3^{L-2}$ paths of length $L-1$, $4\cdot 3^{L-3}$ of length $L-2$, and so on, until we find only 4 paths of length 1. Since

$$1+3+3^2+3^3+\cdots+3^{L-1}=\frac{1-3^L}{1-3}=\frac{3^L-1}{2},$$

this is a total of

$$4 \cdot \frac{3^L - 1}{2} = 2(3^L - 1) = 2(e^{L \log 3} - 1)$$

paths of length $\leq L$.

Similarly, for a 3-petal rose, there are $6 \cdot 5^{L-1}$ possible paths of length L, $6 \cdot 5^{L-2}$ of length $L-1$, $6 \cdot 5^{L-3}$ for $L-2$ ... and 6 paths of length 1. This is a total of

$$6 \cdot \frac{5^L - 1}{4} = \frac{3}{2}(5^L - 1) = \frac{3}{2}(e^{L \log 5} - 1)$$

paths of length $\leq L$.

The pattern holds for the n-petal rose, which has exactly $2n \cdot (2n-1)^{L-1}$ paths of length L. This yields a total of

$$2n \cdot \frac{(2n-1)^L - 1}{(2n-2)} = \frac{n}{n-1}((2n-1)^L - 1) = \frac{n}{n-1}(e^{L \log(2n-1)} - 1)$$

paths of length $\leq L$.

Now since we know that $N_G(L)$ grows exponentially in L, i.e. $N_G(L) \sim e^{K \cdot L}$, and that $\lim_{n \to \infty} \frac{n}{n-1} = 1$, we can conclude that $K = \log(2n-1)$ for the n-petal rose. □

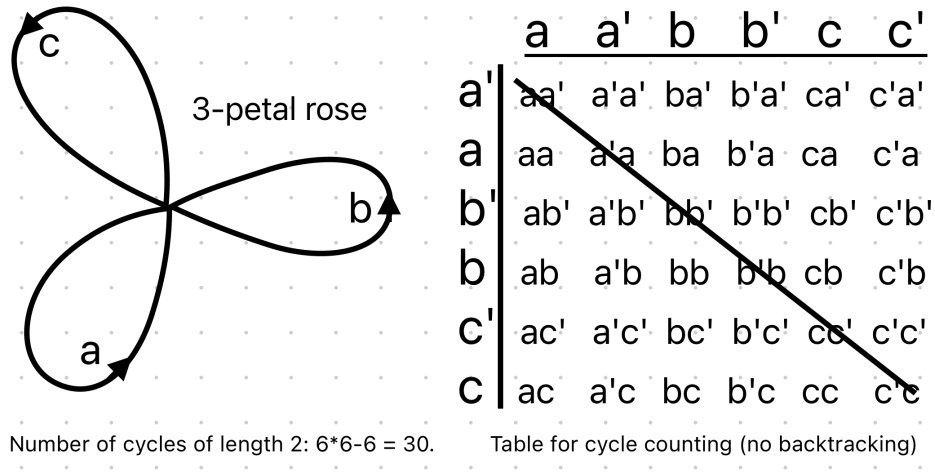

Figure 3: reduced cycles on a 3-petal rose.

We would now like to place upper and lower bounds on the quantity $N_G(L)$.

Theorem 4. *If rank$(G) \geq 2$, then there exists constants $c_1, c_2 > 0$ (depending on G) such that $e^{c_1 \cdot L} \leq N_G(L) \leq e^{c_2 \cdot L}$, for sufficiently large L. (Note: it turns out that $\lim_{L\to\infty} N_G(L)/e^{K\cdot L} = 1$).*

Proof. For a graph G with directed edges $E = \{e_1, e_1^{-1}, \ldots, e_n, e_n^{-1}\}$, the lengths of the edges are $\{l(e_1), \ldots, l(e_n)\}$. If e_{min} is the edge such that $l(e_{min}) \leq l(e_i)$ for all i, then any reduced cycle of length $\leq L$ has no more than $\lceil L/l(e_{min})\rceil$ edges, as the number of edges we can squeeze into this length is maximized when $l(e_i)$ is at its minimum. Consequently, the upper bound for $N_G(L)$ is given by the sum of all reduced cycles that pass through k total $l(e_{min})$-length edges (for all k s.t. $k \leq L$), i.e.

$$n + n^2 + n^3 + \cdots + n^{\lceil L/l(e_{min})\rceil} = \frac{n}{n-1} \cdot n^{\lceil L/l(e_{min})\rceil - 1}.$$

As $n^{L/l(e_{min})} = e^{\log(1/l(e_{min}))\cdot L}$, we let $c_2 = \log(1/l(e_{min}))$.

Now to obtain the lower bound, first notice that if these petals have length 1, $h(G') = \log(2n-1)$, where n is the number of petals (the rank). But more generally, we must observe that if we take G' to be the contraction of G without alterations in edge lengths (apart from contracted edges), we have $h(G') \geq h(G)$. In contrast, increasing the lengths of our edges beyond a certain extent results in $h(G') \leq h(G)$, since we can make the entropy as small as we would like via the fact that $h(\lambda G') = \frac{1}{\lambda} h(G')$.

To proceed, we introduce the fact that the entropy of a graph G with the set of edges E can be interpreted as a function $h : \mathbb{R}_{>0}^{|E|} \to \mathbb{R}$ ($|E|$ being the size of E or number of edges), as the metrics on $G \longleftrightarrow \mathbb{R}_{>0}^{|E|}$. That is, h is function that takes the edges of a graph G as an input, and returns the entropy of that graph as a real number. Furthermore,

Theorem 5. *The entropy of a graph is in fact an infinitely differentiable and also real-analytic function [1].*

And since entropy is continuous as a result, we can use the intermediate value theorem to conclude that there must exist an intermediate G', such that $h(G') = h(G)$, which is a rose we obtain by shrinking (contracting) the non-rose edges of G while simultaneously expanding the self-loop-edges. Therefore,

Lemma 3. *Given any metric graph G, there exists a rose R such that $h(G) = h(R)$.*

This implies that given $\epsilon > 0$, there is an $l > 0$ s.t. for all $L > l$, $|\frac{N_G(L)}{N_R(L)} - 1| < \epsilon$. So, for this L, $N_G(L) \geq (1-\epsilon)e^{\log(2n-1)\cdot L/\lambda}$, where λ is the

petal length for the intermediate rose. So,

$$(1-\epsilon)e^{\log(2n-1)\cdot L/\lambda} \leq N_G(L) \leq e^{\log(1/l(e_{min}))\cdot L}. \tag{2}$$

In fact, the limit of $N_G(L)$ always exists [**3**, 6], and is bounded precisely as stated above. □

3 The Gibbs Measure

To be able to analyze the behaviors of cycles on more complicated metric graphs, we need a method of describing the profile of a randomly chosen reduced cycle on G. Thus we introduce a vector μ, which we define as a probability measure corresponding to the amount of length that a generic cycle spends on each edge in the graph in question.

Theorem 6. *Let G be a metric graph with rank ≥ 2. Then there exists a vector $\mu = (\mu_{e_1}, \mu_{e_2}, \ldots, \mu_{e_n}) \in \mathbb{R}^n$ s.t.*

(i) $\mu_{e_i} > 0$, $\forall i$, *and*

(ii) $\sum_{i=1}^{n} \mu_{e_i} = 1$*: so that for each i,*

$$\lim_{L\to\infty} \frac{\sum_{\gamma|\ell(\gamma)\leq L} \frac{\gamma_{e_i}}{\ell(\gamma)}}{N_G(L)} = \frac{\mu_{e_i}\ell(e_i)}{\sum_{i=1}^{n} \mu_{e_i}\ell(e_i)} \tag{3}$$

*[**5**, 5], where γ_{e_i} is the amount of length that the path given by γ spends on the edge e_i in question. We call the quantity μ_{e_i} the **Gibbs measure** for edge e_i.*

Each component of the Gibbs measure vector μ is positive, and the size of each component conveys what proportion of the length of a randomly chosen reduced cycle of that graph is spent on that particular edge. This vector is normalized so that that the components add to one— this way, we can consider it to be a probability measure.

The left-hand side of the above equation represents the limit as L goes to infinity of the ratio of the sum of proportions of reduced cycles of length $\leq L$ spent on e_i, to the number of such cycles— thus yielding the average expenditure on e_i. The fraction on the right-hand side corresponds to the ratio of the total length spent on that edge for *any* reduced cycle on the graph, to the total length of all reduced cycles. This relationship will be more formally proven in a later part of this paper.

Example 1. *Consider a 2-petal rose, with petal a having length 1, and petal b having length 2. For $L = 1$, there are only two reduced cycles, a and a', meaning all expenditure is on a. For $L = 2$, we have $a, a', aa, a'a', b, b'$. Thus 4 of 6 cycles spend all length on a, whereas the other 2 spend all length on b, making the expenditure on a equal to 2/3.*

For $L = 3$, 8 of the new (length 3) reduced cycles involve some combination of one a and one b, and two involve 3 a's. There are 16 total reduced cycles now (6+8+2), two of which do not cross a at all, so we compute the expenditure on a to be $(6 + \frac{8}{3})/16$ or 25/48. Using the same process for $L = 4$, we find 2 each of cycles with four a's and two b's, and a total of 16 cycles that involve one b and two a's. The total length of all length-4 reduced cycles is thus 80, half of which is spent on a. Consequently, the expenditure for $L = 4$ is 65/128.

Now considering the right-hand side of the equation: intuitively, one would expect that about twice as many traversals of a occur than of b in any random reduced cycle on the graph, as b is twice the length of a. So choosing $\mu_a = 2/3$ and $\mu_b = 1/3$, $\frac{\mu_a \ell(a)}{\mu_a \ell(a) + \mu_b \ell(b)} = 1/2$. This would make sense, as the sequence $\{1, 2/3, 25/48, 65/128, \dots\}$ appears to approach 1/2.

Formally, we can determine μ using the process that follows; but to do so, we need a representation of G as a mathematical object that we can manipulate algebraically: a matrix. We define A_G, our edge-adjacency-matrix representation for G, to be a $2|E| \times 2|E|$ matrix ($2|E|$, not simply $|E|$, due to there being two directions for each edge) such that

$$A_G(e, f) = \begin{cases} \exp(-h(G) \cdot \ell(e)) & \text{if } \exists \text{ reduced cycle } \gamma \text{ s.t. } ef \in \gamma; \\ 0 & \text{otherwise.} \end{cases}$$

In contrast to a traditional "adjacency matrix" for a graph, which has nonzero entries (typically ones) wherever there exists a connecting edge between two vertices, this matrix has nonzero entries (which are weighted by graph entropy and edge length) wherever there exists a vertex connecting two edges, or more specifically wherever a reduced cycle exists that goes through the two edges in order.

For instance, a 2-petal rose of entropy 1 and with edges e and f, where $\ell(e) = x$ and $\ell(f) = y$, has the following matrix, with rows and columns following the order $\{e, f, e', f'\}$:

$$A_G = \begin{bmatrix} e^{-x} & e^{-x} & 0 & e^{-x} \\ e^{-y} & e^{-y} & e^{-y} & 0 \\ 0 & e^{-x} & e^{-x} & e^{-x} \\ e^{-y} & 0 & e^{-y} & e^{-y} \end{bmatrix}.$$

Observe that there are 0's corresponding to the edge-sequences $e'e$, $f'f$, ee', and ff', since these pairs violate the no-backtracking rule.

Theorem 7. *[**18**, 2] A_G has all eigenvalues with modulus ≤ 1, with 1 as an eigenvalue. Also, the 1-eigenspace is 1-dimensional, and the right eigenvector u, for which $Au = u$, has all positive components as a result of Perron-Frobenius theory. Additionally, there exists a left eigenvector v s.t. $vA = v$.*

Proof. (Regarding the existence of the left eigenvector.) The matrix A^T has exactly the same properties as A, and so we can assume that there exists a right eigenvector v^T for which $A^T v^T = v^T$. Transposing both sides yields $vA = v$. □

Using this, the components of μ are computed via:

$$\mu_{e_i} = \frac{u_{e_i} v_{e_i}^T}{\sum_e u_e v_e^T}. \tag{4}$$

Now to understand the existence and nature of the eigenvectors u and v as described above, we involve **Perron-Frobenius theory**, a set of tools in linear algebra. Defining a **positive matrix** to be one with entries all greater than zero, the Perron-Frobenius theorem says that if A is a matrix satisfying the property that it has some power k so that A^k is positive, then:

- A has a positive real eigenvalue that has the largest modulus among eigenvalues. Recall that in general, eigenvalues of A can be complex: thus we are asserting that the eigenvalue of A furthest from the origin in the complex plane happens to be real.
- This largest eigenvalue of A has a corresponding eigenspace which is 1-dimensional. In other words, there is only one unit eigenvector for this eigenvalue. Call this eigenvector u.
- In fact, u has all positive entries. Additionally, if x is any vector, then $(1/n) \cdot A^n(x)$ converges to a scalar multiple of u.

Note that positivity (and nonnegativity, i.e. pertaining to matrices whose entries are all nonnegative) is preserved with matrix powers: that is, if $A > 0$, then $A^2 > 0$, $A^3 > 0$, and so on. Further, the exponential growth rate of A^k as $k \to \infty$ is controlled by the eigenvalue of A with the largest modulus. This is stated as the following theorem.

Theorem 8. *[**18**, 2] If A is an $n \times n$ matrix and A has an eigenvalue λ that dominates all other eigenvalues, i.e. $|\lambda| > |\lambda_i|$ for all i., then*

$$\lim_{k\to\infty} A^k/\lambda^k = uv^T, \tag{5}$$

where u is the right eigenvector for λ, v is the left, and the left and right eigenvectors are normalized so that $v^T u = 1$.

Proof. [**22**, 2] Combining the column eigenvectors of A into a matrix T, we know that there exists an upper triangular matrix J (called the Jordan canonical form of A) such that $A = TJT^{-1}$, where the diagonal entries of J are the eigenvalues of A, the entries just above the diagonal are 1's, and all other entries are 0's. Normalizing the eigenvectors such that $v^T u = 1$ (i.e., $\max|\lambda_i| = 1$), taking $A^k = TJ^kT^{-1}$ and computing the limit as $k \to \infty$ results in all non-dominant (less than 1) eigenvalues in J going to zero in J^k, as a result of the fact that $\lim_{k\to\infty} a^k = 0$ for any a s.t. $|a| < 1$. Consequently, we are left with only the dominant eigenvectors u and v^T when we multiply out TJ^kT^{-1}, yielding uv^T. □

The notion of a matrix having a "dominating eigenvalue" is easiest to understand in the case of a 2 by 2 matrix. Setting up the equation

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \lambda \begin{bmatrix} x_1 \\ x_2 \end{bmatrix},$$

with a, b, c, d all nonnegative, we have that $\det(A - \lambda I) = 0$, i.e. $(a-\lambda)(d-\lambda) - bc = 0$ or $\lambda^2 + (a+d)\lambda + ad - bc = 0$. Applying the quadratic formula,

$$\lambda = \frac{a + d \pm \sqrt{(a+d)^2 - 4(ad-bc)}}{2} = \frac{a + d \pm \sqrt{(a-d)^2 + 4bc}}{2}.$$

Notice that in the case where we take $\pm$ to be $-$, the result is generally much closer to zero than when we have $\pm$ as $+$. In fact, when a, b, c, d are roughly the same, we find that

$$\lambda_1 \approx \frac{2a + \sqrt{4a^2}}{2} = 2a, \quad \lambda_2 \approx \frac{2a - \sqrt{4a^2}}{2} = 0.$$

Clearly, λ_1 is the dominating eigenvalue here. Intuitively, we can see that in general, we will always have eigenvalues close to zero when the entries of A are relatively close together in size. The "dominating eigenvalue" is simply the unique one that does not conform to this trend.

In the case of our edge-adjacency matrix, which has all nonnegative entries, the Perron-Frobenius theorem would (if it were applicable to A_G) allow us to guarantee the existence of this eigenvalue, and consequently the corresponding left and right eigenvalues u and v— which we use to compute μ, as illustrated by equation (4).

But notice that in order for us to apply Perron-Frobenius to A_G, we must confirm that some power k of A_G is indeed positive (has all positive entries). To do so, we consider the directed graph of A_G.

Definition 5. *The **directed graph** D_G of an $n \times n$ matrix has n directed edges, which indicate whether we can traverse from one edge to another adjacent edge, based on whether the entries in the matrix are nonzero. Nonzero entries indicate edges; 0-entries indicate no edges.*

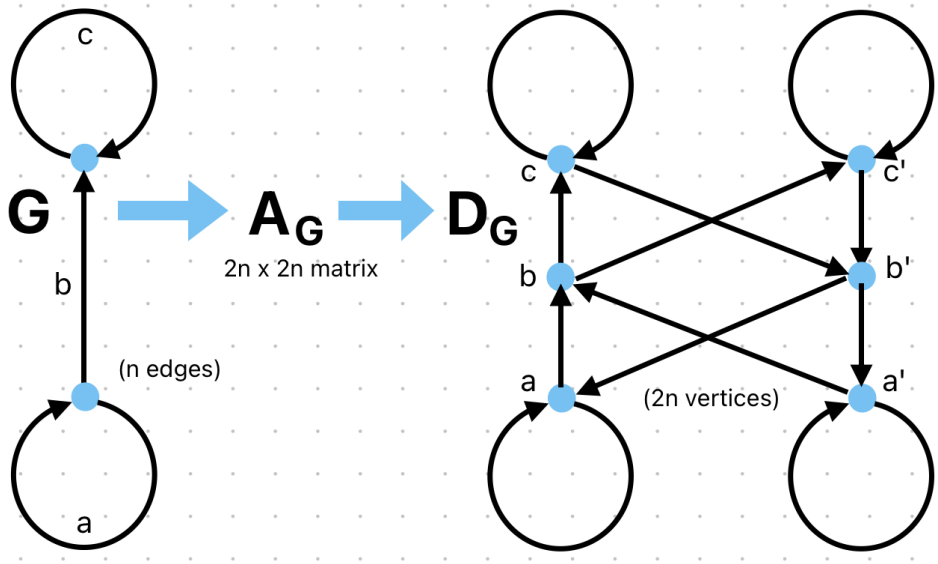

Figure 4: directed graph for a barbell graph

In other words, the directed graph D_G that corresponds to the adjacency matrix A_G for the original graph G has vertices which correspond to the $2|E|$ edges of G (so e.g. if G has 4 edges in total, then D_G has 8 vertices), and

directed edges that convey edge traversals [**19**, 1]. For example, the directed graph for a barbell graph is shown below: while edge transitions like cb', ab, and $a'b$ are valid (as indicated by D_G), ac, cb, and ba are examples of ones that are not (hence the lack of corresponding edges in D_G).

Observe, crucially, that D_G has the property that we can get to any one of its vertices from any other via some sequence of edges— which we accredit to the fact that G is both connected, and has vertices with valencies all greater than or equal to 3. In other words, any two edges e and f of the original graph G can both be found in some cycle. Consequently, there exists some power of A_G, the adjacency matrix of D_G, whose entry in the (e, f) position is nonzero, as we can after some number of intermediate edges arrive at edge f from edge e. (For example, if e and f are 3 edges apart, the (e, f) entry would be nonzero in A_G^3.) It is in this way that we are able to apply Perron-Frobenius to A_G.

To illustrate the conceptual significance of the Perron-Frobenius theorem as it pertains to dynamical systems such as the distribution of reduced cycles in our metric graphs, consider the example of two cities A and B, between which the proportions of each population that migrate to the other city (or remain at the same one) is conveyed in figure 5— α being the proportion of the population that move to city B from A, and β being the proportion of the population that move to city A from B. (This is somewhat parallel to our understanding of reduced cycles on metric graphs, simulating i.e. the proportions of all reduced cycles that land at a particular vertex and proceed to one edge, as opposed to another.)

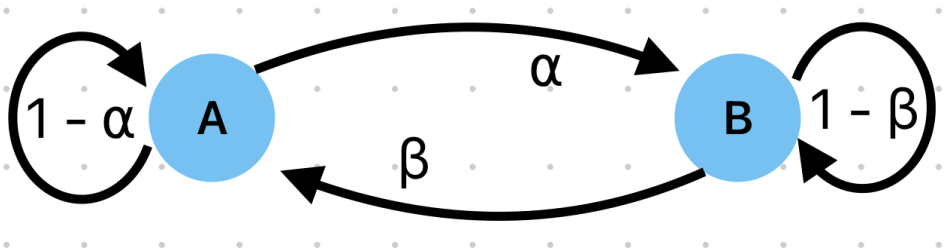

Figure 5: annual migration rates between two cities.

Letting x_n and y_n represent the populations in year n at cities A and B respectively, and defining vector v_n to be $\begin{bmatrix} x_n \\ y_n \end{bmatrix}$, we can model the migration process using the following equation:

$$\begin{bmatrix} x_{n+1} \\ y_{n+1} \end{bmatrix} = \begin{bmatrix} 1-\alpha & \beta \\ \alpha & 1-\beta \end{bmatrix} \begin{bmatrix} x_n \\ y_n \end{bmatrix}$$

or equivalently,

$$v_{n+1} = Av_n. \tag{6}$$

One might ask: for what ratio of populations will the amount of people residing in each city remain constant on a year-to-year basis? And the answer— i.e. an "equilibrium state" of migration, known more generally as the *steady state distribution* for this system, is a pair of values x, y given α, β that satisfy $v = v_n = v_{n+1}$, or $Av = v$. Such a v is exists whenever the matrix A is irreducible: but in fact, A is irreducible if and only if its corresponding graph is strongly connected [**6**, 2]— which in the case of D_G and this system, it is.

Observe that solving for such a vector v, given predetermined values of α and β, is equivalent to determining the right eigenvector of A with eigenvalue 1. Thus, we can begin to get a sense of why the equilibrium or "Gibbs" measure of our given metric graph is directly related to the eigenvectors of our adjacency matrix A.

Furthermore, taking powers of A, i.e. A^k, informs us what will happen in the long term— and as equation (5) suggests, this long-term behavior can be alternatively analyzed using the dominating eigenvectors of A, thus providing some more intuition for equation (4).

In the generalized case $Av = \lambda v$, solving for eigenvectors via the equation $\det(A - \lambda I) = 0$ gives, with some expanding and regrouping of terms,

$$\begin{aligned} (1-\alpha-\lambda)(1-\beta-\lambda) - \alpha\beta &= 0 \\ \lambda^2 + (\alpha+\beta-2)\lambda + 1 - \alpha - \beta &= 0 \\ (1-\lambda)(\lambda+\alpha+\beta-1) &= 0. \end{aligned}$$

Thus we find the eigenvalues to be $\lambda_1 = 1$ and $\lambda_2 = 1 - \alpha - \beta$— and clearly $\lambda_2 < 1$, since $\alpha, \beta > 0$. But the larger eigenvalue λ_1 is more meaningful here in the immediate (year-by-year) sense, since this is what provides us with v such that $Av = v$.

Notice, additionally, that A is *column-stochastic*, meaning each of its columns sum to 1. This property of such matrices (which contributes the fact that A has an eigenvalue of 1) will be useful in determining exactly which sets of metrics (edge lengths) on our graphs will result in an entropy of 1.

Consider, for instance, a two petal rose R_2 with petals both of length 1: then each nonzero entry of A_{R_2} is $e^{-\log(3)} = 1/3$, and so

$$A_{R_2} = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & 0 & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \\ 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & 0 & \frac{1}{3} & \frac{1}{3} \end{bmatrix}.$$

$\det(A_{R_2} - \lambda I)$ is

$$\begin{vmatrix} \frac{1}{3}-\lambda & \frac{1}{3} & 0 & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3}-\lambda & \frac{1}{3} & 0 \\ 0 & \frac{1}{3} & \frac{1}{3}-\lambda & \frac{1}{3} \\ \frac{1}{3} & 0 & \frac{1}{3} & \frac{1}{3}-\lambda \end{vmatrix} = \lambda^4 - \frac{4\lambda^3}{3} + \frac{2\lambda^2}{9} + \frac{4\lambda}{27} - \frac{1}{81},$$

and solving for λ in $\lambda^4 - \frac{4\lambda^3}{3} + \frac{2\lambda^2}{9} + \frac{4\lambda}{27} - \frac{1}{81} = 0$ yields eigenvalues $\lambda = -\frac{1}{3}, \frac{1}{3}, 1$. Choosing $\lambda = 1$, the dominating left and right eigenvectors are $v = (1,1,1,1)^T$ and $u = (1,1,1,1)$, respectively. Therefore,

$$\mu_{e_i} = \frac{1 \cdot 1}{1 \cdot 1 + 1 \cdot 1 + 1 \cdot 1 + 1 \cdot 1} = \frac{1}{4}$$

for all edges e_i, and so $\mu = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$. This is precisely what we would expect it to be, as all four edge directions (counting each petal once forward and once in reverse) are equally likely to be used in any given reduced cycle.

4 The Entropy Gradient

We now introduce an entirely different way to think of μ. First, we note that since the function h is an infinitely differentiable function, we can consider its gradient vector ∇h, which is never 0 on the set $X \in \mathbb{R}^{|E|}$ so that $h(X) = 1$. This set $X = h^{-1}(1)$, i.e. the set of graphs whose entropy is 1, is an $(|E|-1)$-dimensional sub-manifold of $\mathbb{R}^{|E|}$ [3].

Definition 6. *An **m-manifold** in $\mathbb{R}^n$ ($n \geq m$) is a subset M of $\mathbb{R}^n$: it has the property that there exists a collection $\{B_\alpha\}_{\alpha \in A}$ of open balls in $\mathbb{R}^n$ s.t. (1) $M \subseteq \cup_{\alpha \in A} B_\alpha$, and (2) $M \cap B_\alpha$ is the graph of a smooth function $F_\alpha : U \to \mathbb{R}^n$, where $U \subseteq \mathbb{R}^{n-1}$ is an open set.*

We additionally make the distinction that if M and N are both manifolds in $\mathbb{R}^n$ and $N \subset M$, we call N a **sub-manifold** of M. In particular, M is a sub-manifold of $\mathbb{R}^n$.

This perspective is significant because it allows us to consider the normal vector to the level set of $h^{-1}(1)$, which is, in fact, equal to the gradient vector of the function h that corresponds to entropy. The number of inputs to said function h, which spits out the entropy of the graph itself, is equal to the number of edges of the graph. The vector containing the partial derivatives of this function, i.e. the gradient vector, is essentially the same as this Gibbs measure, in the sense that they are in fact proportional. We should now like to show that this relationship between the Gibbs measure and the gradient vector for the entropy function does indeed hold, i.e. $\mu \propto \nabla h$.

To be more specific,

Theorem 9. $\frac{\nabla h}{h} = \frac{-\mu}{\sum_e \mu_e \lambda_e}$ *[**35**, 7].*

Corollary 2. $\mu \propto \nabla h$*, i.e. there exists a scalar* k *for which* $\mu = k\nabla h$*.*

Proof. [**38**, 7] We want to show that $k = \frac{-\sum_e \mu_e \lambda_e}{h}$.

To begin, fix an edge e; we'll show that $\frac{(\frac{\partial h}{\partial e})}{h} = -\frac{\mu_e}{\langle \mu, e \rangle}$. Consider the "path of metrics" to be defined as $\rho_t = (l_{e_0}, l_{e_1}, ..., l_e + \Delta t, ..., l_{e_n})$: then we have $\frac{d}{dt} \mid_{t=0} h(\rho_t) = \frac{\partial h}{\partial e} \mid_{\rho_0}$, where we let Δt denote a small value of t.

Originally, a generic reduced cycle γ spends a fraction of $\frac{\mu_e l_e}{\langle \mu, l \rangle}$ of its length on edge e; equivalently, if we choose a generic length L, we spend $\frac{\mu_e l_e}{\langle \mu, l \rangle} L$ of the length on e. Now when we add Δt to the length of e, a generic reduced cycle spends an additional $\frac{\mu_e l_e \Delta t}{\langle \mu, l \rangle} L$ on e; so the new length is $L(1 + \frac{\mu_e l_e \Delta t}{\langle \mu, l \rangle})$. As a result,

$$\frac{h + \Delta h}{h} = \frac{1}{1 + \frac{\mu_e \Delta t}{\langle \mu, l \rangle}}.$$

So $\frac{h}{h + \Delta h} = \frac{\langle \mu, l \rangle + \mu_e \Delta t}{\langle \mu, l \rangle}$, which becomes

$$h\langle \mu, l \rangle = h\langle \mu, l \rangle + h\mu_e \Delta t + \Delta h \langle \mu, l \rangle + \Delta h \mu_e \Delta t.$$

Factoring, $\frac{-h\mu_e}{\langle \mu, l \rangle + \mu_e \Delta t} = \frac{\Delta h}{\Delta t}$. Now taking the limit as Δt goes to 0, we have that

$$\frac{\partial h}{\partial e} \mid_{t=0} = -\frac{h\mu_e}{\langle \mu, l \rangle}.$$

Dividing both sides by h gives us the desired result. □

Now with this in mind, we would like to determine *why* there is this limiting distribution vector μ that encodes the life of a random reduced cycle on the graph, and how the connection between the linear algebra, calculus,

and dynamics approaches to reduced cycles provides an explanation for the existence of this vector. In particular, we would like to prove equation (3).

In summary of the perspectives we have, the Gibbs measure is as follows:

(A) It is the gradient vector for the entropy function, ∇h;

(B) It encodes the relative frequencies with which a generically chosen cycle traverses each edge, as suggested by equation (3);

(C) It is the vector one obtains from multiplying the left and right Perron-Frobenius eigenvectors together of the transition matrix.

We showed that $(B) \Rightarrow (A)$. $(C) \Rightarrow (B)$ should follow from the Perron-Frobenius theorem: to proceed with the proof of why it does, we apply the Perron-Frobenius theorem to the matrix $A_G = 2|E| \times 2|E|$, as defined. This will yield that $\lim A^k = uv^T$. Then, we relate μ_e to the $(e,e)^{th}$ entry of uv^T.

Thus, begin by recalling theorem 8:

$$\lim_{k\to\infty} A^k/\lambda^k = uv^T.$$

But since $\lambda = 1$, we have that $\lim A^k = uv^T$. Notice, in particular, that the $(e,e)^{th}$ entry of A^k tends to $u_e v_e^T$. Now by equation (5),

$$\mu_{e_i} = \frac{u_{e_i} v_{e_i}^T}{\sum_e u_e v_e^T},$$

and so the diagonal entries of uv^T are proportional to the entries of μ_e.

To interpret the (e,f) entry of A^k: recall that $A(e,f)$ is nonzero when there is a cycle through edges e, f. As a result, we can interpret $A^k(e,f)$ as

$$\sum_{\Gamma} = \exp(-h(G) \cdot [\ell(e) + \ell(e_2) + \cdots + \ell(e_k)]),$$

with Γ being all paths through $\{e, e_2, \ldots, e_{k-1}, e_k\}$, such that the terminal vertex of e_k is the initial vertex of f— thus yielding a cycle.

Now as $L \to \infty$, the proportion of cycles in $N_G(L)$ that go over every edge goes to 1. Furthermore, for k very large, nearly every summand in the sum above corresponds to a cycle that goes over every edge. Thus the magnitude of this sum is dictated by the number of summands, not the value of the summands themselves.

And since the number of times a given summand appears in the sum equals the number of times a cycle goes through the edge in question, the size of the $(e,e)^{th}$ entry of uv^T is proportional to the number of times a cycle goes over the edge e. Hence, $(C) \Rightarrow (B)$.

5 Subgraphs on Roses

From here, we proceed in one of three ways:

1. [Akanda] Fix n at least 3, and define the *loop n-gon* to be the graph obtained by starting with an n-sided polygon and then adding in a self-loop at each vertex. Consider a family of unit entropy metrics on the loop n-gon for which the edges of the n-gon go to infinity in length and the edges of each self-loop are the same and are all shrinking to 0. Measure how the largest entropy of a proper subgraph changes along this family.

2. [Shigeno] Fix n at least 3, and consider R_n, the rose on n petals. Do some experiments on the largest entropy of a proper sub-rose by picking a bunch of different entropy-1 metrics on computing the entropy's of various subgraphs. It would be particularly interesting to determine whether we can find a metric where the largest entropy sub-rose has entropy less than $\log(2n-3)/\log(2n-1)$.

3. [LaRochelle] Consider a graph composed of one non-loop edge e, one loop edge on one side of e, and two loop edges on the other side of e, called a *lamppost*. We can choose an entropy-1 metric on this graph by starting with a barbell that has entropy very nearly 1 and for which e is extremely long and two of the self-loops (one on either side of e) are both short. Then the length of the remaining self-loop has to be very long. Alternatively, we could start by picking a metric on a figure 8 where the entropy is very nearly 1, and then extending it to an entropy-1 metric on the whole graph where the length of e is very long. For the first type of metric, the largest entropy subgraph will be a barbell; in the second type, it will be a figure 8. Try and nail down the "critical metrics" at which the types switch from one to the other.

To begin with question 2, we illustrate how one could go about restricting the metrics of a graph, given its entropy; in particular, we wish to examine entropy-1 roses.

Returning to the example of a 2-petal rose with entropy 1 and petal lengths x and y, recall that columns of A_G either contain one 0, one e^{-x}, and two e^{-y}'s, or one 0, two e^{-x}'s, and one e^{-y}. Using the fact, then, that the largest eigenvalue of A_G is 1, setting $\det(A-I)=0$ yields with some

manipulation and rearranging of terms that that

$$y = \log\left(\frac{e^x + 3}{e^x - 1}\right).$$

So for example, $x = \log(3)$ yields $y = \log(3)$, and $x = \log(5)$ results in $y = \log(2)$. Using a similar process for a 3-petal rose with lengths y, y, x and $h = 1$, we must have that

$$x = \log\left(\frac{e^y + 5}{e^y - 3}\right),$$

where $y > \log(3)$, since entropy of 1 is achieved already in two petals (third petal not needed). For example, when $y = \log(4)$, $x = \log(9)$. And more generally for 3-petal rose with lengths x, y, z, we let $x = \log(5)$: then

$$y = \log\left(\frac{2e^z + 5}{e^z - 2}\right).$$

Observe also that for an n-petal rose with all lengths equal and $h = 1$, $l(e_i) = \log(2n-1)$ for all i: so removing 1 petal from this n-petal rose yields subgraph entropy

$$\frac{\log(2n-3)}{\log(2n-1)},$$

which approaches 1 from the left as $n \to \infty$. For example, for $n = 3$, removing any petal gives $h_{\text{subrose}} = \frac{\log(3)}{\log(5)} \approx 0.683$.

Now let us take this a step further, and determine the exact length of the third petal of a 3-petal rose with entropy 1 in terms of the lengths of the other two petals. Let $x = \log(i)$ and $y = \log(j)$; solving for z yields

$$z = \log\left(\frac{5 + 3i + 3j + ij}{-3 - i - j + ij}\right).$$

Notice that the constraints here are that $-3 - i - j + ij > 0$ or $j > \frac{i+3}{i-1}$, and so $i > 1$.

Conjecture 1. *When all petals are of equal length, this yields the minimum value of largest-entropy sub-rose. In other words,*

Let $\mathcal{M}^1_{R_3} \subseteq \mathbb{R}^3$ be the set consisting of all entropy 1 metrics on 3-petal rose: so for instance, $(\log 5, \log 5, \log 5), (\log 4, \log 4, \log 9), (\log 3, \log 5, \log 11) \in \mathcal{M}^1_{R_3}$. *Define $S : \mathcal{M}^1_{R_3} \subseteq \mathbb{R}^3$ s.t.*

$$S(x, y, z) = \max\{h(x, y), h(x, z), h(y, z)\}.$$

We conjecture that:

$$\min\{S(x,y,z) \mid (x,y,z) \in \mathcal{M}^1_{R_3}\} = h(\log 5, \log 5).$$

To gather evidence for the truth of this conjecture, we pick a wide range of values $(x = \log(i), y = \log(j), z)$ such that $i > 1, j > \frac{i+3}{i-1}$, and $z = \log\left(\frac{5+3i+3j+ij}{-3-i-j+ij}\right)$, as desired. Then,

1. Compute $S(x,y,z)$ for each of the tuples.
2. Plot points (i,j,S).

We want to show that $h(\log 5, \log 5)$ is the global minimum of this graph.

Doing so results in the graph below, computed via Mathematica. We see that the minimum is a achieved at $h(\log 5, \log 5) = \frac{\log(3)}{\log(5)}$, as expected. An in general: we predict that the minimum largest subrose entropy of R_n is exactly $\frac{\log(2n-3)}{\log(2n-1)}$.

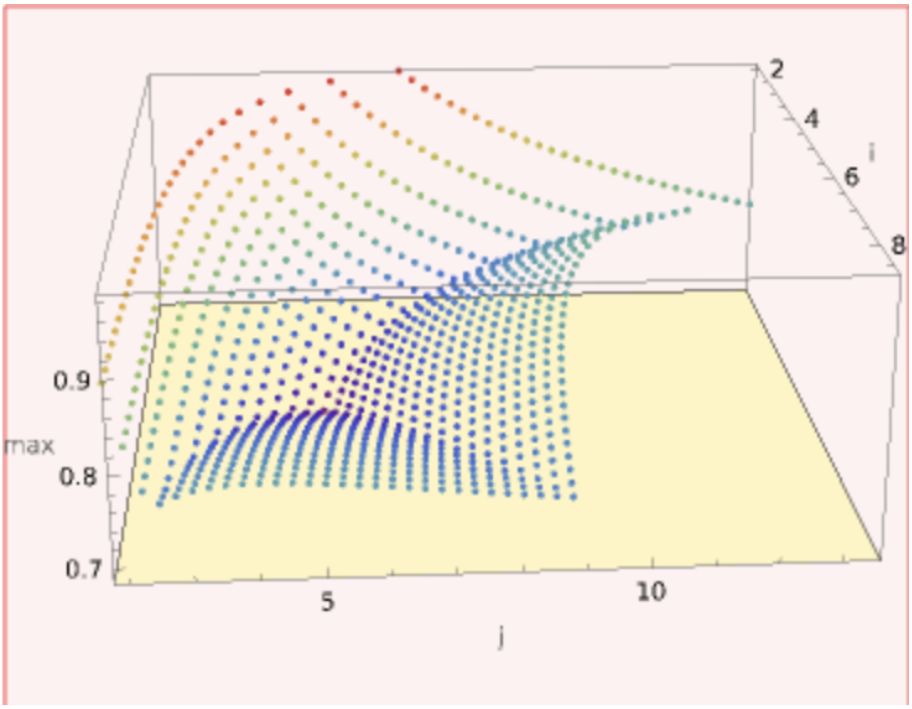

Figure 6: scatterplot of $S(x,y,z)$ in its domain.

Now consider our method of solving for entropy given petal lengths, which is quite tedious due to our unknown variable h being tucked into the exponents of a sum of quantities. E.g. given a 2-petal rose with lengths 1 and 2, we want to compute the entropy. Nonzero entries are e^{-h} and e^{-2h}, so $\det(A - I) = 0 = 1 - 3e^{-6h} + 2e^{-5h} + 3e^{-4h} - e^{-2h} - 2e^{-h}$. From here we obtain that $h \approx 0.7563$.

We conclude this paper by detailing a somewhat more straightforward method to compute the entropy of a graph, i.e. one does not involve solving convoluted exponential equations. Define the **scale-transition matrix** E_G for graph G as

$$E_G(e, f) = \begin{cases} \exp(-\ell(e)) & \text{if } \exists \text{ reduced cycle } \gamma \text{ s.t. } ef \in \gamma; \\ 0 & \text{otherwise.} \end{cases}$$

Additionally, let

$$Q_G(e, f) = \frac{v_f}{\lambda v_e} E_G(e, f),$$

where v is the right Perron-Frobenius eigenvector of E_G associated with λ. Then let $u \in \mathbb{R}^{|E|}$ be the positive left eigenvector with eigenvalue 1 and $\sum_{e \in E} u_e = 1$: the **measure-theoretic entropy** of G is given by

$$h_{meas} = -\sum_{e \in E} u_e Q_G(e, f) \log Q_G(e, f),$$

where if $Q_G(e, f) = 0$ for some $e, f \in E$ then we take the entire term to be 0. Finally, we can compute entropy h via the function below.

$$h(G) = \frac{h_{meas}}{\sum_{e \in E} u_e \ell(e)}. \tag{7}$$

[1][2][3][4][5][6][7]

References

[1] Tarik Aougab, Matt Clay, and Yo'av Rieck. *Thermodynamic metrics on outer space*. 2020. arXiv: 2009.13314 [math.GT].

[2] Karl Eriksson. *Perron-Frobenius' Theory and Applications*. Department of Mathematics, Linköping University, 2023.

[3] Victor Guillemin and Alan Pollack. *Differential Topology*. Prentice-Hall, 1974.

[4] Allen Hatcher. *Algebraic topology*. Cambridge: Cambridge University Press, 2002, pp. xii+544.

[5] Lien-Yung Kao. "Pressure type metrics on spaces of metric graphs". In: *Geometriae Dedicata* 187.1 (Sept. 2016), 151–177. DOI: 10.1007/s10711-016-0194-9.

[6] Seonhee Lim. *Minimal volume entropy on graphs*. 2005. arXiv: math/0506216 [math.GR].

[7] Curtis T. McMullen. "Entropy and the Clique Polynomial". In: *Journal of Topology* (2013). DOI: 10.1112/jtopol/jtu022.