

MAGNITUDE, HOMOLOGY, AND THE WHITNEY TWIST

EMILY ROFF

(communicated by Peter Bubenik)

Abstract

Magnitude is a numerical invariant of metric spaces and graphs, analogous, in a precise sense, to Euler characteristic. Magnitude homology is an algebraic invariant constructed to categorify magnitude. Among the important features of the magnitude of graphs is its behaviour with respect to an operation known as the Whitney twist. We give a homological account of magnitude's invariance under Whitney twists, extending the previously known result to encompass a substantially wider class of gluings. As well as providing a new tool for the computation of magnitudes, this is the first new theorem about magnitude to be proved using magnitude homology.

1. Introduction

Magnitude is an isometric invariant of metric spaces, so-named for its web of connections to 'size-like' quantities of significance in various corners of mathematics [14]. The magnitude of a space X is a function

$$\text{Mag}(X)(-): [0, \infty) \rightarrow \mathbb{R}$$

whose parameter can be thought of as controlling the scale of the metric in X . For any given choice of the parameter, magnitude behaves in many ways like the cardinality of finite sets: it is multiplicative with respect to ℓ_1 -products, additive with respect to disjoint unions, and satisfies an inclusion-exclusion formula with respect to a more general class of unions. Yet, unlike cardinality, magnitude is sensitive to the distances in a space.

The three-point space X in Figure 1 is indicative of the general situation for finite spaces. When the scale parameter t is very small, the magnitude function is close to 1, as if recognizing that 'from far away' X resembles a one-point space. As t increases, it becomes possible to distinguish the point on the right, and $\text{Mag}(X)(t)$ lingers close to 2. Once t is sufficiently large, all three points in the space are easily distinguished and $\text{Mag}(X)(t)$ approaches 3. Accordingly, the magnitude function is often referred to as counting 'the effective number of points' in a finite space as the scale of the metric is varied.

Received November 4, 2022, revised March 26, 2023; published on February 21, 2024.

2020 Mathematics Subject Classification: 55N35, 05C31, 05C76.

Key words and phrases: Magnitude, graph, Whitney twist.

Article available at <http://dx.doi.org/10.4310/HHA.2024.v26.n1.a7>

Copyright © 2024, Emily Roff. Permission to copy for private use granted.




Figure 1: The magnitude function of a three-point space.
(Example 6.4.6 in Leinster [16], originally due to Willerton.)

Magnitude can be defined for many non-finite compact metric spaces, too, and in this context it has been shown to be related to a remarkable variety of more classical size-related quantities. Under suitable conditions, the magnitude function records information about the intrinsic volumes of integral geometry [3, 5, 7, 22, 24], curvature and the Willmore energy [6], Minkowski dimension [23] and even a family of indices used by biologists to quantify the diversity of ecological communities [17, 18, 20]. (These references are far from exhaustive.)

In this paper, though, we will be concerned specifically with the magnitude of graphs. Every undirected finite graph can be regarded as a metric space in which the distance between any pair of vertices is the number of edges in a shortest edge-path joining them. The magnitude of a graph is defined to be its magnitude with respect to this metric.

In this context, magnitude is most naturally interpreted as a formal power series with integer coefficients. Leinster [15] gives a combinatorial formula for the coefficients in this series, revealing them to count—with appropriate signs—the paths of given length in a graph. (This terminology will be explained in Section 2.) Observing that the alternating sum computing each coefficient resembles the Euler characteristic of a chain complex, Hepworth and Willerton in [10] construct a bigraded homology theory for graphs which ‘categorifies’ their magnitude. That theory, *magnitude homology*, has since been extended by Leinster and Shulman [21] to encompass all metric spaces. In what follows we denote the magnitude homology functor by MH_{\bullet}^* , with \bullet and $*$ standing for the two distinct gradings; its construction is described in Section 2.

The magnitude homology of metric spaces has been well studied in recent years, and has been shown to capture subtle information about the convexity of a space [11] and the existence and uniqueness of geodesics [1, 8]. Even in the restricted setting of graphs it has proved to be a rich invariant. Particular attention has been paid to a class of graphs which are termed *diagonal*: graphs X with the property that $MH_k^\ell(X) = 0$ whenever $k \neq \ell$ [10, 9]. Most recently, Asao [2] has established a close relationship between magnitude homology and Grigor’yan–Muranov–Yau *path homology*, exploiting this relationship to prove that every diagonal graph has trivial reduced path homology.

Meanwhile, various results about the formal cardinality-like properties of magnitude have been recovered as the numerical ‘shadows’ of facts about magnitude homology. In particular, Bottinelli and Kaiser [4] (following Hepworth and Willerton [10]) have proved a Künneth-type theorem and a Mayer–Vietoris theorem for the magnitude homology of metric spaces—categorifying, respectively, the multiplicativity of magnitude and the inclusion–exclusion formula. This paper takes a similar form, offering a homological account of magnitude’s behaviour with respect to an operation on graphs known as the Whitney twist.

MAGNITUDE AND THE WHITNEY TWIST In Leinster’s paper [15] initiating the study of magnitude for graphs, there are three central results concerning magnitude’s behaviour under combinations such as products or gluings. The first of these, Lemma 3.6, says that magnitude is multiplicative with respect to the operation that graph theorists call the *cartesian product*, \square . That is, for any graphs X and Y ,

$$\text{Mag}(X \square Y) = \text{Mag}(X) \cdot \text{Mag}(Y). \quad (1)$$

The second such result, Theorem 4.9 in [15], gives an inclusion–exclusion formula for the magnitude of a union of graphs, subject to conditions on the subgraphs involved. Each of these results is a specialization to graphs of a corresponding property of metric space magnitude.

The third combination result in [15] is different. It concerns an operation known as the *Whitney twist*, which can be described as follows. Given two graphs G and H , each with two (distinct) distinguished vertices, there are two ways to form a new graph: we can glue the distinguished vertices together one way round, or the other. If X and Y are the graphs thus formed, one says that X and Y differ by a Whitney twist. Leinster proves that magnitude is preserved under Whitney twists—*provided* the gluing vertices happen to be adjacent.

Theorem 1.1 ([15], Theorem 5.2) *Let X and Y be graphs differing by a Whitney twist, and suppose that the two gluing vertices are adjacent in X (or equivalently in Y). Then*

$$\text{Mag}(X) = \text{Mag}(Y).$$

It can be seen from Example 4.16 in the same paper that magnitude need *not* be preserved under Whitney twists in which the gluing vertices are not adjacent. From this it follows that the magnitude of a graph—unlike many other graph invariants—is not determined by the cycle matroid or the Tutte polynomial.

While multiplicativity and the inclusion–exclusion formula both lend weight to the interpretation of magnitude as a ‘size-like’ invariant, the fact that magnitude is preserved under certain Whitney twists but not others is intriguing, and thus far poorly understood. In particular, it is not easy to extract from the proof of Leinster’s Theorem 5.2 a clear understanding of what exactly it is about the adjacency of the two gluing vertices that guarantees the result, leading one to wonder whether the theorem might be merely an instance of a more general statement about gluings. In this paper we prove that, in fact, it is.

Moreover, whereas the multiplicativity of magnitude and the inclusion–exclusion formula have been categorified, respectively, by the Künneth theorem and the Mayer–Vietoris sequence for magnitude homology, magnitude’s behaviour under Whitney

twists has until now resisted a homological explanation. Indeed, the behaviour of magnitude homology under Whitney twists is listed as an open question in [10].

That question is not settled here; instead we prove, by homological means, an extension of Theorem 1.1 to a wider class of ‘twisted’ graph gluings which we term *sycamore twists* (Definition 3.3). Our main theorem runs as follows.

Theorem 6.5 *Let X and Y be graphs which differ by a sycamore twist. Then*

$$\text{Mag}(X) = \text{Mag}(Y).$$

As well as providing a new tool for computing magnitudes of gluings, Theorem 6.5 can be regarded as clarifying the picture of graph magnitude in two senses—and in doing so it opens up new questions about magnitude homology.

First, the definition of a sycamore twist distills the important features of a Whitney twist with adjacent gluing vertices and elucidates the role played by this adjacency. Second, the proof of the theorem makes it clear that the behaviour of magnitude under twists is intimately related to the excision theorem and the Mayer–Vietoris sequence for magnitude homology and thereby to the inclusion-exclusion formula for magnitude.

Indeed, the result opens up the possibility that the excision and Mayer–Vietoris theorems might yet be strengthened, implying a more general inclusion-exclusion principle for magnitude. This would afford efficiency in computations, but it could also prove valuable for theoretical reasons. As things stand, the excision theorem is not strong enough to provide an axiomatization of magnitude homology for any class of graphs or metric spaces large enough to be interesting. A stronger excision theorem might allow for an axiomatic theory of magnitude homology and thereby of magnitude for finite spaces.

Finally, while various authors have obtained categorifications of known results about magnitude—and there are many interesting results about magnitude homology which do not appear to have immediate consequences for magnitude—to the best of our knowledge Theorem 6.5 is the first *new* result about magnitude to be proved using magnitude homology. We believe this offers reason to hope that other mysteries in the theory of magnitude may eventually be resolvable using homological techniques.

THE STRUCTURE OF THIS PAPER We begin, in Section 2, with an overview of magnitude and magnitude homology. Section 3 introduces the main object of interest: a pair of graphs related by a sycamore twist. Sections 4–7 are dedicated to the proof of the main theorem.

Acknowledgments

I am grateful to Tom Leinster for discussions throughout the development of this work, and to Richard Hepworth for his comments on the early draft of this paper which forms part of my doctoral thesis. Thanks are due also to Sebastian Schlegel Mejia for helpful conversations concerning homological algebra, and to the anonymous referee, whose attentive reading and suggestions have improved the exposition in several places.

2. Magnitude and magnitude homology

Throughout this paper, a **graph** is an undirected finite graph with no loops or multiple edges. We endow a graph X with the metric in which the distance $d(u, v)$ (or $d_X(u, v)$) from a vertex u to a vertex v is the number of edges in a minimal edge-path connecting them, or ∞ if no such edge-path exists.

Let $\mathbb{Z}[q]$ denote the ring of polynomials over \mathbb{Z} in one variable, q . Given a graph X , we construct a square matrix Z_X whose rows and columns are indexed by the vertices of X and whose (u, v) entry is

$$Z_X(u, v) = q^{d(u, v)} \in \mathbb{Z}[q],$$

adopting the convention that $q^\infty = 0$.

The diagonal entries of the matrix Z_X are all equal to 1, and each of its off-diagonal entries is either zero or q^n for some natural number $n \neq 0$. The determinant of Z_X is thus a polynomial in q with constant term 1, so is a unit in the ring $\mathbb{Z}[[q]]$ of power series in q with integer coefficients. This ensures that Z_X is invertible over $\mathbb{Z}[[q]]$, allowing us to make the following definition.

Definition 2.1 (Leinster [15], Definition 2.1) The **magnitude** of a graph X is

$$\text{Mag}(X) = \sum_{u, v \in X} Z_X^{-1}(u, v) \in \mathbb{Z}[[q]].$$

The magnitude of a graph is a specialization of an invariant defined in the vastly greater generality of finite enriched categories ([14], Section 1). Specialized in a different direction, to ordinary categories, magnitude turns out to have close links to topological Euler characteristic ([13], Propositions 2.11 and 2.12); interpreted for posets, it extends the theory of Möbius inversion. Perhaps its most fertile environment, though, is that of metric spaces regarded as categories enriched in the poset $([0, \infty], \geq)$ with monoidal structure given by addition. (For the classical account of this perspective on metric spaces, see Lawvere [12].) An overview of the properties of metric space magnitude can be found in [19].

Evaluating $\text{Mag}(X)$ at $q = e^{-t}$ for all nonnegative real numbers t yields the **magnitude function** discussed in the introduction:

$$\text{Mag}(X)(t) = \text{Mag}(X)|_{q=e^{-t}}.$$

In general this function may diverge at some t ; Example 2.2.7 in [14] describes a graph exhibiting this behaviour. In what follows, we will deal exclusively with the power series $\text{Mag}(X)$.

Given magnitude's relationship to Euler characteristic, it is natural to ask whether there exists an algebraic invariant which 'categorifies' magnitude in the sense that singular homology categorifies the Euler characteristic of a topological space. In 2015 Hepworth and Willerton answered this question affirmatively in the case of graphs, constructing a homology theory whose Euler characteristic recovers their magnitude [10]; in 2017 Leinster and Shulman extended that construction to a broad class of enriched categories, including all metric spaces [21]. In this paper we will be working exclusively with Hepworth and Willerton's *magnitude homology* for graphs, which we now describe.

Their construction depends on a combinatorial formula for the coefficients in the power series $\text{Mag}(X)$, given by Leinster in [15]. To state the formula, it will be helpful to make the following definition.

Definition 2.2 A **path**, or a **k -path**, in a graph X is a tuple $\mathbf{x} = (x_0, \dots, x_k)$ of vertices in X . We call \mathbf{x} **non-degenerate** if $x_0 \neq x_1 \neq \dots \neq x_k$. The **length** of a path \mathbf{x} is

$$L(\mathbf{x}) = \sum_{i=0}^{k-1} d(x_i, x_{i+1}).$$

Remark 2.3 Notice we do not require that consecutive vertices in a path be connected by an edge; thus, every non-degenerate k -path in a graph X has length at least k , and possibly greater. Since all our graphs are finite, there can be only finitely many non-degenerate paths in X of any given length; moreover, there can be *no* non-degenerate k -paths of length ℓ for $k > \ell$.

Leinster's formula says that calculating the magnitude of a graph X comes down to counting non-degenerate paths in X of every possible length.

Proposition 2.4 ([15], Proposition 3.9) *For any graph X ,*

$$\text{Mag}(X) = \sum_{\ell=0}^{\infty} c_{\ell}(X)q^{\ell} \quad (2)$$

where the coefficients are given by

$$c_{\ell}(X) = \sum_{k=0}^{\ell} (-1)^k \#\{\text{non-degenerate } k\text{-paths of length } \ell \text{ in } X\}. \quad (3)$$

Hepworth and Willerton proceed from Proposition 2.4, constructing an \mathbb{N} -graded chain complex whose Euler characteristic in grading $\ell \in \mathbb{N}$ computes the coefficient c_{ℓ} in Leinster's formula.

Definition 2.5 ([10], Definition 2) The **magnitude chain complex** of a graph X is the direct sum of chain complexes

$$MC(X) = \bigoplus_{\ell \in \mathbb{N}} MC_{\bullet}^{\ell}(X)$$

where the chain complex $MC_{\bullet}^{\ell}(X)$ is freely generated in degree $k \geq 0$ by the set of non-degenerate k -paths in X whose length is ℓ . The boundary operator

$$d: MC_k^{\ell}(X) \rightarrow MC_{k-1}^{\ell}(X)$$

is an alternating sum $d = \sum_{i=1}^{k-1} (-1)^i \delta_i$ where δ_i is defined on generators by

$$\delta_i(x_0, \dots, x_k) = \begin{cases} (x_0, \dots, \widehat{x}_i, \dots, x_k) & \text{if } L(x_0, \dots, \widehat{x}_i, \dots, x_k) = \ell \\ 0 & \text{otherwise.} \end{cases}$$

Here, $(x_0, \dots, \widehat{x}_i, \dots, x_k)$ denotes the path $(x_0, \dots, x_{i-1}, x_{i+1}, \dots, x_k)$.

Definition 2.6 ([10], Definition 3) The **magnitude homology** of a graph X is the bigraded abelian group defined for $(k, \ell) \in \mathbb{N} \times \mathbb{N}$ by

$$MH_k^\ell(X) = H_k(MC_\bullet^\ell(X)).$$

We will refer to ℓ as the **length grading** and k as the **homological degree** of the abelian groups $MC_k^\ell(X)$ and $MH_k^\ell(X)$.

That the magnitude of a graph can be recovered from its magnitude chain complex follows immediately from Proposition 2.4.

Theorem 2.7 ([10], Theorem 2.8) *Let X be a graph. Then*

$$\text{Mag}(X) = \sum_{\ell \in \mathbb{N}} \chi(MC_\bullet^\ell(X)) q^\ell.$$

The construction of the magnitude chain complex—and thus magnitude homology—is functorial with respect to morphisms of graphs which preserve or contract edges; equivalently, vertex functions $f: X \rightarrow Y$ satisfying

$$d_Y(f(x), f(x')) \leq d_X(x, x')$$

for all pairs of vertices x, x' in X . These **distance-decreasing** maps are especially natural to consider when regarding a metric space as an enriched category.

The chain map induced by a distance-decreasing function $f: X \rightarrow Y$ is given on each generator $\mathbf{x} = (x_0, \dots, x_k)$ of $MC(X)$ by

$$MC(f)(\mathbf{x}) = \begin{cases} f(\mathbf{x}) & \text{if } L(f(\mathbf{x})) = L(\mathbf{x}) \\ 0 & \text{otherwise,} \end{cases}$$

where $f(\mathbf{x})$ denotes the path $(f(x_0), \dots, f(x_k))$ in Y . That is,

$$MC(f)(\mathbf{x}) = \begin{cases} f(\mathbf{x}) & \text{if } d_Y(f(x_i), f(x_{i+1})) = d_X(x_i, x_{i+1}) \text{ for } 0 \leq i < k \\ 0 & \text{otherwise.} \end{cases}$$

In other words, the induced map $MC(f): MC(X) \rightarrow MC(Y)$ retains information only about those paths in X which are mapped into Y by f in a strictly length-preserving manner.

On the other hand, lifting magnitude from an element of a ring of formal power series to an object of a category of chain complexes grants us the ability to study it homologically. By Remark 2.3 the chain complex $MC_\bullet^\ell(X)$ vanishes in homological degrees $k > \ell$, while in degrees $k \leq \ell$ it is finitely generated. Standard facts of homological algebra imply, then, that the Euler characteristic of the complex $MC_\bullet^\ell(X)$ coincides with that of its homology. Thus, we can compute the ℓ^{th} coefficient in the power series $\text{Mag}(X)$ as

$$c_\ell = \sum_{k=0}^{\ell} (-1)^k \text{rk}(MH_k^\ell(X)).$$

That is,

$$\text{Mag}(X) = \sum_{\ell \in \mathbb{N}} \chi(MH_\bullet^\ell(X)) q^\ell. \quad (4)$$

With this formula, homological techniques can be used to simplify the calculation of

the coefficients in $\text{Mag}(X)$. This is the principle we will be applying in later sections of this paper.

Before introducing the main object of interest, we briefly review the inclusion-exclusion theorem for magnitude and the Mayer–Vietoris theorem for the magnitude homology of graphs. Both theorems hold under the same conditions, captured by the next three definitions.

Definition 2.8 ([15], Definition 4.2) A subgraph $W \subseteq X$ is called **convex** if it satisfies $d_W(u, v) = d_X(u, v)$ for all pairs of vertices $u, v \in W$. In other words, W is convex if the inclusion $W \hookrightarrow X$ is an isometric embedding.

Definition 2.9 Let $W \subseteq X$ be a convex subgraph. We say that a vertex $v \in X$ **projects** to W if v is connected by an edge-path to some vertex in W and there exists a vertex $\pi(v) \in W$ such that

$$d(v, w) = d(v, \pi(v)) + d(\pi(v), w)$$

for every vertex w in W . We say a subgraph $Y \subseteq X$ **projects** to W if $Y \cap W \neq \emptyset$ and every vertex in Y which is connected by an edge-path to W projects to W .

Note that, if a vertex v projects to W , then $\pi(v)$ is the *unique* vertex of W closest to v . Thus, writing U_W for the set of vertices in X which project to W , we have a function $\pi: U_W \rightarrow W$.

Definition 2.10 ([10], Definition 26) A **projecting decomposition** is a triple $(X; G, H)$ where X is a graph with subgraphs G and H such that the following properties hold.

- $X = G \cup H$.
- $G \cap H$ is convex in X .
- H projects to $G \cap H$.

The inclusion-exclusion formula for the magnitude of graphs (Leinster [15], Theorem 4.9) says that if $(X; G, H)$ is a projecting decomposition then

$$\text{Mag}(X) = \text{Mag}(G) + \text{Mag}(H) - \text{Mag}(G \cap H).$$

The Mayer–Vietoris theorem for magnitude homology (Hepworth and Willerton [10], Theorem 29) categorifies this formula by a split short exact sequence

$$0 \rightarrow MH_{\bullet}^*(G \cap H) \rightarrow MH_{\bullet}^*(G) \oplus MH_{\bullet}^*(H) \rightarrow MH_{\bullet}^*(X) \rightarrow 0.$$

Hepworth and Willerton do not attempt to categorify Leinster’s theorem concerning Whitney twists (Theorem 1.1 in this paper). Instead, they pose the question: do two graphs related by a Whitney twist along adjacent gluing vertices have isomorphic magnitude homology?

That question will not be answered here. Rather, we will give a homological proof of Leinster’s theorem which does not depend on the existence of an isomorphism of magnitude homologies. Our main theorem will extend Leinster’s result to encompass a wider class of twisted gluings we term *sycamore twists*. The next section introduces these.

3. Sycamore twists

We begin by stating the definition of a Whitney twist more formally.

Definition 3.1 Let G be a graph with two distinct distinguished vertices g_+ and g_- , and H a graph with distinct distinguished vertices h_+ and h_- . Form a new graph X by taking the disjoint union of G and H and identifying g_+ with h_+ and g_- with h_- , then identifying any double edges that result. Form another graph Y in the same manner, this time identifying g_+ with h_- and g_- with h_+ . The graphs X and Y then differ by a **Whitney twist**.




Figure 2: A Whitney twist along adjacent gluing vertices. Dashed edges and their incident vertices belong to G ; dotted edges and their incident vertices belong to H ; and the solid edge and its incident vertices belong to $G \cap H$. Notice that the graphs X and Y are non-isometric and that neither $(X; G, H)$ nor $(Y; G, H)$ is a projecting decomposition, since the vertices g_1, g_2, h_1 and h_2 do not project to $G \cap H$.

Consider the following straightforward generalization of Definition 3.1.

Definition 3.2 Let G, H and K be graphs equipped with induced subgraph inclusions

$$G \xleftarrow{\iota_G} K \xrightarrow{\iota_H} H$$

and let $\alpha: K \rightarrow K$ be an isometry. Form a new graph X by taking the disjoint union $G \sqcup H$ and identifying the vertices $\iota_G(v)$ and $\iota_H(v)$ for each $v \in K$; wherever a double edge is created, identify the two edges. Form another graph Y in the same way, but identifying $\iota_G(v)$ with $\iota_H(\alpha(v))$ for each $v \in K$. We will say that X and Y differ by a **generalized Whitney twist**, or just by a **twist**, and specify the twist by the tuple (G, H, K, α) .

There are two scenarios in which we can already be sure that two graphs differing by a twist will have the same magnitude—indeed, the same magnitude homology. One relates to the Mayer–Vietoris sequence. If K is convex in X and Y , and H projects to K , then $(X; G, H)$ is a projecting decomposition and so is $(Y; G, H)$. In that case the Mayer–Vietoris sequence tells us that the homology of both X and Y is determined by that of G, H and K ; in particular, we have

$$MH_{\bullet}^*(X) \cong MH_{\bullet}^*(Y) \tag{5}$$

and thus $\text{Mag}(X) = \text{Mag}(Y)$.

On the other hand, suppose α extends to a self-isometry of H which fixes every vertex in $H \setminus K$ —that is, suppose

$$d_H(v, k) = d_H(v, \alpha(k)) \quad (6)$$

for all $v \in H \setminus K$ and all $k \in K$. In this case X and Y are isometric, so certainly their magnitude homology agrees—though in general we cannot expect to compute it from the homologies of G , H and K .

For graphs differing by a Whitney twist—even one with adjacent gluing vertices—neither of these scenarios necessarily applies (see Figure 2). However, a Whitney twist along adjacent gluing vertices does satisfy a hybrid of the two conditions. In fact, Leinster’s proof of the invariance of magnitude under such twists relies on the observation that the vertex sets of X and Y can be partitioned into those vertices which lie equidistant from g_+ and g_- , and those which lie closer to one gluing vertex than the other. The first subset consists precisely of those v satisfying (6); the second—thanks to the presence of the edge joining g_+ and g_- —projects to the gluing set.

Our main theorem concerns twists that possess a relaxed version of this hybrid property: every vertex of the subgraph H either projects to the gluing set or else satisfies (6).

Definition 3.3 A **sycamore twist** is a generalized Whitney twist (G, H, K, α) satisfying two additional conditions:

- K is convex in X and Y .
- Every vertex $h \in H$ which does *not* project to K satisfies $d_H(h, k) = d_H(h, \alpha(k))$ for every $k \in K$.

Not every pair of graphs that differ by sycamore twist can be related by a Whitney twist, as the following example shows.

Example 3.4 Consider the graphs X and Y depicted in Figure 3. They differ by a sycamore twist (G, H, K, α) : dashed edges belong to G ; dotted edges to H ; and solid edges to $K = G \cap H$. The map α interchanges v_5 and v_6 , fixing v_4 and v_7 . The vertices v_8 and v_{11} project to K , while v_9 and v_{10} satisfy equation (6) with respect to each vertex in K .

The graphs X and Y cannot be related by a Whitney twist. If they could, then the two gluing vertices would form a vertex cut in X (and in Y): deleting those two vertices and their incident edges would disconnect the graph. The graph X contains exactly 12 two-element vertex cuts. (One can count them by hand, or with a few lines of code.) Thus, there are in principle 12 graphs to which it can be related by a Whitney twist. However, each of those graphs is in fact isometric to X ; in particular, none of them is isometric to Y .

Equally, not every Whitney twist is a sycamore twist. An example is given by any Whitney twist in which the gluing vertices are non-adjacent yet are connected by an edge-path in X (or equivalently in Y)—for in this case, the subgraph on the gluing vertices is not convex. However, in a Whitney twist along *adjacent* vertices, this subgraph is guaranteed to be convex. Indeed, by the discussion preceding Definition 3.3,




Figure 3: A sycamore twist which is not a Whitney twist. The sycamore twist is named after the winged seed pod of the sycamore tree, which twists in the air as it falls. (British schoolchildren call these ‘helicopters’.) The idea is that the vertices which project to K resemble wings projecting from a central core.

any such twist is a sycamore twist with respect to the subgraph K comprising the two gluing vertices and the edge between them; the map α flips the vertices of K .

This paper’s main Theorem (Theorem 6.5) says that magnitude is invariant under sycamore twists. Example 3.4 shows that this is a proper generalization of Leinster’s result. We will prove it using a homological argument, but the proof will not imply that magnitude homology is invariant under sycamore twists. Instead, the homological algebra serves to simplify the counting of paths in X and Y . After establishing, in Section 4, a few basic facts about the metrics on X and Y , we partition the set of paths in each graph into those which are ‘twistable’—meaning that there is a particular bijection between the vertices of X and Y which preserves their length—and those which are not. (This is the subject of Section 5.) By showing that non-twistable paths generate a chain complex which is contractible, we discover (in Sections 6 and 7) that we can discount them when it comes to calculating magnitude.

4. Properties of sycamore twists

This section establishes those properties of a sycamore twist which will facilitate our analysis of its magnitude. We begin by fixing notation and recording basic information about the distance functions on the twisted graphs X and Y . We will see that the set of non-gluing vertices in H can be partitioned into two subsets, each with a convenient property derived from the defining properties of a sycamore twist. Finally, we relate X and Y by two functions on their vertex sets which restrict locally to isometries; in Section 5 these functions will be used to establish a relationship between the magnitude complexes of X and Y .

Definition 4.1 (Notation for gluing vertices) Let (G, H, K, α) be a sycamore twist. Given a vertex $v \in K$, we denote its image $\iota_G(v)$ in G simply by v , and do the same for its image $\iota_H(v)$ in H . We also use v to denote the vertex $\iota_G(v) = \iota_H(v)$ in X and the vertex $\iota_G(v) = \iota_H(\alpha(v))$ in Y .

Using this notation, for all $g \in G$, $v \in K$ and $h \in H$ we have

$$d_X(g, v) = d_G(g, v) \text{ and } d_X(h, v) = d_H(h, v)$$

while for $g \in G$, $v \in K$ and $h \in H \setminus K$ we have

$$d_Y(g, v) = d_G(g, v) \text{ and } d_Y(h, v) = d_H(h, \alpha(v)).$$

It will be useful to record a description of the other distances in X and in Y .

Lemma 4.2 *Given vertices u and v in X we have*

$$d_X(u, v) = \begin{cases} d_G(u, v) & u, v \in G \\ d_H(u, v) & u, v \in H \setminus K \\ \min_{k \in K} (d_G(u, k) + d_H(k, v)) & u \in G \text{ and } v \in H \setminus K. \end{cases}$$

Given vertices u and v in Y we have

$$d_Y(u, v) = \begin{cases} d_G(u, v) & u, v \in G \\ d_H(u, v) & u, v \in H \setminus K \\ \min_{k \in K} (d_G(u, k) + d_H(\alpha(k), v)) & u \in G \text{ and } v \in H \setminus K. \end{cases}$$

Proof. The statements concerning $u, v \in G$ and $u, v \in H \setminus K$ follow from Lemma 4.3 in [15], which says that if $K = G \cap H$ is convex in X (or in Y) then both G and H are also convex in X (respectively in Y). Take $u \in G$ and $v \in H \setminus K$; we want to see that

$$d_X(u, v) = \min_{k \in K} (d_G(u, k) + d_H(k, v)) \quad (7)$$

and

$$d_Y(u, v) = \min_{k \in K} (d_G(u, k) + d_H(\alpha(k), v)). \quad (8)$$

Let k be any vertex in K . Suppose there exists an edge-path in G between u and k , and an edge-path in H between k and v : the concatenation of any such pair of paths gives an edge path from u to v in X , so all three distances are finite and we have

$$d_X(u, v) \leq d_G(u, k) + d_H(k, v).$$

On the other hand, if there is *no* edge-path in G between u and k then $d_G(u, k) = \infty$ and the inequality still holds; the same is true if there is no edge-path in H between k and v . Thus

$$d_X(u, v) \leq \min_{k \in K} (d_G(u, k) + d_H(k, v)).$$

For the reverse inequality, we may assume $d_X(u, v)$ is finite. Let $\mathbf{x} = (x_0, \dots, x_n)$ be a minimal edge path from u to v in X . Any path between a vertex in G and one in H must pass through at least one vertex in K (this is Lemma 4.4 in [15]). Let x_i be such a vertex in \mathbf{x} . Then (x_0, \dots, x_i) is a minimal edge path between u and x_i in X , and as G is convex in X this implies $d_G(u, x_i) = d_X(u, x_i) = i$. Similarly, $d_H(x_i, v) = d_X(x_i, v) = n - i$. So $d_X(u, v) = d_G(u, x_i) + d_H(x_i, v)$ for some vertex $x_i \in K$, and hence

$$d_X(u, v) \geq \min_{k \in K} (d_G(u, k) + d_H(k, v)).$$

This proves that (7) holds; the same argument establishes (8), after making use of the fact that $d_Y(k, v) = d_H(\alpha(k), v)$ for all $k \in K$ and $v \in H \setminus K$. \square

Definition 4.3 Let (G, H, K, α) be a sycamore twist. The vertices of K are the **gluing** vertices. Vertices in $H \setminus K$ which project to K will be called **biased** vertices; we will denote the set of biased vertices by H_* . A vertex v in $H \setminus K$ which does *not*

project to K will be called α -**neutral** or just **neutral**; we will denote the set of neutral vertices by H_0 .

The terminology is motivated by Lemma 4.4, which says that neutral vertices do not notice the difference between the metrics in X and in Y . Lemma 4.6 tells us that biased vertices have a complementary property: in both X and Y the vertices in H_* project to G .

Lemma 4.4 *Let (G, H, K, α) be a sycamore twist and let v be a neutral vertex in H . For every $u \in G \cup H$ we have*

$$d_X(u, v) = d_Y(u, v).$$

Proof. By the definition of a sycamore twist, each neutral vertex v satisfies

$$d_H(k, v) = d_H(\alpha(k), v)$$

for all $k \in K$. The result follows upon comparing the expressions for $d_X(u, v)$ and $d_Y(u, v)$ in Lemma 4.2. \square

Definition 4.5 A vertex w in a graph G is **between** vertices u and v if

$$d_G(u, v) = d_G(u, w) + d_G(w, v).$$

Lemma 4.6 *Let (G, H, K, α) be a sycamore twist. Then $h \in H$ projects to K in H if and only if h projects to G in X and in Y .*

Proof. The “if” statement is clear; we must prove the converse. Suppose $h \in H$ projects to K in H , and let g be a vertex in G . If g is a gluing vertex then

$$d_X(h, g) = d_H(h, g) = d_H(h, \pi(h)) + d_H(\pi(h), g) = d_X(h, \pi(h)) + d_X(\pi(h), g),$$

so we need to see that in X the vertex $\pi(h)$ lies between h and every non-gluing vertex of G . Take $g \in G \setminus K$. Via Lemma 4.2 we have

$$\begin{aligned} d_X(g, h) &= \min_{k \in K} (d_G(g, k) + d_H(k, h)) \\ &= \min_{k \in K} (d_G(g, k) + d_H(k, \pi(h)) + d_H(\pi(h), h)) \\ &= \min_{k \in K} (d_G(g, k) + d_H(k, \pi(h))) + d_H(\pi(h), h) \\ &= d_X(g, \pi(h)) + d_X(\pi(h), h) \end{aligned}$$

as required. Thus, h projects to G in X .

The argument for Y is essentially the same. If $g \in G$ is a gluing vertex, then

$$\begin{aligned} d_Y(g, h) &= d_H(\alpha(g), h) = d_H(\alpha(g), \pi(h)) + d_H(\pi(h), h) \\ &= d_Y(g, \alpha^{-1}(\pi(h))) + d_Y(\alpha^{-1}(\pi(h)), h) \end{aligned}$$

so we need to see that $\alpha^{-1}(\pi(h))$ lies between h and every non-gluing vertex of G .

Take $g \in G \setminus K$. We have

$$\begin{aligned}
d_Y(g, h) &= \min_{k \in K} (d_G(g, k) + d_H(\alpha(k), h)) \\
&= \min_{k \in K} (d_G(g, k) + d_H(\alpha(k), \pi(h)) + d_H(\pi(h), h)) \\
&= \left(\min_{k \in K} (d_G(g, k) + d_H(\alpha(k), \pi(h))) \right) + d_Y(\alpha^{-1}(\pi(h)), h) \\
&= \left(\min_{k \in K} (d_Y(g, k) + d_Y(k, \alpha^{-1}(\pi(h)))) \right) + d_Y(\alpha^{-1}(\pi(h)), h) \\
&= d_Y(g, \alpha^{-1}(\pi(h))) + d_Y(\alpha^{-1}(\pi(h)), h),
\end{aligned}$$

where the final line holds by the triangle inequality in Y , since $\alpha^{-1}(\pi(h))$ belongs to K . Thus, if h projects to K in H , it projects to G in X and in Y . \square

Whenever X and Y differ by a twist (G, H, K, α) we can define bijective functions τ_G, τ_H from the vertex set of X to that of Y by

$$\tau_G(v) = v \text{ for all } v \in X$$

and

$$\tau_H(v) = \begin{cases} v & \text{if } v \text{ is not a gluing vertex} \\ \alpha^{-1}(v) & \text{if } v \text{ is a gluing vertex.} \end{cases}$$

(Here, as above, we are using the notation for gluing vertices established in Definition 4.1.)

The function τ_G can be thought of as fixing the subgraph G and twisting H , while τ_H fixes H and twists G . Except in degenerate cases, neither τ_G nor τ_H defines a graph homomorphism on the whole of X . Rather, they are constructed so that, provided the twist is a *sycamore* twist, each map is a graph homomorphism—in fact, a bijective isometry—when restricted to one part of the cover $\{G \cup H_0, H\}$.

Lemma 4.7 *Let (G, H, K, α) be a sycamore twist. Then the function τ_G restricts to an isometry on $G \cup H_0$ and τ_H restricts to an isometry on H .*

Proof. It is clear that τ_G restricts to an isometry on G and on H_0 separately, and Lemma 4.4 tells us that for all $g \in G$ and $h \in H_0$ we have

$$d_Y(\tau_G(g), \tau_G(h)) = d_X(g, h),$$

so τ_G restricts to an isometry on $G \cup H_0$.

Similarly, it is clear that τ_H restricts to an isometry on $H \setminus K$. It also restricts to an isometry on K : given any $k, k' \in K$, we have

$$\begin{aligned}
d_Y(\tau_H(k), \tau_H(k')) &= d_Y(\alpha^{-1}(k), \alpha^{-1}(k')) \\
&= d_H(\alpha(\alpha^{-1}(k)), \alpha(\alpha^{-1}(k'))) \\
&= d_H(k, k') \\
&= d_X(k, k')
\end{aligned}$$

as required. Finally, if $h \in H \setminus K$ and $k \in K$ then

$$d_Y(\tau_H(h), \tau_H(k)) = d_Y(h, \alpha^{-1}(k)) = d_H(h, \alpha(\alpha^{-1}(k))) = d_H(h, k) = d_X(h, k).$$

Thus, τ_H restricts to an isometry on H . □

Section 5 will make use of one more lemma which, although we state it in the notation of our present setup, is a general and elementary fact about projection.

Lemma 4.8 *Suppose $u, v \in G$ and $w \in H_*$. Then v lies between u and w if and only if it lies between u and $\pi(w)$.*

Proof. Since $\pi(w)$ is between u and w , we have

$$d(u, \pi(w)) = d(u, w) - d(\pi(w), w).$$

If, and only if, v is between u and w , we can rewrite the right hand side as $d(u, v) + d(v, w) - d(\pi(w), w)$. Since $\pi(w)$ is between v and w , this yields

$$d(u, \pi(w)) = d(u, v) + d(v, \pi(w)) + d(\pi(w), w) - d(\pi(w), w) = d(u, v) + d(v, \pi(w)).$$
□




Figure 4: The subgraphs G' (shaded in light gray) and H' (shaded in dark gray) do not cover X or Y as graphs: the edges (h_0, h_1) and (h_2, h_3) are missing from $G' \cup H'$.

Remark 4.9 Given a sycamore twist (G, H, K, α) , let G' denote the subgraph of X whose vertices are those in $G \cup H_0$ and let H' denote the subgraph whose vertices are those in $H_* \cup K$. Then the vertex sets of G' and H' cover the vertex set of X , and every vertex in H' projects to $G' \cap H' = K$.

If $(X; G', H')$ were a projecting decomposition, and the same were true for the similar triple $(Y; G', H')$, then our main theorem would follow from the inclusion-exclusion formula. But this is not the case. In general G' and H' do not cover X as graphs: there may exist edges (u, v) in X such that u belongs to $G' \setminus (G' \cap H')$ and v belongs to $H' \setminus (G' \cap H')$, in which case (u, v) does not appear in $G' \cup H'$. Figure 4 provides an example.

5. Twistable and non-twistable paths

Consider the graphs X and Y in Figure 5. There are evident differences in the sorts of paths that exist in each graph: for instance, Y has two 1-paths of length 3—the paths (g_0, h_1) and (h_1, g_0) —while X has none. On the other hand, the two graphs also have many paths in common. For example, any path which never visits the biased

vertex h_1 evidently has the same length whether it is regarded as belonging to X or to Y . And the same is true for any path which never visits h_1 without visiting the neutral vertex h_0 immediately beforehand and immediately afterwards—this follows from the fact, characteristic of neutral vertices, that $d_X(h_0, -) = d_Y(h_0, -)$ (Lemma 4.4).




Figure 5: The graphs X and Y differ by a sycamore twist along the graph k_0 – k_1 .

Proposition 5.6, below, extrapolates from examples of this type. We use the vertex functions τ_G and τ_H defined in Section 4 to construct a length-preserving bijection between those paths in X and Y which never cross from G to H_* (or *vice versa*) without passing through a neutral vertex. Paths with this property will be called ‘twistable’; the next few definitions formalize this notion.

Definition 5.1 Let X be any graph and U a subset of its vertices. A path $\mathbf{x} = (x_0, \dots, x_k)$ in X will be said to **visit n vertices in U** if there are n indices, i_1, \dots, i_n , such that the vertex x_{i_j} belongs to U . We will say \mathbf{x} is **contained in U** if every vertex x_0, \dots, x_k belongs to U . A **subpath** of \mathbf{x} is a string $(x_i, x_{i+1}, \dots, x_j)$ of consecutive vertices in \mathbf{x} , where $0 \leq i \leq j \leq k$.

Definition 5.2 Let (G, H, K, α) be a sycamore twist, and let $\mathbf{x} = (x_0, \dots, x_k)$ be a path in X (or in Y). We will say \mathbf{x} is **flat** if it is contained in $G \cup H_0$ or contained in H .

Every flat path can be mapped in a length-preserving manner from X to Y using one or other of the vertex functions τ_G and τ_H .

Lemma 5.3 Suppose X and Y differ by a sycamore twist. For each $k \geq 0$ there is a length-preserving bijection between flat k -paths in X and flat k -paths in Y .

Proof. Given a flat path \mathbf{x} in X , let $T(\mathbf{x})$ be the path in Y defined by

$$T(\mathbf{x}) = \begin{cases} \tau_G(\mathbf{x}) & \text{if } \mathbf{x} \text{ is contained in } G \cup H_0 \\ \tau_H(\mathbf{x}) & \text{otherwise.} \end{cases}$$

Clearly, if \mathbf{x} is a k -path then $T(\mathbf{x})$ is a k -path too. Since τ_G restricts to an isometry on $G \cup H_0$ and τ_H restricts to an isometry on H (by Lemma 4.7), the path $T(\mathbf{x})$ is flat and its length is the same as that of \mathbf{x} .

Thus, we have a length-preserving function

$$T: \{\text{flat } k\text{-paths in } X\} \rightarrow \{\text{flat } k\text{-paths in } Y\}.$$

We can construct a function T' in the other direction in the same manner, replacing

τ_G by τ_G^{-1} and τ_H by τ_H^{-1} . The maps T and T' are mutually inverse, proving the lemma. \square




Figure 6: On the left we see two flat paths in the graph X from Figure 2. The path \mathbf{w} (dashed) is contained in $G \cup H_0$ and the path \mathbf{x} (solid) is contained in H . On the right we see their images $T(\mathbf{w})$ and $T(\mathbf{x})$ in Y .

Suppose we are given a flat path $\mathbf{x} = (x_0, \dots, x_k)$ which is contained in $G \cup H_0$ and a flat path $\mathbf{x}' = (x'_0, \dots, x'_n)$ which is contained in H , with the property that $x_k = x'_0$ and this vertex is neutral. Then we can concatenate the paths \mathbf{x} and \mathbf{x}' in X , and since $\tau_G(x_k) = \tau_H(x'_0)$ (as τ_G and τ_H agree on all but the gluing vertices) we can also concatenate the paths $T(\mathbf{x})$ and $T(\mathbf{x}')$ in Y . Length sums over concatenation of paths, so we have that

$$L(T(\mathbf{x}) * T(\mathbf{x}')) = L(T(\mathbf{x})) + L(T(\mathbf{x}')) = L(\mathbf{x}) + L(\mathbf{x}') = L(\mathbf{x} * \mathbf{x}').$$

By the same reasoning, T can be extended in a length-preserving manner to any sequence of finitely many flat paths concatenated at neutral vertices. We now give a name to such concatenations.

Definition 5.4 Let (G, H, K, α) be a sycamore twist. A path \mathbf{x} in X (or in Y) is **twistable** if it can be decomposed as a concatenation of paths

$$\mathbf{x} = \mathbf{x}_0 * \mathbf{x}_2 * \dots * \mathbf{x}_m$$

such that each path \mathbf{x}_i is flat and each point of concatenation is a neutral vertex. We will call such a decomposition a **maximal decomposition into flat subpaths** if every neutral vertex visited by \mathbf{x} is a point of concatenation.

In order to extend the function T in a canonical manner from flat to twistable paths, we need to specify a canonical decomposition of each twistable path into flat subpaths. The following lemma says that this is possible.

Lemma 5.5 Every twistable path has a unique maximal decomposition into flat subpaths.

Proof. Let \mathbf{x} be a twistable path in X . If \mathbf{x} visits no neutral vertices then \mathbf{x} itself must be flat, and in this case there is nothing to show. On the other hand, if \mathbf{x} visits at least one neutral vertex, we can decompose it as follows.

Extract the ordered list of neutral vertices visited by \mathbf{x} —say,

$$(x_{i_0}, \dots, x_{i_k}).$$

Let \mathbf{x}_0 be the subpath (x_0, \dots, x_{i_0}) of \mathbf{x} , let \mathbf{x}_k be the subpath (x_{i_k}, \dots, x_n) , and for each $0 < j < k$ let \mathbf{x}_j be the subpath $(x_{i_j}, \dots, x_{i_{j+1}})$. Then

$$\mathbf{x} = \mathbf{x}_0 * \dots * \mathbf{x}_k \tag{9}$$

and in this decomposition each point of concatenation is neutral, while every neutral vertex visited by \mathbf{x} occurs as a point of concatenation. We just need to see that every subpath \mathbf{x}_j is flat.

As \mathbf{x} is twistable we can find *some* decomposition of \mathbf{x} into flat subpaths:

$$\mathbf{x} = \mathbf{x}'_0 * \dots * \mathbf{x}'_m \tag{10}$$

Since each point of concatenation in (10) is neutral, each of the paths $\mathbf{x}_0, \dots, \mathbf{x}_k$ in (9) must be a subpath of \mathbf{x}'_i for some $0 \leq i \leq m$. And since each subpath \mathbf{x}'_i in (10) is flat—that is, contained either in $G \cup H_0$ or in H —the same must be true for each \mathbf{x}_j . Thus, (9) is a maximal decomposition into flat subpaths, and this decomposition is evidently unique as such. \square

Lemma 5.5 lets us extend the bijection T established in Lemma 5.3 to a bijection between the twistable paths in X and those in Y .

Proposition 5.6 *Suppose X and Y differ by a sycamore twist. Then for each $k \geq 0$ there is a length-preserving bijection between twistable k -paths in X and twistable k -paths in Y .*

Proof. Given a twistable k -path \mathbf{x} in X , let

$$\mathbf{x} = \mathbf{x}_0 * \dots * \mathbf{x}_m$$

be its maximal decomposition into flat subpaths. Let $T(\mathbf{x})$ be the k -path in Y defined by

$$T(\mathbf{x}) = T(\mathbf{x}_0) * \dots * T(\mathbf{x}_m) \tag{11}$$

where $T(\mathbf{x}_i)$ is specified, for each flat subpath \mathbf{x}_i , as in the proof of Lemma 5.3.

Since τ_G preserves the vertex set $G \cup H_0$ and τ_H preserves the vertex set H , and both maps fix the neutral vertices, (11) is a decomposition of $T(\mathbf{x})$ into flat subpaths and thus $T(\mathbf{x})$ is twistable. Indeed, since all its neutral vertices are points of concatenation, (11) is the maximal decomposition of $T(\mathbf{x})$ into flat subpaths.

Moreover, we have

$$L(T(\mathbf{x})) = \sum_{i=0}^m L(T(\mathbf{x}_i)) = \sum_{i=0}^m L(\mathbf{x}_i) = L(\mathbf{x}),$$

which says that the map

$$T: \{\text{twistable } k\text{-paths in } X\} \rightarrow \{\text{twistable } k\text{-paths in } Y\}$$

is length-preserving. We can construct map T' in the other direction by extending, in the same fashion, the function T' of Lemma 5.3. The maps T and T' are mutually inverse, proving the proposition. \square

Having established a bijection between the twistable paths in X and those in Y , our aim in Sections 6 and 7 is to prove that every path which is *not* twistable can be disregarded for the purposes of calculating magnitude. For this, we will need a more positive characterization of the ‘non-twistable’ paths. The remainder of this section provides that characterization.

Definition 5.7 A path $\mathbf{x} = (x_0, \dots, x_k)$ in X or Y is said to be **sticky** if x_0 is a vertex in $G \setminus K$ and x_k is a vertex in H_* , or *vice versa*, and x_1, \dots, x_{k-1} are all gluing vertices.

Since a sticky path always visits both $G \setminus K$ and H_* , no flat path can have a subpath which is sticky. A partial converse also holds:

Lemma 5.8 *Let \mathbf{x} be a path which is not flat and which visits no neutral vertices except perhaps its endpoints. Then \mathbf{x} has a sticky subpath.*

Proof. As it is not flat, the path $\mathbf{x} = (x_0, \dots, x_k)$ must visit at least one vertex in $G \setminus K$ and at least one vertex in H_* . Identify vertices $x_i \in G \setminus K$ and $x_j \in H_*$ in \mathbf{x} such that $|j - i|$ is minimal. We may assume without loss of generality that $i < j$. As $|j - i|$ is minimal, the vertices x_{i+1}, \dots, x_{j-1} must all belong to $K \cup H_0$; in fact, as \mathbf{x} visits no neutral vertices (except perhaps its endpoints), they must all belong to K . Thus, the subpath (x_i, \dots, x_j) is sticky. \square

Using the lemma, we can characterize the ‘non-twistable’ paths in X and Y as those which possess at least one sticky subpath.

Proposition 5.9 *A path is twistable if and only if has no sticky subpath.*

Proof. First, suppose \mathbf{x} is a twistable path, and let $\mathbf{x} = \mathbf{x}_0 * \dots * \mathbf{x}_k$ be a decomposition into flat subpaths. If \mathbf{x} were to have a sticky subpath, say \mathbf{w} , then—as a sticky path visits no neutral vertices— \mathbf{w} must be a subpath of \mathbf{x}_i for some $0 \leq i \leq k$. But no flat path can have a sticky subpath; thus, \mathbf{x} has no sticky subpath.

Conversely, suppose \mathbf{x} has no sticky subpath. If \mathbf{x} visits no neutral vertices, then Lemma 5.8 implies that \mathbf{x} must be flat. On the other hand, if \mathbf{x} visits at least one neutral vertex, we can decompose it by concatenating at every neutral vertex as in Lemma 5.5 (equation (9)). To see that every subpath \mathbf{x}_j in this decomposition is flat, observe that for each j we either have $\mathbf{x}_j = (x_{j_0}, x_{j_1})$ where x_{j_0} and x_{j_1} are both neutral, or else $\mathbf{x}_j = (x_{j_0}, \dots, x_{j_m})$ where $m > 1$ and none of the vertices $x_{j_1}, \dots, x_{j_{m-1}}$ is neutral. In the first case \mathbf{x}_j is contained in H_0 , so is flat. In the second case, \mathbf{x}_j is flat by Lemma 5.8.

Thus we have a decomposition into flat subpaths, so \mathbf{x} is twistable. \square

6. Magnitude is invariant under sycamore twists

Having established that the twistable paths of a given length in X correspond bijectively with those in Y , we would like to be able to disregard the paths that are not twistable; by Proposition 5.9, the non-twistable paths are those which possess at least one sticky subpath.

Ideally, one might hope to show that the non-twistable paths span an acyclic subcomplex of the magnitude complex and thus do not contribute to its homology—or, therefore, to its Euler characteristic. The trouble is that the non-twistable paths *do not* span a subcomplex of $MC(X)$ or $MC(Y)$, acyclic or otherwise. A sticky subpath can easily be destroyed, for instance if removing one of its endpoints brings its gluing vertices into contact with a neutral vertex; Figure 7 gives an example. Consequently, the subset of generators which contain a sticky subpath is not closed under the boundary operator.




Figure 7: Here we see two paths (dashed) in the graph X from Figure 2. In the path \mathbf{x} on the left there is a sticky subpath, (g_0, k_0, h_3) . On the right is the path $\delta_3(\mathbf{x})$, which is flat: it is contained in $G \cup H_0$.

To circumvent this problem we are going to filter the magnitude complex in such a way that, when we consider any term-by-term subquotient of the filtration, the non-twistable paths *do* span a subcomplex. That subcomplex will turn out to be acyclic. The proof of its acyclicity, which is inspired by Hepworth and Willerton’s proof of the excision theorem for magnitude chains ([10], Theorem 28), is given in Section 7.

The idea behind the filtration is as follows. Since a sticky subpath may be destroyed by removing one of its endpoints, and since the endpoints of a sticky subpath always belong to the vertex set $X \setminus (H_0 \cup K)$, we would like to engineer things so that removing a vertex in $X \setminus (H_0 \cup K)$ from a generator always has the effect of sending that generator to zero. To achieve this, we filter the magnitude chains by the number of vertices they visit in $X \setminus (H_0 \cup K)$.

Formally, for each $m \in \mathbb{N}$, let T_m^X denote the subcomplex of $MC(X)$ spanned in degree $k \in \mathbb{N}$ by k -paths that visit at most m vertices in $X \setminus (H_0 \cup K)$. We can filter $MC(X)$ by these subcomplexes—

$$0 \subseteq T_0^X \subseteq T_1^X \subseteq \dots \subseteq T_m^X \subseteq \dots \subseteq MC(X) \tag{12}$$

—and we can filter $MC(Y)$ in a similar manner. We have $T_{m,k}^X = MC_k(X)$ for all $k < m$, and the same for Y . What’s more, in any given length grading ℓ the magnitude chain complex vanishes above homological degree $k = \ell$ (since any k -path in a graph has length at least k), so in every length grading ℓ this filtration stabilizes after the ℓ^{th} term.

For each m , write $Q_m^X = T_m^X / T_{m-1}^X$. The complex Q_m^X is freely generated in degree k by paths $\mathbf{x} = (x_0, \dots, x_k)$ that visit *exactly* m vertices in $X \setminus (H_0 \cup K)$. The boundary map is

$$d = \sum_{i=1}^{k-1} (-1)^i \delta_i$$

where δ_i removes the vertex x_i unless x_i belongs to $X \setminus (H_0 \cup K)$ or removing it reduces the length of the path; in either of those cases, $\delta_i(\mathbf{x}) = 0$.

Lemma 6.1 *For each $m \in \mathbb{N}$, the set of non-twistable paths in X which visit exactly m vertices in $X \setminus (H_0 \cup K)$ spans a subcomplex of Q_m^X .*

Proof. Let $\mathbf{x} = (x_0, \dots, x_k)$ be a generator of Q_m^X in degree k . By Proposition 5.9, \mathbf{x} has at least one sticky subpath; let (x_i, \dots, x_j) be a sticky subpath.

Removing any of the vertices x_p for $p < i$ or $p > j$ will return a path still containing the sticky subpath (x_i, \dots, x_j) , provided it does not send the generator to zero. Meanwhile, removing any of the vertices x_{i+1}, \dots, x_{j-1} will return a sticky subpath with one fewer vertex, unless it sends the generator to zero. And since the endpoints x_i and x_j belong to $X \setminus (H_0 \cup K)$, removing either one of them will certainly send the generator to zero. Thus, for every $p \in \{1, \dots, k-1\}$, either $\delta_p(\mathbf{x})$ has a sticky subpath or $\delta_p(\mathbf{x}) = 0$. \square

Exactly same argument proves the same lemma for Y :

Lemma 6.2 *For each $m \in \mathbb{N}$, the set of non-twistable paths in Y which visit exactly m vertices in $Y \setminus (H_0 \cup K)$ spans a subcomplex of Q_m^Y .* \square

At heart of the main theorem is the following proposition, which says that the subcomplexes spanned by the non-twistable paths are acyclic. The proof is given in Section 7.

Proposition 6.3 *Let E_m^X denote the subcomplex of Q_m^X spanned by non-twistable paths, and define E_m^Y similarly for Y . For every $m \in \mathbb{N}$, the complexes E_m^X and E_m^Y are both acyclic.* \square

Corollary 6.4 *In each length grading ℓ we have $\chi(Q_m^{X,\ell}) = \chi(Q_m^{Y,\ell})$.*

Proof. Proposition 6.3 tells us that for each m and ℓ the subcomplex $E_m^{X,\ell}$ of $Q_m^{X,\ell}$ is acyclic; consequently, its Euler characteristic vanishes. This lets us write

$$\chi(Q_m^{X,\ell}) = \chi(Q_m^{X,\ell}) - \chi(E_m^{X,\ell})$$

and thus, counting generators,

$$\begin{aligned} \chi(Q_m^{X,\ell}) &= \sum_{k=0}^{\infty} (-1)^k \# \left\{ \begin{array}{l} k\text{-paths of length } \ell \text{ in } X \text{ which visit} \\ \text{exactly } m \text{ vertices in } X \setminus (H_0 \cup K) \end{array} \right\} \\ &\quad - \sum_{k=0}^{\infty} (-1)^k \# \left\{ \begin{array}{l} \text{non-twistable } k\text{-paths of length } \ell \text{ in } X \text{ which} \\ \text{visit exactly } m \text{ vertices in } X \setminus (H_0 \cup K) \end{array} \right\} \\ &= \sum_{k=0}^{\infty} (-1)^k \# \left\{ \begin{array}{l} \text{twistable } k\text{-paths of length } \ell \text{ in } X \text{ which} \\ \text{visit exactly } m \text{ vertices in } X \setminus (H_0 \cup K) \end{array} \right\}. \end{aligned} \quad (13)$$

By the same reasoning,

$$\chi(Q_m^{Y,\ell}) = \sum_{k=0}^{\infty} (-1)^k \# \left\{ \begin{array}{l} \text{twistable } k\text{-paths of length } \ell \text{ in } Y \text{ which} \\ \text{visit exactly } m \text{ vertices in } Y \setminus (H_0 \cup K) \end{array} \right\}. \quad (14)$$

Since τ_G and τ_H each fix every vertex outside K and map K to itself, both functions

preserve the vertex set $H_0 \cup K$ and its complement. Consequently, the bijection

$$T: \{\text{twistable } k\text{-paths in } X\} \rightarrow \{\text{twistable } k\text{-paths in } Y\}$$

constructed in Proposition 5.6 preserves, as well as the length of a path, the number of vertices it visits outside $H_0 \cup K$. Using T to compare each summand in (13) and (14), we can conclude that $\chi(Q_m^{X,\ell}) = \chi(Q_m^{Y,\ell})$. \square

We are now equipped to prove the main theorem.

Theorem 6.5 *Let X and Y be graphs which differ by a sycamore twist. Then*

$$\text{Mag}(X) = \text{Mag}(Y).$$

Proof. By Theorem 2.7 we have

$$\text{Mag}(X) = \sum_{\ell \in \mathbb{N}} \chi(MC_{\bullet}^{\ell}(X))q^{\ell} \quad (15)$$

and

$$\text{Mag}(Y) = \sum_{\ell \in \mathbb{N}} \chi(MC_{\bullet}^{\ell}(Y))q^{\ell}. \quad (16)$$

As the filtration

$$0 \subseteq T_0^X \subseteq T_1^X \subseteq \cdots T_m^X \subseteq \cdots \subseteq MC(X)$$

and the similar filtration of $MC(Y)$ both stabilize, in any given length grading ℓ , after the ℓ^{th} term, the additivity of Euler characteristic implies that we can calculate the ℓ^{th} coefficient in (15) as

$$\chi(MC_{\bullet}^{\ell}(X)) = \chi(T_{\ell}^{X,\ell}) = \sum_{k=0}^{\ell} \chi(Q_k^{X,\ell})$$

the ℓ^{th} coefficient in (16) as

$$\chi(MC_{\bullet}^{\ell}(Y)) = \chi(T_{\ell}^{Y,\ell}) = \sum_{k=0}^{\ell} \chi(Q_k^{Y,\ell}).$$

Now Corollary 6.4 tells us that for every $k, \ell \in \mathbb{N}$ we have $\chi(Q_k^{X,\ell}) = \chi(Q_k^{Y,\ell})$ and hence $\chi(MC_{\bullet}^{\ell}(X)) = \chi(MC_{\bullet}^{\ell}(Y))$ for every ℓ . The theorem follows on comparing coefficients in (15) and (16). \square

7. Proof of Proposition 6.3

It remains to prove that for each $m \in \mathbb{N}$ the subcomplex of Q_m^X spanned by the non-twistable paths is acyclic. The main idea of this proof is adapted from Hepworth and Willerton's Lemma 9.2 in [10], a central component of their excision theorem: it exploits the projection property of vertices in H_* to construct a contracting chain homotopy.

Essentially, given a non-twistable generator \mathbf{x} whose first sticky subpath (x_i, \dots, x_j) **crosses from H_* to $G \setminus K$** —meaning that x_i belongs to H_* and x_j belongs to $G \setminus K$ —the homotopy s inserts the vertex $\pi(x_i)$ immediately after x_i . By definition of stickiness, the vertex x_{i+1} must belong to G , so $\pi(x_j)$ lies between x_i and x_{i+1} and thus we have

$$\delta_{i+1}(s(\mathbf{x})) = \mathbf{x}. \quad (17)$$

If, instead, the first sticky subpath in \mathbf{x} crosses from $G \setminus K$ to H_* , then s inserts $\pi(x_j)$ immediately before x_j , and we have

$$\delta_j(s(\mathbf{x})) = \mathbf{x}. \quad (18)$$

Ultimately the fact that s is a contraction comes down to equations (17) and (18).

Proposition 6.3 *Let E_m^X denote the subcomplex of Q_m^X spanned by non-twistable paths, and define E_m^Y similarly. For every $m \in \mathbb{N}$, the complexes E_m^X and E_m^Y are both acyclic.*

We will prove the statement for E_m^X . The same proof goes through for E_m^Y , since it depends only on the common properties of X and Y established in Section 4—in particular, the fact that vertices in H_* project, in both X and Y , to G .

Proof. First, observe that $E_m^X \cong E_m^X(G, H) \oplus E_m^X(H, G)$ where $E_m^X(G, H)$ is spanned by paths whose first sticky subpath crosses from $G \setminus K$ to H_* , and $E_m^X(H, G)$ is spanned by paths whose first sticky subpath crosses from H_* to $G \setminus K$. We are going to establish that for every $m \in \mathbb{N}$ the complexes $E_m^X(G, H)$ and $E_m^X(H, G)$ are both acyclic. We prove it in full for $E_m^X(H, G)$ and sketch the proof for $E_m^X(G, H)$.

Given a generator \mathbf{x} of $E_m^X(H, G)$, denote the index of the first point of its first sticky subpath by $F_{\mathbf{x}}$. For each $i \geq 0$, let $N_{\bullet}(i)$ be the subcomplex of $E_m^X(H, G)$ spanned by paths \mathbf{x} such that $F_{\mathbf{x}} \leq i$. (That is, such that the subpath (x_0, \dots, x_i) overlaps with the first sticky subpath.) For every i we have $N_k(i) = E_{m,k}^X(H, G)$ for all $k \leq i$, and for each i there is an inclusion $N_{\bullet}(i) \subseteq N_{\bullet}(i+1)$. Indeed, we can filter $E_m^X(H, G)$ as

$$0 = N_{\bullet}(-1) \subseteq N_{\bullet}(0) \subseteq N_{\bullet}(1) \subseteq \dots \subseteq N_{\bullet}(i) \subseteq \dots \subseteq E_m^X(H, G)$$

and again this filtration stabilizes in every length grading. We will show that for each $i \geq 0$ the quotient $N_{\bullet}(i)/N_{\bullet}(i-1)$ is contractible. Thus, on passing to homology, each of these inclusions becomes an isomorphism and we can conclude that $E_m^X(H, G)$ is acyclic.

The complex $N_{\bullet}(i)/N_{\bullet}(i-1)$ is freely generated in degree k by paths $\mathbf{x} = (x_0, \dots, x_k)$ such that $F_{\mathbf{x}} = i$; that is, paths whose first sticky subpath begins with the i^{th} vertex. To understand the boundary operator, notice that removing any one of the vertices x_0, \dots, x_{i-1} from \mathbf{x} yields a path whose first sticky subpath begins with the $(i-1)^{\text{th}}$ vertex; this goes to zero in the quotient. Since x_i is the endpoint of a sticky subpath, it belongs to $X \setminus (K \cup H_0)$, so removing it also sends the generator to zero. Thus, $d = \sum_{j=i+1}^{k-1} (-1)^j \delta_j$ where δ_j removes x_j unless x_j belongs to $X \setminus (K \cup H_0)$ or removing it reduces the length of the path; in either of those cases $\delta_j(\mathbf{x}) = 0$.

Define a map $s: N_k(i)/N_k(i-1) \rightarrow N_{k+1}(i)/N_{k+1}(i-1)$ on generators (x_0, \dots, x_k) as follows. As x_i belongs to H_* , it projects to G in X . Let $\pi(x_i)$ denote the gluing

vertex it projects through, and put

$$s(x_0, \dots, x_k) = \begin{cases} (-1)^{i+1}(x_0, \dots, x_i, \pi(x_i), x_{i+1}, \dots, x_k) & \text{if } x_{i+1} \neq \pi(x_i) \\ 0 & \text{if } x_{i+1} = \pi(x_i). \end{cases}$$

I claim that this is a contracting chain homotopy: that $s \circ d + d \circ s$ is the identity on $N_k(i)/N_k(i-1)$, or equivalently, for every generator $\mathbf{x} = (x_0, \dots, x_k)$ we have

$$\sum_{j=i+1}^{k-1} (-1)^j s \delta_j^k(\mathbf{x}) + \sum_{j=i+1}^k (-1)^j \delta_j^{k+1} s(\mathbf{x}) = \mathbf{x}. \quad (19)$$

For $j = i+2, \dots, k-1$ we have $s \delta_j^k(\mathbf{x}) = \delta_{j+1}^{k+1} s(\mathbf{x})$, so most terms on the left of (19) cancel, leaving

$$(-1)^{i+1} s \delta_{i+1}^k(\mathbf{x}) + (-1)^{i+1} \delta_{i+1}^{k+1} s(\mathbf{x}) + (-1)^{i+2} \delta_{i+2}^{k+1} s(\mathbf{x}). \quad (20)$$

If $x_{i+1} = \pi(x_i)$, the second and third terms of (20) vanish, leaving

$$\begin{aligned} (-1)^{i+1} s \delta_{i+1}^k(x_0, \dots, x_k) &= (-1)^{i+1} s(x_0, \dots, x_i, x_{i+2}, \dots, x_k) \\ &= (-1)^{2(i+1)}(x_0, \dots, x_i, \pi(x_i), x_{i+2}, \dots, x_k) \\ &= (x_0, \dots, x_i, x_{i+1}, x_{i+2}, \dots, x_k). \end{aligned}$$

Here, the first equation holds since x_{i+2} is in G —it's either a gluing vertex, or it's the other end of the sticky subpath—so $x_{i+1} = \pi(x_i)$ lies between it and x_i .

Suppose $x_{i+1} \neq \pi(x_i)$; then x_{i+1} may be a gluing vertex, or it may be the other end of the sticky subpath. In the latter case it belongs to $G \setminus K$, so $\delta_{i+1}^k(\mathbf{x}) = 0 = \delta_{i+2}^{k+1} s(\mathbf{x})$ and (20) reduces to

$$\begin{aligned} (-1)^{i+1} \delta_{i+1}^{k+1} s(x_0, \dots, x_k) &= (-1)^{2(i+1)} \delta_{i+1}^{k+1}(x_0, \dots, x_i, \pi(x_i), x_{i+1}, \dots, x_k) \\ &= (x_0, \dots, x_i, x_{i+1}, \dots, x_k) \end{aligned}$$

where the second equation holds because $x_{i+1} \in G$.

Finally, suppose x_{i+1} is a gluing vertex (but not the vertex $\pi(x_i)$); then x_{i+2} is either a gluing vertex or an element of $G \setminus K$. There are two cases to consider: either $x_{i+1} \in K$ lies between $x_{i+2} \in G$ and $x_i \in H_*$, or it doesn't. In the former case, by Lemma 4.8 and Lemma 4.6, x_{i+1} must also lie between x_{i+2} and $\pi(x_i)$; in the latter case, it must not. Thus, in the former case we have

$$\begin{aligned} &(-1)^{i+1} s \delta_{i+1}^k(\mathbf{x}) + (-1)^{i+1} \delta_{i+1}^{k+1} s(\mathbf{x}) + (-1)^{i+2} \delta_{i+2}^{k+1} s(\mathbf{x}) \\ &= (-1)^{i+1} s(x_0, \dots, x_i, x_{i+2}, \dots, x_k) \\ &\quad + (-1)^{2(i+1)} \delta_{i+1}^{k+1}(x_0, \dots, x_i, \pi(x_i), x_{i+1}, \dots, x_k) \\ &\quad + (-1)^{2i+3} \delta_{i+2}^{k+1}(x_0, \dots, x_i, \pi(x_i), x_{i+1}, \dots, x_k) \\ &= (-1)^{2i+2}(x_0, \dots, x_i, \pi(x_i), x_{i+2}, \dots, x_k) \\ &\quad + (-1)^{2(i+1)}(x_0, \dots, x_i, x_{i+1}, \dots, x_k) \\ &\quad + (-1)^{2i+3}(x_0, \dots, x_i, \pi(x_i), x_{i+2}, \dots, x_k) \\ &= (x_0, \dots, x_k) \end{aligned}$$

as required, while in the latter case we have $\delta_{i+1}^k(\mathbf{x}) = 0 = \delta_{i+2}^{k+1} s(\mathbf{x})$, and (20) reduces

again to

$$\begin{aligned} (-1)^{i+1} \delta_{i+1}^{k+1} s(x_0, \dots, x_k) &= (-1)^{2(i+1)} \delta_{i+1}^{k+1} (x_0, \dots, x_i, \pi(x_i), x_{i+1}, \dots, x_k) \\ &= (x_0, \dots, x_k). \end{aligned}$$

In every case equation (19) holds, and we can conclude that the complex $E_m^X(H, G)$ is acyclic.

The proof for $E_m^X(G, H)$ is very similar; we will sketch the set-up to highlight where things differ. This time, given a generator \mathbf{x} , denote the index of the *last* point of its first sticky subpath by $L_{\mathbf{x}}$, and for each $i \geq 0$ let $M_{\bullet}(i)$ be the subcomplex of $E_m^X(G, H)$ spanned by paths \mathbf{x} such that $L_{\mathbf{x}} \geq k - i$. (That is, such that the subpath (x_{k-i}, \dots, x_k) overlaps with the first sticky subpath.) Thus we obtain a filtration

$$0 = M_{\bullet}(-1) \subseteq M_{\bullet}(0) \subseteq M_{\bullet}(1) \subseteq \dots \subseteq M_{\bullet}(i) \subseteq \dots \subseteq E_m^X(G, H), \quad (21)$$

which stabilizes in every length grading as before.

Now, for each i , define a map $s: M_k(i)/M_k(i-1) \rightarrow M_{k+1}(i)/M_{k+1}(i-1)$ on generators (x_0, \dots, x_k) as follows. As x_{k-i} belongs to H_* , it projects to G in X . Let $\pi(x_{k-i})$ denote the gluing vertex it projects through, and put

$$s(x_0, \dots, x_k) = \begin{cases} (-1)^{k-i} (x_0, \dots, x_{k-i-1}, \pi(x_{k-i}), x_{k-i}, \dots, x_k), & x_{k-i-1} \neq \pi(x_{k-i}) \\ 0, & x_{k-i-1} = \pi(x_{k-i}). \end{cases}$$

An argument mirroring the one above proves that s is a contracting chain homotopy. Thus, on passing to homology, each inclusion in (21) becomes an isomorphism, and we can conclude that $E_m^X(G, H)$ is acyclic. \square

References

- [1] Yasuhiko Asao. Magnitude homology of geodesic metric spaces with an upper curvature bound. *Algebr. Geom. Topol.*, 21:647–664, 2021.
- [2] Yasuhiko Asao. Magnitude homology and path homology. *Bull. Lond. Math. Soc.*, 55(1):375–398, 2023.
- [3] Juan Antonio Barceló and Anthony Carbery. On the magnitudes of compact sets in Euclidean spaces. *Amer. J. Math.*, 140(2):449–494, 2018.
- [4] Rémi Bottinelli and Tom Kaiser. Magnitude homology, diagonality, and median spaces. *Homol. Homotopy Appl.*, 23(2):121–140, 2021.
- [5] Heiko Gimperlein and Magnus Goffeng. On the magnitude function of domains in Euclidean space. *Amer. J. Math.*, 143:939–967, 2021.
- [6] Heiko Gimperlein and Magnus Goffeng. The Willmore energy and the magnitude of Euclidean domains. *Proc. Amer. Math. Soc.*, 151:897–906, 2023.
- [7] Heiko Gimperlein, Magnus Goffeng, and Nikoletta Louca. The magnitude and spectral geometry. Preprint [arXiv:2201.11363](https://arxiv.org/abs/2201.11363), 2022.
- [8] Kiyonori Gomi. Magnitude homology of geodesic space. Preprint [arXiv:1902.07044](https://arxiv.org/abs/1902.07044), 2019.
- [9] Yuzhou Gu. Graph magnitude homology via algebraic Morse theory. Preprint [arXiv:1809.07240](https://arxiv.org/abs/1809.07240), 2018.

- [10] Richard Hepworth and Simon Willerton. Categorifying the magnitude of a graph. *Homol. Homotopy Appl.*, 19:31–60, 2017.
- [11] Ryuki Kaneta and Masahiko Yoshinaga. Magnitude homology of metric spaces and order complexes. *Bull. Lond. Math. Soc.*, 53(3):893–905, 2021.
- [12] William Lawvere. Metric spaces, generalized logic, and closed categories. *Rend. Semin. Mat. Fis. Milano*, 43:135–166, 1974. Reprinted as *Repr. Theory Appl. Categ.* 1:1–37, 2002.
- [13] Tom Leinster. The Euler characteristic of a category. *Doc. Math.*, 13:21–49, 2008.
- [14] Tom Leinster. The magnitude of metric spaces. *Doc. Math.*, 18:857–905, 2013.
- [15] Tom Leinster. The magnitude of a graph. *Math. Proc. Cambridge Philos. Soc.*, 166:247–264, 2019.
- [16] Tom Leinster. *Entropy and Diversity: The Axiomatic Approach*. Cambridge University Press, 2021.
- [17] Tom Leinster and Christina A. Cobbold. Measuring diversity: the importance of species similarity. *Ecology*, 93(3):477–489, 2012.
- [18] Tom Leinster and Mark Meckes. Maximizing diversity in biology and beyond. *Entropy*, 18(3), 2016.
- [19] Tom Leinster and Mark Meckes. The magnitude of a metric space: From category theory to geometric measure theory. In Nicola Gigli, editor, *Measure Theory in Non-Smooth Spaces*. de Gruyter Open, 2017.
- [20] Tom Leinster and Emily Roff. The maximum entropy of a metric space. *Q. J. Math.*, 72(4):1271–1309, 2021.
- [21] Tom Leinster and Michael Shulman. Magnitude homology of enriched categories and metric spaces. *Algebr. Geom. Topol.*, 21:1001–1047, 2021.
- [22] Tom Leinster and Simon Willerton. On the asymptotic magnitude of subsets of Euclidean space. *Geom. Dedicata*, 164:287–310, 2013.
- [23] Mark Meckes. Magnitude, diversity, capacities, and dimensions of metric spaces. *Potential Anal.*, 42(2):549–572, 2015.
- [24] Mark Meckes. On the magnitude and intrinsic volumes of a convex body in Euclidean space. *Mathematika*, 66:343–355, 2020.

Emily Roff emily.roff@ed.ac.uk

The School of Mathematics, The University of Edinburgh, James Clerk Maxwell Building, Peter Guthrie Tait Road, Edinburgh EH9 3FD, Scotland