New invariants of metric spaces: magnitude and maximum entropy

Tom Leinster

Edinburgh

Republic of Scotland

Friday the 13th

Plan

1. The big categorical picture
2. The magnitude of a metric space
3. The magnitude homology of a metric space
4. The maximum entropy of a metric space

1. The big categorical picture

The idea

For many types of mathematical object, there is a canonical notion of size.

- Sets have cardinality. It satisfies

$$|S \cup T| = |S| + |T| - |S \cap T|$$
$$|S \times T| = |S| \times |T|.$$

- Subsets of $\mathbb{R}^n$ have volume. It satisfies

$$\mathrm{vol}(S \cup T) = \mathrm{vol}(S) + \mathrm{vol}(T) - \mathrm{vol}(S \cap T)$$
$$\mathrm{vol}(S \times T) = \mathrm{vol}(S) \times \mathrm{vol}(T).$$

- Topological spaces have Euler characteristic. It satisfies

$$\chi(S \cup T) = \chi(S) + \chi(T) - \chi(S \cap T) \quad \text{(under hypotheses)}$$
$$\chi(S \times T) = \chi(S) \times \chi(T).$$

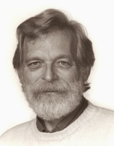

Stephen Schanuel:
Euler characteristic is the topological analogue of cardinality.

The idea

For many types of mathematical object, there is a canonical notion of size.

- Sets have cardinality. It satisfies

$$
\begin{aligned}
|S \cup T| &= |S| + |T| - |S \cap T| \\
|S \times T| &= |S| \times |T|.
\end{aligned}
$$

- Subsets of $\mathbb{R}^n$ have volume. It satisfies

$$
\begin{aligned}
\mathrm{vol}(S \cup T) &= \mathrm{vol}(S) + \mathrm{vol}(T) - \mathrm{vol}(S \cap T) \\
\mathrm{vol}(S \times T) &= \mathrm{vol}(S) \times \mathrm{vol}(T).
\end{aligned}
$$

- Topological spaces have Euler characteristic. It satisfies

$$
\begin{aligned}
\chi(S \cup T) &= \chi(S) + \chi(T) - \chi(S \cap T) \quad \text{(under hypotheses)} \\
\chi(S \times T) &= \chi(S) \times \chi(T).
\end{aligned}
$$

Challenge Find a general definition of 'size', including these and other examples.

One answer The magnitude of an enriched category.

The magnitude of a matrix

Let Z be a matrix.

If Z is invertible, the magnitude of Z is

$$|Z| = \sum_{i,j} (Z^{-1})_{ij}$$

—the sum of all the entries of Z^{-1}.

(The definition can be extended to many non-invertible matrices...
but we won't need this refinement today.)

Enriched categories (informally)

A monoidal category is a category **V** equipped with some kind of product.

A category enriched in **V** is like an ordinary category, with a set/class of objects, but the 'hom-sets' $\mathrm{Hom}(A, B)$ are now objects of **V**.

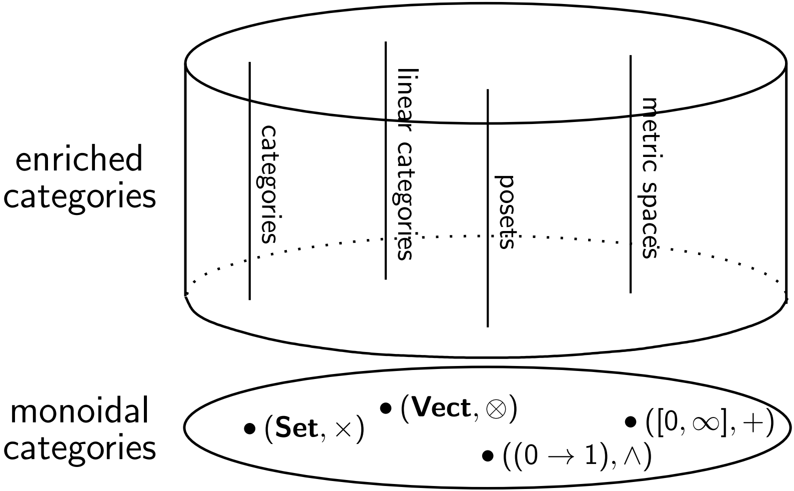

The magnitude of an enriched category (informally)

Let $\mathbf{V}$ be a monoidal category.

Suppose we have a notion of the 'size' of each object of $\mathbf{V}$:
a multiplicative function $|\cdot|$ from ob $\mathbf{V}$ to some field k.

E.g. $\mathbf{V} = \mathbf{FinSet}$, $k = \mathbb{Q}$, $|\cdot| =$ cardinality;
$\mathbf{V} = \mathbf{FDVect}$, $k = \mathbb{Q}$, $|\cdot| =$ dimension.

Then we get a notion of the 'size' of a category $\mathbf{A}$ enriched in $\mathbf{V}$:

- write $Z_{\mathbf{A}}$ for the matrix $\big(|\mathrm{Hom}(A, B)|\big)_{A,B \in \mathrm{ob}\,\mathbf{A}}$ over k
- define the magnitude of the enriched category $\mathbf{A}$ to be

$$|\mathbf{A}| = |Z_{\mathbf{A}}| \in k$$

—i.e. the magnitude of the matrix $Z_{\mathbf{A}}$.

(Here assume $\mathbf{A}$ has only finitely many objects and $Z_{\mathbf{A}}$ is invertible.)

Examples not involving metric spaces

Ordinary finite categories (i.e. $\mathbf{V} = \mathbf{FinSet}$):

- For a finite category $\mathbf{A}$ satisfying mild conditions, $|\mathbf{A}|$ is $\chi(B\mathbf{A}) \in \mathbb{Z}$, the Euler characteristic of the classifying space of $\mathbf{A}$.
- For a finite group G seen as a one-object category, $|G| = 1/\mathrm{order}(G)$.
- For a finitely triangulated manifold X, its poset $\mathbf{A}$ of simplices has magnitude $|\mathbf{A}| = \chi(X) \in \mathbb{Z}$.
- For a finitely triangulated *orbifold* X, its *category* $\mathbf{A}$ of simplices has magnitude $|\mathbf{A}| = \chi(X) \in \mathbb{Q}$. (Joint result with Ieke Moerdijk.)

Linear categories (i.e. $\mathbf{V} = \mathbf{Vect}$):

- For a suitably finite associative algebra E, let $\mathbf{IP}(E)$ denote the linear category of indecomposable projective E-modules.
  Then the magnitude of $\mathbf{IP}(E)$ is a certain Euler form associated with E.
  (Joint result with Joe Chuang and Alastair King.)

Metric spaces as enriched categories

There's at least an *analogy* between categories and metric spaces:

A category has:	A metric space has:
objects $a, b, \ldots$	points $a, b, \ldots$
sets $\mathrm{Hom}(a, b)$	numbers $d(a, b)$
composition operations	triangle inequalities
$\quad \mathrm{Hom}(a, b) \times \mathrm{Hom}(b, c) \to \mathrm{Hom}(a, c)$	$\quad d(a, b) + d(b, c) \geq d(a, c)$

In fact, both are special cases of the concept of enriched category.

(A metric space is a category enriched in the poset $([0, \infty], \geq)$ with $\otimes = +$.)

2. The magnitude of a metric space

The magnitude of a finite metric space (concretely)

To compute the magnitude of a finite metric space $A = \{a_1, \ldots, a_n\}$:

- write down the $n \times n$ matrix with (i, j)-entry $e^{-d(a_i, a_j)}$
- invert it
- add up all n^2 entries.

And that's the magnitude $|A|$.

The magnitude of a finite metric space: first examples

- $|\emptyset| = 0$.
- $|\bullet| = 1$.
- $\left|\bullet \xleftrightarrow{\ \ell\ } \bullet\right| = \text{sum of entries of } \begin{pmatrix} e^{-0} & e^{-\ell} \\ e^{-\ell} & e^{-0} \end{pmatrix}^{-1} = \dfrac{2}{1 + e^{-\ell}}$

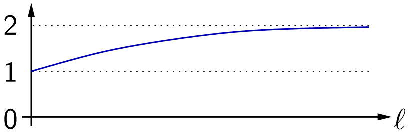

- If $d(a, b) = \infty$ for all $a \neq b$ then $|A| = \text{cardinality}(A)$.

> Slogan: Magnitude is the 'effective number of points'

Magnitude functions

Magnitude assigns to each metric space not just a *number*, but a *function*.

For $t > 0$, write tA for A scaled up by a factor of t.

The magnitude function of a metric space A is the partial function

$$\begin{array}{ccc} (0,\infty) & \to & \mathbb{R} \\ t & \mapsto & |tA|\,. \end{array}$$

E.g.: the magnitude function of $A = (\bullet \xleftarrow{\quad} \ell \xrightarrow{\quad} \bullet)$ is

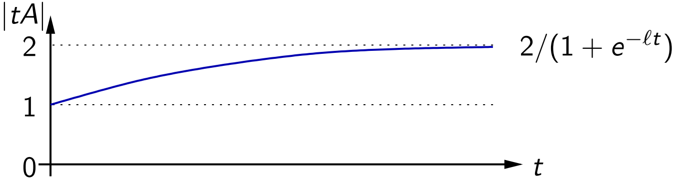

A magnitude function has only finitely many singularities (none if $A \subseteq \mathbb{R}^n$).

It is increasing for $t \gg 0$, and $\lim_{t\to\infty} |tA| = \text{cardinality}(A)$.

The magnitude of a compact metric space

In principle, magnitude is only defined for enriched categories *with finitely many objects* — here, *finite* metric spaces.

Can the definition be extended to, say, compact metric spaces?

Theorem (Mark Meckes)

All sensible ways of extending the definition of magnitude from finite metric spaces to compact 'positive definite' spaces are equivalent.

Proof Uses functional analysis.

Positive definite spaces include all subspaces of $\mathbb{R}^n$ with Euclidean or ℓ^1 (taxicab) metric, and many other common spaces.

The magnitude of a compact positive definite space A is

$$|A| = \sup\{|B| \ : \ \text{finite } B \subseteq A\}.$$

First examples

E.g. Line segment: $|t[0, \ell]| = 1 + \frac{1}{2}\ell \cdot t$.

E.g. Let $A \subseteq \mathbb{R}^2$ be an axis-parallel rectangle with the ℓ^1 (taxicab) metric. Then

$$|tA| = \chi(A) + \tfrac{1}{4}\text{perimeter}(A) \cdot t + \tfrac{1}{4}\text{area}(A) \cdot t^2.$$

Magnitude encodes geometric information

Theorem (Meckes) Let A be a compact subset of $\mathbb{R}^n$, with Euclidean metric.

From the magnitude function of A, you can recover its Minkowski dimension.

Proof Uses a deep theorem from potential analysis, plus the notion of maximum diversity.

Theorem (Willerton) Let A be a homogeneous Riemannian n-manifold. Then as $t \to \infty$,

$$|tA| = a_n \operatorname{vol}(A) \cdot t^n + b_n \operatorname{tsc}(A) \cdot t^{n-2} + O(t^{n-4}),$$

where a_n and b_n are constants and tsc is total scalar curvature.

Proof Uses some asymptotic analysis.

Magnitude encodes geometric information

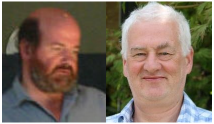

Theorem (Barceló and Carbery) From the magnitude function of A, you can recover the volume of A.

Proof Uses PDEs and Fourier analysis.

Theorem (Barceló and Carbery) The magnitude function of the Euclidean ball B^n (for odd n) is a rational function over $\mathbb{Q}$.

Specifically:

$$\left|tB^1\right| = 1 + t$$

$$\left|tB^3\right| = 1 + 2t + t^2 + \tfrac{1}{6}t^3$$

$$\left|tB^5\right| = \frac{360 + 1080t + 1080t^2 + 525t^3 + 135t^4 + 18t^5 + t^6}{120(3+t)}$$

Magnitude encodes geometric information


Theorem (Gimperlein and Goffeng) From the magnitude function of A, you can recover the surface area of A.

(Needs n odd and some regularity hypotheses.)

Proof Uses heat trace asymptotics (techniques related to heat equation proof of Atiyah–Singer index theorem) and treats t as a *complex* parameter.

Theorem (Gimperlein and Goffeng) Let $A, B \subseteq \mathbb{R}^n$, subject to technical hypotheses. Then

$$|t(A \cup B)| + |t(A \cap B)| - |tA| - |tB| \to 0$$

as $t \to \infty$.

Magnitude of metric spaces doesn't *literally* obey inclusion-exclusion, as that would make it trivial. But it *asymptotically* does.

3. The magnitude homology of a metric space

The idea in brief

Find a homology theory for enriched categories that categorifies magnitude.

This was first done for graphs (seen as metric spaces via shortest paths) by Hepworth and Willerton in 2015: given a graph G,

- they defined a group $H_{n,\ell}(G)$ for all integers $n, \ell \geq 0$ (a *graded* homology theory);
- writing $\chi_\ell(G) = \sum_n (-1)^n \operatorname{rank}(H_{n,\ell}(G))$, the magnitude function of G equals

$$t \mapsto \sum_{\ell} \chi_\ell(G) e^{-\ell t}.$$

So: the Euler characteristic of magnitude homology is magnitude.

The definition was extended to enriched categories in work with Mike Shulman in 2017.

General definition omitted...

The definition for metric spaces (concretely)

In the case of metric spaces A, magnitude homology looks like this.

For each integer $n \geq 0$ and real $\ell \geq 0$, put

$$C_{n,\ell}(A) = \mathbb{Z} \cdot \{(a_0, \ldots, a_n) : d(a_0, a_1) + \cdots + d(a_{n-1}, a_n) = \ell, a_0 \neq \cdots \neq a_n\}.$$

Then $C_{*,\ell}(A)$ is a chain complex for each ℓ, with $\partial = \sum_{0<i<n} (-1)^i \partial_i$ and

$$\partial_i(a_0, \ldots, a_n) = \begin{cases} (a_0, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n) & \text{if } a_i \text{ is between } a_{i-1} \text{ and } a_{i+1} \\ 0 & \text{otherwise.} \end{cases}$$

(Between means $d(a_{i-1}, a_i) + d(a_i, a_{i+1}) = d(a_{i-1}, a_{i+1})$.)

The magnitude homology of A at scale $\ell \in [0, \infty)$ is $H_{*,\ell}(A) = H(C_{*,\ell}(A))$.

Properties of magnitude homology

- For finite metric spaces, magnitude homology categorifies magnitude:

$$|tA| = \sum_{\ell \in [0,\infty)} \chi_\ell(A) e^{-\ell t}$$

(interpreted suitably), where $\chi_\ell(A) = \sum_n (-1)^n \operatorname{rank}(H_{n,\ell}(A))$ as before.

- Magnitude homology detects convexity: for closed $A \subseteq \mathbb{R}^n$,

$$A \text{ is convex} \iff H_{1,\ell}(A) = 0 \text{ for all } \ell > 0.$$

- 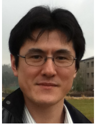 While ordinary homology detects holes, magnitude homology detects the *diameter* of holes (Kaneta and Yoshinaga).

-  There is a precise relationship between magnitude homology and persistent homology—but they detect different information (Otter).

4. The maximum entropy of a metric space

joint with Emily Roff

How spread out is a probability distribution?

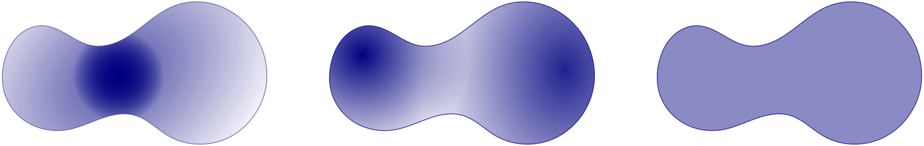

Let A be a compact metric space and μ a (Radon) probability measure on A.
The typicality of a point $a \in A$ is

$$(Z\mu)(a) = \int_A e^{-d(a,b)}\, d\mu(b).$$

It measures how much mass is nearby.

Here a is more typical than b:

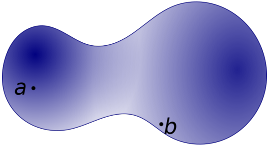

The atypicality of a is $\dfrac{1}{(Z\mu)(a)}$.

How spread out is a probability distribution?

Let A be a compact metric space and μ a probability measure on A.

We quantify spread as the average atypicality of a point in A.

Here 'average' *could* be the ordinary arithmetic mean

$$\int_A \frac{1}{Z\mu}\, d\mu,$$

but it's useful to consider *all* power means:

Definition Let $q \in [0, \infty]$. The diversity of μ of order q is

$$D_q(\mu) = \left(\int_A \left(\frac{1}{Z\mu} \right)^{1-q} d\mu \right)^{\frac{1}{1-q}}$$

taking limits in q for the values $q = 1, \infty$ where this is undefined.

The entropy of μ of order q is $H_q(\mu) = \log D_q(\mu)$.

Special cases

- Let A be a finite set; give it the metric $d(a, b) = \infty$ for all $a \neq b$. Then $H_q(\mu)$ is the Rényi entropy of a probability distribution μ on A, and $H_1(\mu)$ is its Shannon entropy.
- 

  (Joint work with Christina Cobbold.)

  Let A be a finite set, with any metric. Can interpret points of A as species, with distance measured e.g. genetically.

  Then $D_q(\mu)$ measures the diversity of a community with species abundances μ.

  This subsumes many of the biodiversity measures used by ecologists.
- Any compact metric space A, with $q = 2$:

$$D_2(\mu) = \frac{1}{\int_A \int_A e^{-d(a,b)}\, d\mu(a)\, d\mu(b)}.$$

Denominator is expected 'similarity' ($e^{-\text{distance}}$) between a random pair of points.

The role of q

In the definition of diversity $D_q(\mu)$ and entropy $H_q(\mu)$ there is a real parameter q. What does it do?

Example Take $A = \{1, \ldots, 8\}$ with $d(a, b) = \infty$ for all $a \neq b$.

Take μ to be the frequencies of the eight species of great ape on the planet.

Let ν be the 50-50 distribution of chimpanzees and bonobos only.

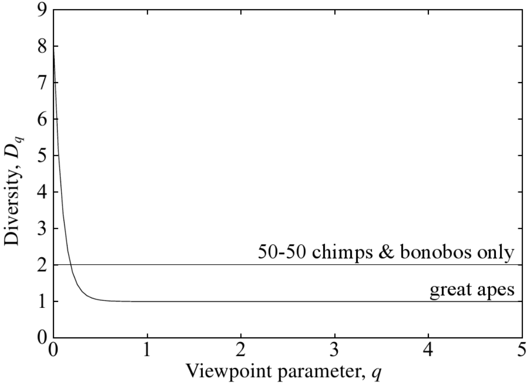

Moral: You can't ask whether one probability measure has higher diversity/entropy than another.

The answer may depend on q.

The maximum diversity theorem

Let A be a compact metric space.

What is the maximum possible diversity (or entropy) achievable by a probability measure on A? What *is* that maximum?

In principle, both answers depend on q.

Theorem (with Mark Meckes [finite case] and Emily Roff [general case])
Both answers are independent of q. That is:

- there is a probability measure μ maximizing $D_q(\mu)$ for all $q \in [0, \infty]$ simultaneously
- $\sup_\mu D_q(\mu)$ is independent of q

(and the same for entropy H_q).

Hence on a compact metric space, we have:

- a canonical probability measure (the maximizer, which is usually unique)
- a canonical real number, $D_{\max}(A) = \sup_\mu D_q(\mu)$.

A little on maximum diversity

Maximum diversity is closely related to magnitude. In fact, for any compact metric space A,

$$D_{\max}(A) = |B|$$

for some closed $B \subseteq A$.

As for magnitude, the large-t asymptotics of $D_{\max}(tA)$ encode geometric information about the space A. E.g.:

- The asymptotic growth rate of $D_{\max}(tA)$ is the Minkowski dimension of A (Meckes).
- For $A \subseteq \mathbb{R}^n$,

  $$\mathrm{vol}(A) = c_n \lim_{t \to \infty} \frac{D_{\max}(tA)}{t^n},$$

  where c_n is a known constant.

But the maximum diversity of even some very simple spaces is unknown, e.g. Euclidean balls of dimension > 1.

The obvious probability measure

Some spaces carry an obvious probability measure—the first one everyone thinks of. E.g.:

- For finite spaces, it's the uniform probability measure.
- For homogeneous spaces, it's Haar probability measure.
- For subsets of $\mathbb{R}^n$, it's normalized Lebesgue measure.

We'll give a formal definition of the 'uniform' measure on a space, capturing all these examples.

Maximizing entropy on an interval

The unique probability measure on $[0, t]$ that maximizes all the entropies H_q (or equivalently all the diversities D_q) is the normalization of

$$\delta_0 + \lambda_{[0,t]} + \delta_t.$$

Here λ_A means the restriction of Lebesgue measure to A.

For large t, the two deltas make no contribution to μ_t.

So 'in the limit', the maximizing measure is just normalized Lebesgue.

Formally: recall that for a metric space $A = (A, d)$, we write tA for the set A with the metric td. (E.g. $t[0, 1] \cong [0, t]$.)

Let μ_t denote the unique entropy-maximizing measure on $t[0, 1]$. Then

$$\lim_{t\to\infty} \mu_t = \text{normalization of } \lambda_{[0,1]}.$$

The uniform measure

Definition Let A be a compact metric space. Suppose that tA has a unique entropy-maximizing measure μ_t for all $t \gg 0$.

The uniform measure on A is $\mu_A = \lim_{t\to\infty} \mu_t$, if it exists.

Examples

- When A is finite, μ_A is the uniform probability measure in the usual sense.
- When A is homogeneous (isometry group acts transitively), μ_A is the Haar probability measure.
- When $A \subseteq \mathbb{R}^n$, with nonzero finite volume, μ_A is normalized Lebesgue.

And the uniform measure is scale-invariant ($\mu_A = \mu_{tA}$).

Summary

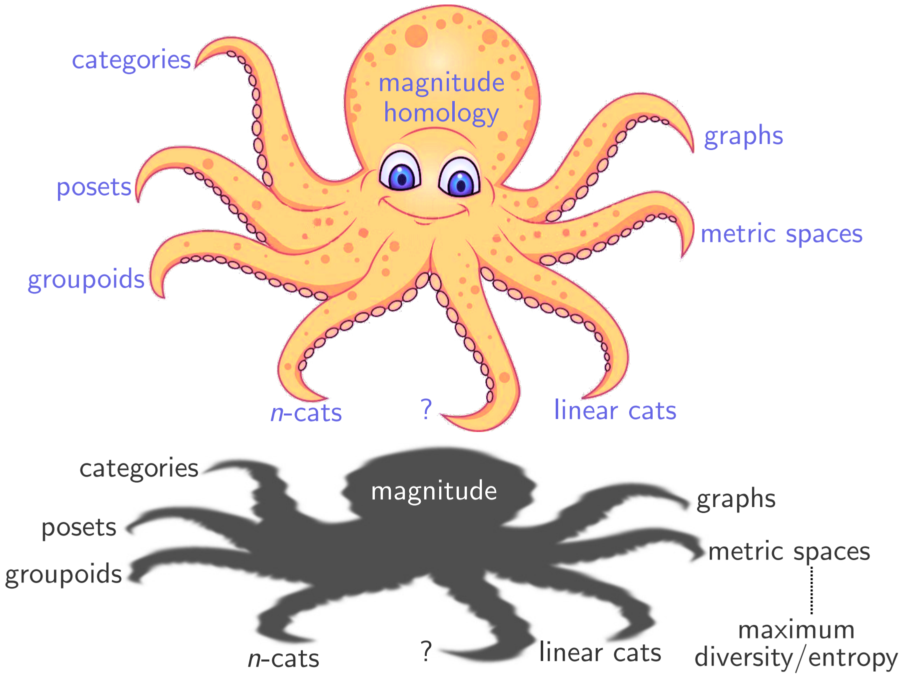

categories
magnitude
homology
graphs
posets
metric spaces
groupoids
n-cats
?
linear cats
categories
magnitude
graphs
posets
metric spaces
groupoids
maximum
diversity/entropy
n-cats
?
linear cats