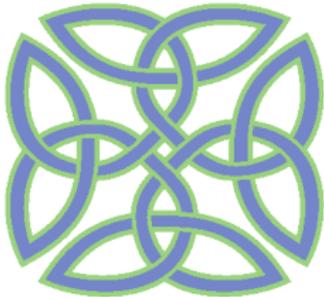


The mathematics of diversity

Tom Leinster

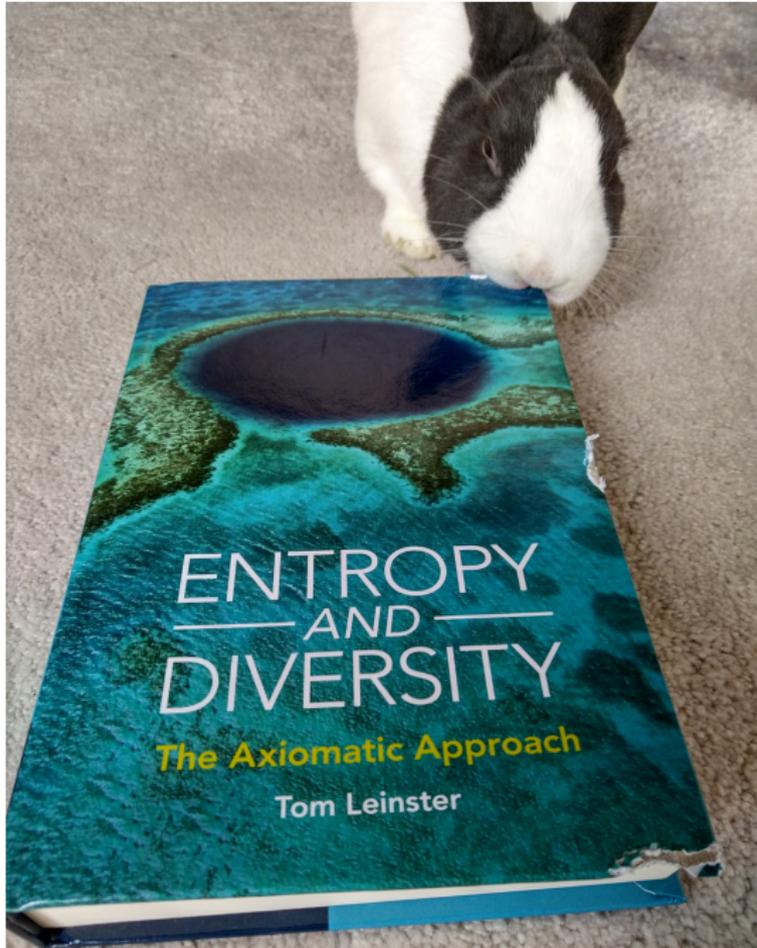


School of Mathematics
University of Edinburgh

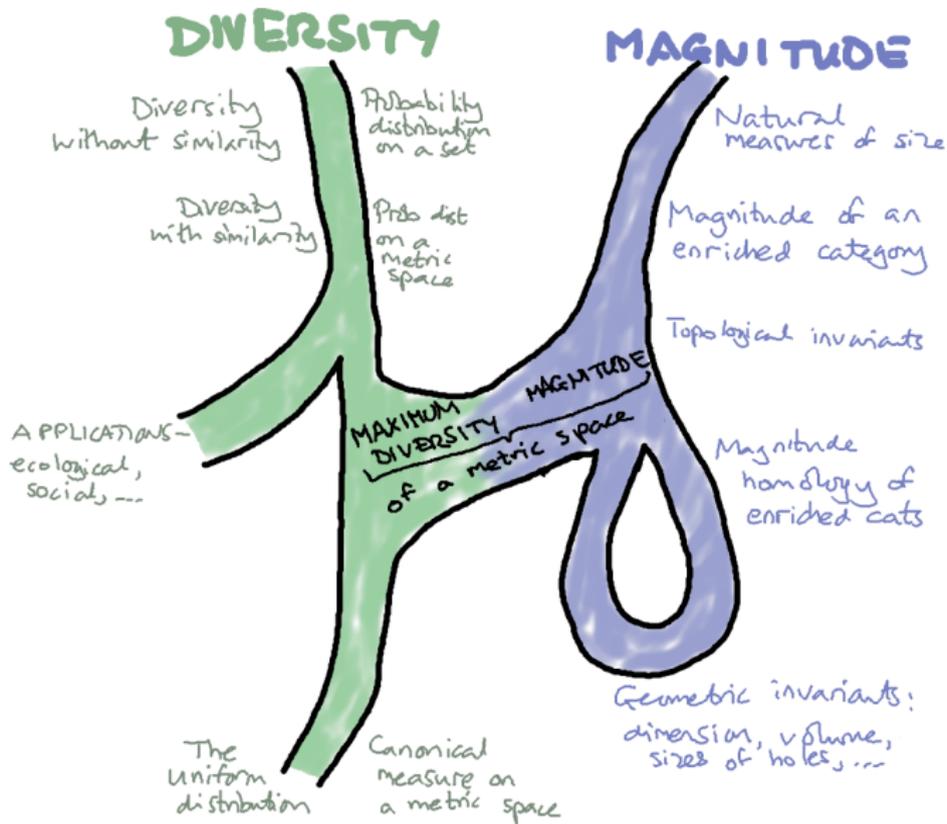


Boyd Orr Centre
for Population and Ecosystem Health
University of Glasgow

Main reference



Plan



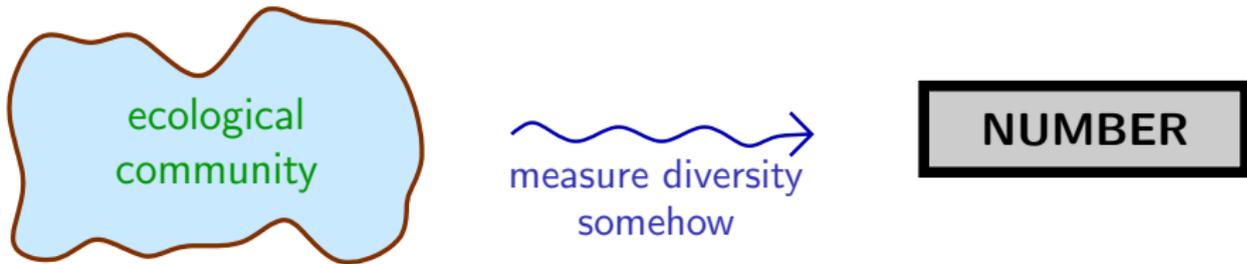
Just the left-hand half today. (Right-hand half [here](#).)

Recurrent question

What is the $\left\{ \begin{array}{l} \text{unique} \\ \text{universal} \\ \text{canonical} \end{array} \right\}$ such-and-such?

What is the best/canonical way to measure diversity?

Basic challenge:



Take a community divided into species.

Crudest diversity measure: the number of species present.

But this is often misleading.

Example There are 8 species of great ape on the planet. . .

. . . but 99.99% of ape individuals are from a single species.

Ecologists have proposed and used many, many diversity measures...

Whittaker's index
of association

Percentage difference
(*alias* Bray-Curtis)

Wishart coefficient =
(1 - similarity ratio)

D = (1 - Kulczynski
coefficient)

Abundance-based
Jaccard

Abundance-based
Sørensen

Abundance-based
Ochiai

Species richness $x \equiv \sum_{i=1}^S p_i^0$

Shannon entropy $x \equiv - \sum_{i=1}^S p_i \ln p_i$

Simpson concentration $x \equiv \sum_{i=1}^S p_i^2$

Gini-Simpson index $x \equiv 1 - \sum_{i=1}^S p_i^2$

HCDDT entropy $x \equiv \left(1 - \sum_{i=1}^S p_i^q \right) / (q - 1)$

Renyi entropy $x \equiv \left(-\ln \sum_{i=1}^S p_i^q \right) / (q - 1)$

Ecologists have proposed and used many, many diversity measures...

Whittaker's index of association

$${}^q D_{Tj} = 1 / \hat{p}_{(ij)all} \quad \gamma_j = \hat{p}_{(ij)all}^{1/(1-q)}$$

$$\hat{p}_{(ij)all} = \frac{1}{\sqrt{\sum_{i=1}^N \sum_{l=1}^S \rho_{il} \rho_{lj}^{q-1}}}$$

$$\alpha_i = {}^q D_{Tj} = \hat{\gamma}_j \quad 1 / \hat{p}_{(ij)all}$$

$$\alpha_d = {}^q D_{\alpha} \quad \alpha_i / CU$$

$$\alpha_R = {}^q D_{\gamma_{var}/\alpha} \quad {}^q D_{\gamma_{var}} / {}^q D_{\alpha}$$

$$\beta_{Mtd} = {}^q D_{\beta} = {}^q D_{\gamma} / {}^q D_{\alpha} \quad \gamma / \alpha_d$$

$$\beta_{Mtr} = {}^q D_{\gamma/\hat{\gamma}_j} = {}^q D_{\gamma} / {}^q D_{\hat{\gamma}_j} \quad \gamma / \alpha_i$$

$$\beta_R = {}^q D_{\gamma} / {}^q D_{\alpha} / {}^q D_{\gamma_{var}} \quad \gamma / \alpha_R$$

$$\beta_{\alpha_i} = {}^q D_{\gamma} - {}^q D_{\hat{\gamma}_j} \quad \gamma - \alpha_i$$

$$\beta_{Mtr-1} = \gamma / \alpha_i - 1 \quad (\gamma - \alpha_i) / \alpha_i$$

$$\beta_{Pr} = 1 - \alpha_i / \gamma \quad (\gamma - \alpha_i) / \gamma$$

$$H_{\beta}^* = H_{\gamma}^* - H_{\alpha}^* \quad \log^1(\beta_{Mtd}) = \log(\gamma) - \log(\alpha_d)$$

$$\tilde{H}_{\gamma-\hat{\gamma}_j}^* = H_{\gamma}^* - \tilde{H}_{\hat{\gamma}_j}^* \quad \log^1(\beta_{Mtr}) = \log(\gamma) - \log(\alpha_i)$$

$${}^2 \tilde{\lambda}_{\gamma-\hat{\gamma}_j} = {}^2 \tilde{\lambda}_{\gamma} - {}^2 \tilde{\lambda}_{\hat{\gamma}_j} \quad (\gamma - \alpha_i) / \gamma \alpha_i$$

${}^1 E$
 ${}^2 E$
 $D_{q/0}$
 $D_{2/1}$
 $J'_{1/max}$ or $J'_{1/0}$ or ${}^1 J'$

gamma unit j (mean) genera
 mean s
 weight
 true al
 sampli
 (measu
 effectiv
 sampli
 true be
 (measu
 region:
 * ${}^2 H'$
 * ${}^2 D$ or ${}^1 D$
 two-w:
 region:
 (measu
 Whitta
 multipl
 (measu
 propor
 as a pr
 * E_{D_2}
 beta Shann
 entropy me
 measurement
 unit: base of
 the logarithm)
 regional Shannon entropy excess (measurement unit: depends on the base of the logarithm)
 regional variance excess (measurement unit: sp_i/sp_j^2)

Sheldon 1969, Buzas and Gibson 1969, Buzas and H
 McCarthy 2002, Camargo 2008

Weiber and Keddy 1999, Wilsey and Potvin 2000, Mc
 2003, Ma 2005, Martin et al. 2005, Bock et al. 200
 Camargo 2008

Alatalo 1981, Taillie 1979, Patil and Taillie 1982, Ricc
 Rotenberry 1978, Alatalo 1981, Ricotta and Avena 20
 Sheldon 1969, Tramer 1969, Kricher 1972, DeBenedic
 Wills et al. 1997, Rex et al. 2000, Wilsey and Potvi
 Miranda et al. 2002, Woodd-Walker et al. 2002, Ols
 et al. 2005, Kimbro and Grosholz 2006, Wilsey and
 Anticamara et al. 2010, Castro et al. 2010, Kardol
 Hill 1973, Ricotta 2003, Kindt et al. 2006

Hurlbert 1971, Fager 1972, DeJong 1975, Smith and
 Jabot and Chave 2009, Anticamara et al. 2010

Weiber and Keddy 1999, Stevens and Willig 2002, D.
 Vellend 2005, Ulrich and Zalewski 2007, Jarvis et
 Walker and Cyr 2007

Gardezi and Gonzales 2008, Anticamara et al. 2010
 Mouillot and Wilson 2002, Stevens and Willig 2002

Camargo 1992a, 1993, Drobner et al. 1998, Mouillot
 2002, Mouillot and Wilson 2002, Stevens and Will
 Ghera 2011

Drobner et al. 1998, Weiber and Keddy 1999, Mouill
 Symonds and Johnson 2008, Bernhardt-Römerman
 Weiber and Keddy 1999

Drobner et al. 1998, Mouillot and Wilson 2002, Ma z

$$\sum_{i=1}^S P_i / (1 - P_i)$$

A very simple model of an ecological community

Take a community whose organisms are divided into n species.

Let p_i be the relative abundance of the i th species. So $p_1 + \cdots + p_n = 1$.

Write $\mathbf{p} = (p_1, \dots, p_n)$.

Mathematically: a community is a probability distribution on a finite set.

A spectrum of viewpoints

Species

are what matter

Rare species

count for as much

as common ones

—every species is precious

Communities

are what matter

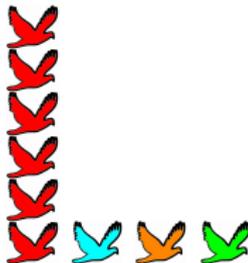
Common species

are the really

important ones

—they shape the community

This →



← This

is more diverse than

is less diverse than

that →



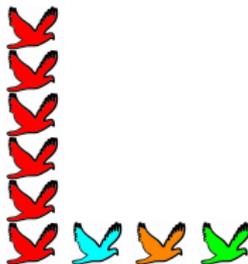
← that

A spectrum of viewpoints

*Rare species are
important*

*Rare species are
unimportant*

This →



← This

is more diverse than

is less diverse than

that →



← that

How to acknowledge the spectrum of viewpoints

In 1973, the ecologist Mark Hill defined a family of diversity measures acknowledging the spectrum of viewpoints.

Let $0 \leq q \leq \infty$. The **diversity** or **Hill number of order q** of the community is

$$D_q(\mathbf{p}) = \left(\sum_i p_i^q \right)^{1/(1-q)}$$

(taking limits to get the definitions for $q = 1, \infty$).

The parameter q controls the relative emphasis placed on rare and common species—in other words, where on the spectrum you are.

Examples

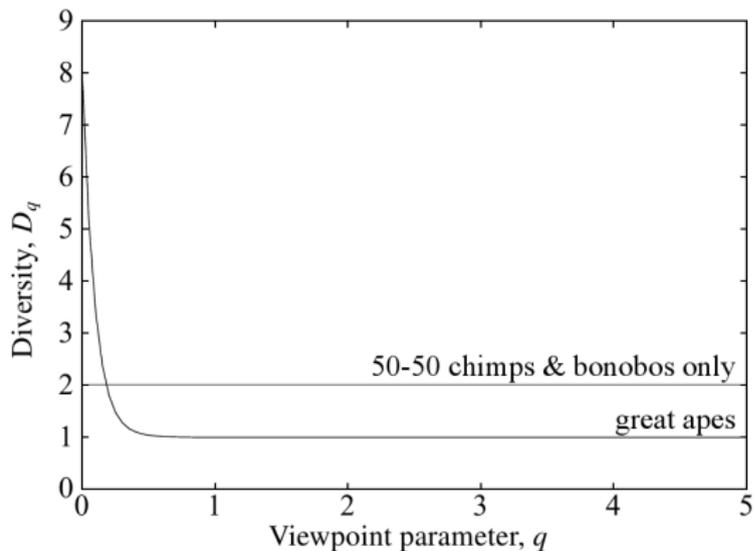
- $D_0(\mathbf{p}) =$ number of species present.
- $D_q(1/n, \dots, 1/n) = n$: ‘there are n species in perfect balance’.

The role of q

In the definition of diversity $D_q(\mathbf{p})$, there is a parameter q . What does it do?

Example Take \mathbf{p} to be the frequencies of the eight species of great ape on the planet.

Or take \mathbf{p} to be the 50-50 distribution of chimpanzees and bonobos only.



Moral: You can't always say whether one distribution has higher diversity than another.

The answer may depend on q .

Digression: entropy

$$\text{entropy} = \log(\text{diversity})$$

The logarithm of the Hill number D_q is the **Rényi entropy** of order q :

$$H_q(\mathbf{p}) = \log D_q(\mathbf{p}).$$

When $q = 1$, this is the **Shannon entropy**.

For diversity, there are good reasons to use the exponential form.

For information theory, there are good reasons to use the logarithmic form.

Unique characterization of the Hill numbers

Any measure of diversity should behave in a logical way that reflects biological etc. intuition.

E.g. Consider a group of islands with disjoint species.



Increasing the diversity of one island should increase the diversity of the whole.

Unique characterization of the Hill numbers

Any measure of diversity should behave in a logical way that reflects biological etc. intuition.

E.g. Consider a group of islands with disjoint species.



Increasing the diversity of one island should increase the diversity of the whole.

Theorem *The only diversity measures satisfying seven sensible properties are the Hill numbers D_q ($0 \leq q \leq \infty$).*

A major shortcoming

Intuitively, diversity should reflect how *different* the species (etc.) are, not just their frequencies.

‘Biological diversity’ means the *variability* among living organisms

—UN Environment Programme definition (quoted in Magurran, *Measuring Biological Diversity*, p.6).

... associated with the idea of diversity is the concept of ‘*distance*’,
i.e. some measure of the dissimilarity of the resources in question

—OECD *Handbook of Biodiversity Valuation: A Guide for Policy Makers*.

A slightly less simple model of an ecological community



Tom Leinster and Christina Cobbold,
Measuring diversity: the importance of species similarity,
Ecology 93 (2012), 477–89.

Assume we also have a measure of the similarity between the i th and j th species,

$$0 \leq Z_{ij} \leq 1.$$

Here $Z_{ij} = 0$ means total dissimilarity, and $Z_{ij} = 1$ means identical species.

This defines an $n \times n$ matrix $Z = (Z_{ij})$.

The similarities Z_{ij} can be determined genetically, phylogenetically, functionally, morphologically, taxonomically, You choose!

- E.g. The **naive model** $Z = I$: different species have nothing in common.
- E.g. Given a metric d on $\{1, \dots, n\}$, put $Z_{ij} = e^{-d(i,j)}$.

Roughly: a community is a probability distribution on a finite metric space.

Similarity-sensitive diversity

How typical is the i th species? Multiply the similarity matrix Z by the abundance vector \mathbf{p} :

$$(Z\mathbf{p})_i = \sum_j Z_{ij}p_j.$$

So the *atypicality* of the i th species can be quantified as

$$1/(Z\mathbf{p})_i.$$

Diversity is defined as the average atypicality of an individual.

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

$$\sum_i p_i \cdot \frac{1}{(Z\mathbf{p})_i},$$

but it's useful to consider *all* power means. So for $0 \leq q \leq \infty$, define

$$D_q^Z(\mathbf{p}) = \left(\sum_i p_i \left(\frac{1}{(Z\mathbf{p})_i} \right)^{1-q} \right)^{1/(1-q)}$$

(taking limits at $q = 1, \infty$).

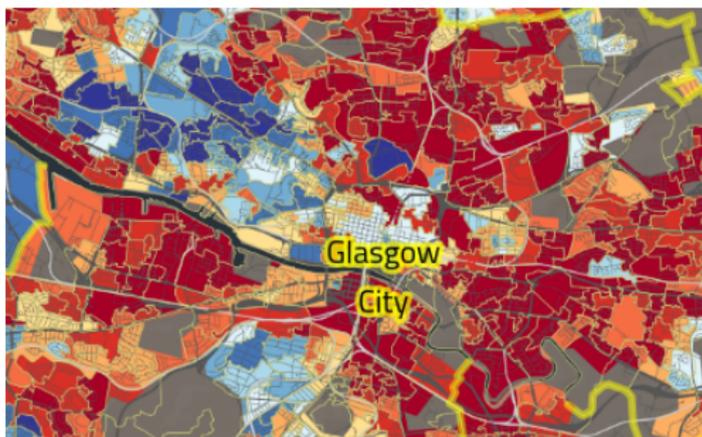
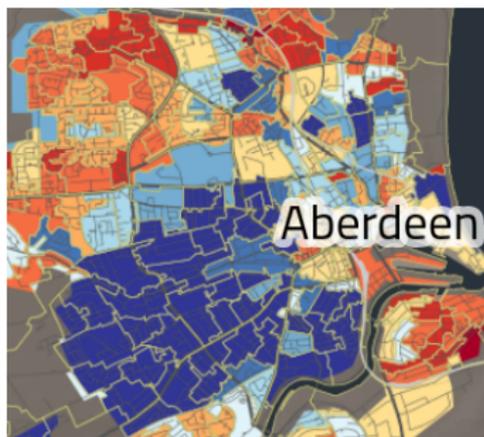
Features of the similarity-sensitive diversity measures

- This definition unifies many of the diversity measures used by ecologists (and elsewhere in the life sciences).
- When $Z = I$ (naive model), we recover the Hill numbers: $D_q^I = D_q$.
- They have been applied at all ecological scales, from microbial to large predators.
- Mathematically: taking logs, we get a notion of the **entropy** of a probability distribution on a metric space.

Digression: social applications

From the **Scottish Index of Multiple Deprivation**.

Red: high deprivation. **Blue:** low deprivation.



How would you quantify the concentration/spread of deprivation?

Can you measure how separated the poor and rich areas are?

A group of Edinburgh undergraduates is working on it. . .

Maximizing diversity



Tom Leinster and Mark Meckes,
Maximizing diversity in biology and beyond,
Entropy 18 (2016), article 18.

Fix a similarity matrix Z , i.e. a list of species with known similarities.

Or if you prefer: fix a finite metric space.

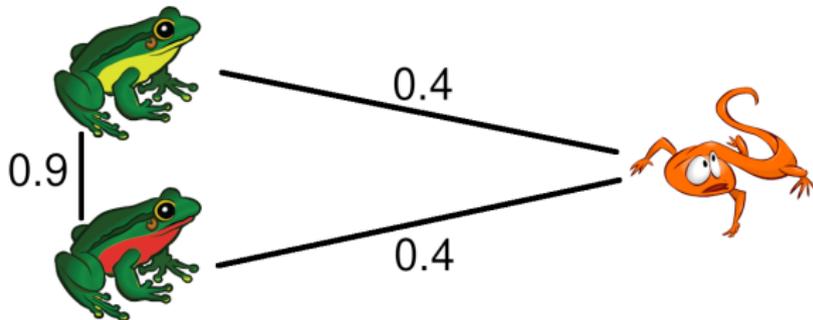
Allow the relative abundance distribution \mathbf{p} to vary.

Which \mathbf{p} achieves the maximum possible diversity? What *is* that maximum?

Example: frogs and newts

Take a three-species system with these similarities:

$$Z = \begin{pmatrix} 1 & 0.9 & 0.4 \\ 0.9 & 1 & 0.4 \\ 0.4 & 0.4 & 1 \end{pmatrix}$$



Which distribution maximizes diversity?

- Not $(1/3, 1/3, 1/3)$, because then we'd have $2/3$ frog and $1/3$ newt.
- Not $(1/4, 1/4, 1/2)$, because the frog species aren't quite identical.
- It should be somewhere in between.

In particular, the maximizing distribution should *not* be uniform.

The maximum diversity theorem

Fix a similarity matrix Z , i.e. a list of species with known similarities.

Or if you prefer: fix a finite metric space.

Allow the relative abundance distribution \mathbf{p} to vary.

Which \mathbf{p} achieves the maximum possible diversity? What *is* that maximum?

In principle, both answers depend on q .

Theorem (with Mark Meckes)

Both answers are independent of q . That is:

- there is a probability measure \mathbf{p} maximizing $D_q^Z(\mathbf{p})$ for all $q \in [0, \infty]$ simultaneously
- $\sup_{\mathbf{p}} D_q^Z(\mathbf{p})$ is independent of q .

If \mathbf{p} maximizes $D_q^Z(\mathbf{p})$ for all q , we call \mathbf{p} a **maximizing measure**.

It is usually unique, making it a *canonical probability measure on a finite metric space*.

From finite to infinite spaces



Tom Leinster and Emily Roff,
The maximum entropy of a metric space,
Quarterly Journal of Mathematics, to appear.

All this can be done for a general *compact* metric space, not necessarily finite.

For a probability measure \mathbb{P} on a compact metric space X , put

$$Z\mathbb{P}(x) = \int e^{-d(x,y)} d\mathbb{P}(y)$$

and

$$D_q^X(\mathbb{P}) = \left(\int \left(\frac{1}{Z\mathbb{P}(x)} \right)^{1-q} d\mathbb{P}(x) \right)^{1/(1-q)}.$$

The maximum diversity theorem holds for all compact metric spaces.

Example: real interval

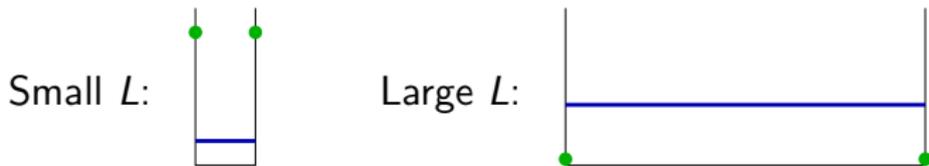
Take a real interval $[0, L]$ of length L with its usual metric.

Which probability measure on $[0, L]$ maximizes diversity?

- One guess: $\frac{1}{2}(\delta_0 + \delta_L)$ (push all the mass to the ends).
- Another guess: the uniform distribution, i.e. the normalization of Lebesgue measure $\lambda_{[0,L]}$.
- In fact, it's a linear combination of these. It's the normalization of

$$\delta_0 + \lambda_{[0,L]} + \delta_L.$$

Warning The maximizing measure is scale-dependent!



What is the uniform distribution?

Given a space X , which probability measure on X deserves to be called the 'uniform distribution' on X ?

The answer is obvious for certain classes of space X . E.g.:

- finite spaces 
- suitably symmetric spaces  : the unique symmetric measure
- subsets of \mathbb{R}^n with $0 < \text{Vol}(X) < \infty$: normalized Lebesgue measure.

And clearly there's no sensible uniform distribution for some X , e.g. \mathbb{R} or \mathbb{Z} .

Is there a good general answer?

What is the uniform distribution on a compact metric space?

- Old idea from statistics: the canonical choice of probability distribution is the one with the maximum entropy.

Should we refer to the maximizing measure as the 'uniform distribution'?

No! It's scale-dependent, and that's bad.

It also gives the wrong answer in examples (e.g. the interval).

- But if we take the large-scale limit, it works. . .

The uniform measure

Let X be a compact metric space.

For $t > 0$, write tX for X with the metric scaled up by a factor of t .

Assume that for $t \gg 0$, there is only one maximizing measure \mathbb{P}_t on tX .

Definition The **uniform measure** on X is

$$\mathbb{U}_X = \lim_{t \rightarrow \infty} \mathbb{P}_t,$$

if this limit exists (in the weak* topology).

This is scale-independent: $\mathbb{U}_{tX} = \mathbb{U}_X$ for all $t > 0$.

The uniform measure is a canonical scale-independent measure on a compact metric space.

Easy examples

- **Finite spaces:** The uniform measure on a finite space is the uniform measure in the usual sense.
- **Symmetric spaces:** Let X be a compact metric space such that for all $x, y \in X$, some self-isometry of X maps x to y .
The Haar measure theorem implies that there is a unique isometry-invariant probability measure on X .
And that's what the uniform measure on X is.

Slightly less easy example: the interval

Consider the real interval $[0, L]$.

We saw that the maximizing measure on $[0, L]$ is the normalization of $\delta_0 + \lambda_{[0,L]} + \delta_L$.

When we scale up by a large factor t , the point masses at the endpoints become negligible.

So the uniform measure on $[0, L]$ is the normalization of Lebesgue measure, $\lambda_{[0,L]}/L$.

That is, it's the uniform distribution in the usual sense.

Much harder case: subsets of \mathbb{R}^n

Let X be a compact subset of \mathbb{R}^n .

Problem Unlike for the interval, we have no description of the maximizing measure on X — even when X is a 2-dimensional disc!

So unlike for the interval, we *can't* find the uniform measure on X by finding the maximizing measure on tX for each finite t , then letting $t \rightarrow \infty$.

Nevertheless, assuming that X has nonzero measure...

Theorem *The uniform measure \mathbb{U}_X is normalized Lebesgue measure on X .*

Proof Lots of analysis ([paper with Emily Roff](#)).

DIVERSITY

MAGNITUDE

Diversity without similarity

Diversity with similarity

Probability distribution on a set

Prob dist on a metric space

Natural measures of size

Magnitude of an enriched category

Topological invariants

APPLICATIONS -
ecological,
social, ...

MAXIMUM DIVERSITY
of a metric space

MAGNITUDE

Magnitude homology of enriched cats

The uniform distribution

Canonical measure on a metric space

Geometric invariants:
dimension, volume,
sizes of holes, ...

Thanks