

Article

Representational Rényi Heterogeneity

Abraham Nunes ^{1,2,*} , Martin Alda ¹ , Timothy Bardouille ³ and Thomas Trappenberg ^{2,*} ¹ Department of Psychiatry, Dalhousie University, Halifax, NS B3H 2E2, Canada; malda@dal.ca² Faculty of Computer Science, Dalhousie University, Halifax, NS B3H 4R2, Canada³ Department of Physics and Atmospheric Sciences, Dalhousie University, Halifax, NS B3H 4R2, Canada; tim.bardouille@dal.ca

* Correspondence: nunes@dal.ca (A.N.); tt@cs.dal.ca (T.T.)

† Current address: 5909 Veterans Memorial Lane (8th Floor), Abbie J. Lane Memorial Building, QE II. Health Sciences Centre, Halifax, NS B3H 2E2, Canada.

Received: 26 March 2020; Accepted: 4 April 2020; Published: 7 April 2020



Abstract: A discrete system's heterogeneity is measured by the Rényi heterogeneity family of indices (also known as Hill numbers or Hannah–Kay indices), whose units are the numbers equivalent. Unfortunately, numbers equivalent heterogeneity measures for non-categorical data require a priori (A) categorical partitioning and (B) pairwise distance measurement on the observable data space, thereby precluding application to problems with ill-defined categories or where semantically relevant features must be learned as abstractions from some data. We thus introduce representational Rényi heterogeneity (RRH), which transforms an observable domain onto a latent space upon which the Rényi heterogeneity is both tractable and semantically relevant. This method requires neither a priori binning nor definition of a distance function on the observable space. We show that RRH can generalize existing biodiversity and economic equality indices. Compared with existing indices on a beta-mixture distribution, we show that RRH responds more appropriately to changes in mixture component separation and weighting. Finally, we demonstrate the measurement of RRH in a set of natural images, with respect to abstract representations learned by a deep neural network. The RRH approach will further enable heterogeneity measurement in disciplines whose data do not easily conform to the assumptions of existing indices.

Keywords: heterogeneity; diversity; Rényi heterogeneity; representation learning; variational autoencoder; functional diversity indices; Hill numbers; Leinster–Cobbold Index; Rao's quadratic entropy

1. Introduction

Measuring heterogeneity is of broad scientific importance, such as in studies of biodiversity (ecology and microbiology) [1,2], resource concentration (economics) [3], and consistency of clinical trial results (biostatistics) [4], to name a few. In most of these cases, one measures the heterogeneity of a discrete system equipped with a probability mass function.

Discrete systems assume that all observations of a given state are identical (zero distance), and that all pairwise distances between states are permutation invariant. This assumption is violated when relative distances between states are important. For example, an ecosystem is not biodiverse if all species serve the same functional role [5]. Although species are categorical labels, their pairwise differences in terms of ecological functions differ and thus violate the discrete space assumptions. Mathematical ecologists have thus developed heterogeneity measures for non-categorical systems, which they generally call “functional diversity indices” [6–11]. These indices typically require a priori discretization and specification of a distance function on the observable space.

The requirement for defining the state space a priori is problematic when the states are incompletely observable: that is, when they may be noisy, unreliable, or invalid. For example,

consider sampling a patient from a population of individuals with psychiatric disorders and assigning a categorical state label corresponding to his or her diagnosis according to standard definitions [12]. Given that psychiatric conditions are not defined by objective biomarkers, the individual's diagnostic state will be uncertain. Indeed, many of these conditions are inconsistently diagnosed across raters [13], and there is no guarantee that they correspond to valid biological processes. Alternatively, it is possible that variation within some categorical diagnostic groups is simply related to diagnostic “noise,” or nuisance variation, but that variation within other diagnostic groups constitutes the presence of sub-strata. Appropriate measurement of heterogeneity in such disciplines requires freedom from the discretization requirement of existing non-categorical heterogeneity indices.

Pre-specified distance functions may fail to capture semantically relevant geometry in the raw feature space. For example, the Euclidean distance between Edmonton and Johannesburg is relatively useless since the straight-line path cannot be traversed. Rather, the appropriate distances between points must account for the data's underlying manifold of support. Representation learning addresses this problem by learning a latent embedding upon which distances are of greater semantic relevance [14]. Indeed, we have observed superior clustering of natural images embedded on Riemannian manifolds [15] (but also see Shao et al. [16]), and preservation of semantic hierarchies when linguistic data are embedded on a hyperbolic space [17].

Therefore, we seek non-categorical heterogeneity indices without requisite a priori definition of categorical state labels or a distance function. The present study proposes a solution to these problems based on the measurement of heterogeneity on learned latent representations, rather than on raw observable data. Our method, representational Rényi heterogeneity (RRH), involves learning a mapping from the space of observable data to a latent space upon which an existing measure (the Rényi heterogeneity [18], also known as the Hill numbers [19] or Hannah–Kay indices [20]) is meaningful and tractable.

The paper is structured as follows. Section 2 introduces the original categorical formulation of Rényi heterogeneity and various approaches by which it has been generalized for application on non-categorical spaces [8,10,21]. Limitations of these indices are highlighted, thereby motivating Section 3, which introduces the theory of Representational Rényi Heterogeneity (RRH), which generalizes the process for computing many indices of biodiversity and economic equality. Section 4 provides an illustration of how RRH may be measured in various analytical contexts. We provide an exact comparison of RRH to existing non-categorical heterogeneity indices under a tractable mixture of beta distributions. To highlight the generalizability of our approach to complex latent variable models, we also provide an evaluation of RRH applied to the latent representations of a handwritten image dataset [22] learned by a variational autoencoder [23,24]. Finally, in Section 5 we provide a summary of our findings and discuss avenues for future work.

2. Existing Heterogeneity Indices

2.1. Rényi Heterogeneity in Categorical Systems

There are many approaches to derive Rényi heterogeneity [18–20]. Here, we loosely follow the presentation of Eliazar and Sokolov [25] by using the metaphor of repeated sampling from a discrete system X with event space $\mathcal{X} = \{1, 2, \dots, n\}$ and probability distribution $\mathbf{p} = (p_i)_{i=1,2,\dots,n}$. The probability that $q \in \mathbb{N}_{>1}$ independent and identically distributed (i.i.d.) realizations of X , sampled with replacement, will be identical is

$$\mathbb{P}_X [x_1 = x_2 = \dots = x_q] = \sum_{i=1}^n p_i^q. \quad (1)$$

Now let X_* be an idealized reference system with a uniform probability distribution over n_* categorical states, $\mathbf{p}_* = (n_*^{-1})_{i=1,2,\dots,n_*}$, and let $(x_{*1}, x_{*2}, \dots, x_{*q})$ be a sample of q i.i.d. realizations of X_* such that

$$\mathbb{P}_X [x_1 = x_2 = \dots = x_q] = \mathbb{P}_{X_*} [x_{*1} = x_{*2} = \dots = x_{*q}] = \sum_{i=1}^{n_*} n_*^{-q}. \tag{2}$$

We call X_* an “idealized” categorical system because its probability distribution is uniform, and it is a “reference” system for X in that the probability of drawing homogeneous samples of q observations from both systems is identical. Substituting Equation (2) into Equation (1) and solving for n_* yields the Rényi heterogeneity of order q ,

$$\Pi_q(\mathbf{p}) = \left(\sum_{i=1}^n p_i^q \right)^{\frac{1}{1-q}} = n_*, \tag{3}$$

whose units are the numbers equivalent of system X [1,26–28], insofar as n_* is the number of states in an “equivalent” (idealized reference) system X_* . Thus far, we have restricted the parameter q to take integer values greater than 1 solely to facilitate this intuitive derivation in a concise fashion. However, the elasticity parameter q in Equation (3) can be any real number (but $q \neq 1$), although in the context of heterogeneity measurement only $q \geq 0$ are used [1,25]. Despite Equation (3) being undefined at $q = 1$ directly, L’Hôpital’s rule can be used to show that the limit $q \rightarrow 1$ exists, wherein it corresponds to the exponential of Shannon’s entropy [28,29], known as perplexity [30].

Equation (3) is the exponential of Rényi’s entropy [18], and is alternatively known as the Hill numbers in ecology [1,19], Hannah–Kay indices in economics [20], and generalized inverse participation ratio in physics [25]. Interestingly, it generalizes or can be transformed into several heterogeneity indices that are commonly employed across scientific disciplines (Table 1).

Table 1. Relationships between Rényi heterogeneity and various diversity or inequality indices for a system X with event space $\mathcal{X} = \{1, 2, \dots, n\}$ and probability distribution $\mathbf{p} = (p_i)_{i=1,2,\dots,n}$. The function $\mathbb{1}[\cdot]$ is an indicator function that evaluates to 1 if its argument is true or to 0 otherwise.

Index	Expression
Observed richness [31]	$\Pi_0(\mathbf{p}) = \sum_{i=1}^n \mathbb{1}[p_i > 0]$
Perplexity [30]	$\Pi_1(\mathbf{p}) = \exp \left\{ -\sum_{i=1}^n p_i \log p_i \right\}$
Inverse Simpson concentration [1]	$\Pi_2(\mathbf{p}) = \left(\sum_{i=1}^n p_i^2 \right)^{-1}$
Berger-Parker Diversity Index [32,33]	$\Pi_\infty(\mathbf{p}) = \left(\max_i p_i \right)^{-1}$
Rényi entropy [18]	$R_q(\mathbf{p}) = \log \Pi_q(\mathbf{p})$
Shannon entropy [29]	$H(\mathbf{p}) = \log \Pi_1(\mathbf{p})$
Tsallis entropy [34]	$T_q(\mathbf{p}) = \frac{1}{q-1} \left(1 - \Pi_q(\mathbf{p})^{1-q} \right)$
Simpson concentration [35]	$\text{Simpson}(\mathbf{p}) = \left(\Pi_2(\mathbf{p}) \right)^{-1}$
Gini-Simpson index [36]	$\text{GSI}(\mathbf{p}) = 1 - \text{Simpson}(\mathbf{p})$
Generalized entropy index [3,37]	$\text{GEI}(\mathbf{p}) = \frac{1}{q(q-1)} \left[\left(\frac{1}{n} \Pi_q(\mathbf{p}) \right)^{1-q} - 1 \right]$

2.1.1. Properties of the Rényi Heterogeneity

Equation (3) satisfies several properties that render it a preferable measure of heterogeneity. These have been detailed elsewhere [1,20,25,28,33,38], but we focus on three properties that are of particular relevance for the remainder of this paper.

First, Π_q satisfies the principle of transfers [39,40] which states that any equality-increasing transfer of probability between states must increase the heterogeneity. The maximal value of Π_q is attained if and only if $p_i = p_j$ for all $(i, j) \in \{1, 2, \dots, n\}$. This property follows from Schur-concavity of Equation (3) [20].

Second, Π_q satisfies the replication principle [1,38,41], which is equivalent to stating that Equation (3) scales linearly with the number of equally probable states in an idealized categorical system [25]. More formally, consider a set of systems X_1, X_2, \dots, X_N with probability distributions

$\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N$ over respective discrete event spaces $\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_N$. These systems are also assumed to satisfy the following properties:

1. Event spaces are disjoint: $\mathcal{X}_i \cap \mathcal{X}_j = \emptyset$ for all $(i, j) \in \{1, 2, \dots, N\}$ where $i \neq j$
2. All systems have equal heterogeneity: $\Pi_q(\mathbf{p}_1) = \Pi_q(\mathbf{p}_2) = \dots = \Pi_q(\mathbf{p}_i) = \dots = \Pi_q(\mathbf{p}_N)$

The replication principle states that if we combine X_1, X_2, \dots, X_N into a pooled system X with probability distribution $\bar{\mathbf{p}}$, then

$$\Pi_q(\bar{\mathbf{p}}) = N\Pi_q(\mathbf{p}_i) \tag{4}$$

must hold (see Appendix A for proof that Rényi heterogeneity satisfies the replication principle).

The replication principle suggests that Equation (3) satisfies a property known as decomposability, in that the heterogeneity of a pooled system can be decomposed into that arising from variation within and between component subsystems. However, we require that this property be satisfied when either (A) subsystems' event spaces are overlapping, or (B) subsystems do not have equal heterogeneity. The decomposability property will be particularly important for Section 3, and so we detail it further in Section 2.1.2.

2.1.2. Decomposition of Categorical Rényi Heterogeneity

Consider a system X defined by pooling subsystems X_1, X_2, \dots, X_N with potentially overlapping event spaces $\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_N$, respectively. The event space of the pooled system is defined as

$$\mathcal{X} = \cup_{i=1}^N \mathcal{X}_i = \{1, 2, \dots, n\}. \tag{5}$$

Furthermore, we define the matrix $\mathbf{P} = (p_{ij})_{i=1,2,\dots,n}^{j=1,2,\dots,n}$ whose i^{th} row is the probability of system X_i being observed in each state $j \in \{1, 2, \dots, n\}$.

It may be the case that some subsystems comprise a larger proportion of X than others. For instance, if the probability distribution for subsystem X_i was estimated based on a larger sample size than that of X_j , one may want to weight the contribution of X_i higher. Thus, we define a column vector of weights $\mathbf{w} = (w_i)_{i=1,2,\dots,N}$ over the N subsystems such that $\sum_{i=1}^N w_i = 1$ and $w_i \geq 0$ for all i . The probability distribution over states in the pooled system X may thus be computed as $\bar{\mathbf{p}} = \sum_{i=1}^N w_i \mathbf{p}_i$, from which the definition of pooled heterogeneity follows:

$$\Pi_q^P(\mathbf{P}, \mathbf{w}) = \left[\sum_{j=1}^n \left(\sum_{i=1}^N w_i p_{ij} \right)^q \right]^{\frac{1}{1-q}}. \tag{6}$$

One can interpret $\Pi_q^P(\mathbf{P}, \mathbf{w})$ as the effective number of states in the pooled categorical system X .

Jost [28] showed that the within-group heterogeneity, which is the effective number of unique states arising from individual component systems, can be defined as

$$\Pi_q^W(\mathbf{P}, \mathbf{w}) = \left[\frac{\sum_{i=1}^N w_i^q \left(\sum_{j=1}^n p_{ij}^q \right)}{\sum_{k=1}^N w_k^q} \right]^{\frac{1}{1-q}}, \tag{7}$$

For example, in the case where all subsystems have disjoint event spaces with heterogeneity equal to constant ν , then they each contribute ν unique states to the pooled system X .

Deriving the between-group heterogeneity $\Pi_q^B(\mathbf{P}, \mathbf{w})$, is thus straightforward. If the effective total number of states in the pooled system is $\Pi_q^P(\mathbf{P}, \mathbf{w})$, and the effective number of unique states contributed by distinct subsystems is $\Pi_q^W(\mathbf{P}, \mathbf{w})$, then

$$\Pi_q^B(\mathbf{P}, \mathbf{w}) = \frac{\Pi_q^P(\mathbf{P}, \mathbf{w})}{\Pi_q^W(\mathbf{P}, \mathbf{w})} \tag{8}$$

is the effective number of completely distinct subsystems in the pooled system X . A word of caution is warranted. If we require that within-group heterogeneity is a lower bound on pooled heterogeneity [42], then (Jost [28], see Proofs 2 and 3) showed that Equation (8) will hold (A) at any value of q when weights are equal (i.e., $w_i = 1/N$ for all $i \in \{1, 2, \dots, N\}$), or (B) only at $q = 0$ and $q = 1$ if weights are unequal.

2.1.3. Limitations of Categorical Rényi Heterogeneity

The chief limitation of Rényi heterogeneity (Equation (3)) is its assumption that all states in a system X (with event space $\mathcal{X} = \{1, 2, \dots, n\}$ and probability distribution $\mathbf{p} = (p_i)_{i=1,2,\dots,n}$) are categorical. More formally, the dissimilarity between a pair of observations $(x, y) \in \mathcal{X}$ from this system is defined by the discrete metric

$$d^*(x, y) = 1 - \delta_{xy}, \tag{9}$$

where δ_{xy} is Kronecker’s delta, which takes a value of 1 if $x = y$ and 0 otherwise. Since the discrete metric assumption is an idealization, we have continued to use the asterisk to qualify an arbitrary distance function $d(\cdot, \cdot)$ as categorical in nature. The resulting expected pairwise distance matrix between states in X is

$$\mathbf{D}^* = [d^*(i, j)]_{i=1,2,\dots,n}^{j=1,2,\dots,n} = \mathbf{1}\mathbf{1}^\top - \mathbf{I}, \tag{10}$$

where $\mathbf{1} = (1)_{i=1,2,\dots,n}$ is a column vector of ones, and $\mathbf{I} = (\delta_{ij})_{i=1,2,\dots,n}^{j=1,2,\dots,n}$ is the $n \times n$ identity matrix.

Clearly, many systems of interest in the real world are not categorical. For example, although we may label a sample of organisms according to their respective species, there may be differences between these taxonomic classes that are relevant to the functioning of the ecosystem as a whole [5]. It is also possible that no valid and reliable set of categorical labels is known a priori for a system whose event space is naturally non-categorical.

2.2. Non-Categorical Heterogeneity Indices

Consider a system X with probability distribution $\mathbf{p} = (p_i)_{i=1,2,\dots,n}$ defined over event space $\mathcal{X} = \{1, 2, \dots, n\}$ and equipped with dissimilarity function $d_X(\cdot, \cdot)$. We assume that d_X is more general than the discrete metric (Equation (9)), and further still need not be a true (metric) distance. For such systems, there are three heterogeneity indices whose units are numbers equivalent, and respect the replication principle [6,8,10,11,21]. Much like our derivation of the Rényi heterogeneity in Section 2.1, these indices quantify the heterogeneity of a non-categorical system as the number of states in an idealized reference system, but differ primarily in how the idealized reference is defined. We begin with a discussion of the Numbers-Equivalent Quadratic Entropy (Section 2.2.1), followed by the Functional Hill Numbers (Section 2.2.2) and the Leinster–Cobbold index [10] (Section 2.2.3).

2.2.1. Numbers Equivalent Quadratic Entropy

Rao [43] introduced the diversity index commonly known as Rao’s quadratic entropy (RQE),

$$Q_1(\mathbf{D}, \mathbf{p}) = \sum_{i=1}^n \sum_{j=1}^n D_{ij} p_i p_j \tag{11}$$

where \mathbf{D} is an $n \times n$ matrix where $D_{ij} = d_X(i, j)$ for states $(i, j) \in \mathcal{X}$.

Ricotta and Szeidl [21] assume that $D_{ij} = 1$ means that states i and j are maximally dissimilar (i.e., categorically different), and that $D_{ij} = 0$ means $i = j$, which occurs when \mathcal{X} is a categorical system. An arbitrary dissimilarity matrix \mathbf{D} can be rescaled to respect this assumption by applying the following transformation:

$$\tilde{\mathbf{D}} = \frac{\mathbf{D} - \min_{ij} D_{ij}}{\max_{ij} D_{ij} - \min_{ij} D_{ij}}. \tag{12}$$

Under this transformation, Ricotta and Szeidl [21] search for an idealized categorical reference system X_* with event space $\mathcal{X}_* = \{1, 2, \dots, n_*\}$, probability distribution $\mathbf{p}_* = (n_*^{-1})_{i=1,2,\dots,n_*}$, and RQE equal to that of X . For a column vector of ones, $\mathbf{1} = (\mathbf{1})_{i=1,2,\dots,n_*}$, and the identity matrix $\mathbf{I} = (\delta_{ij})_{i=1,2,\dots,n_*}^{j=1,2,\dots,n_*}$, this is

$$Q_1(\tilde{\mathbf{D}}, \mathbf{p}) = Q_1(\mathbf{1}\mathbf{1}^\top - \mathbf{I}, \mathbf{p}_*). \tag{13}$$

Expanding the right-hand side, we have

$$Q_1(\tilde{\mathbf{D}}, \mathbf{p}) = \sum_{i=1}^{n_*} \sum_{j=1}^{n_*} n_*^{-2} (1 - \delta_{ij}) = 1 - \frac{1}{n_*}. \tag{14}$$

Recalling that $\Pi_q(\mathbf{p}_*) = n_*$ and substituting into Equation (14) yields

$$\Pi_q(\mathbf{p}_*) = [1 - Q_1(\tilde{\mathbf{D}}, \mathbf{p})]^{-1}, \tag{15}$$

which establishes the units of $[1 - Q_1(\tilde{\mathbf{D}}, \mathbf{p})]^{-1}$ as numbers equivalent.

For consistency, we require that $\Pi_q(\mathbf{p}_*) = \Pi_q(\mathbf{p})$ if $\tilde{\mathbf{D}}$ were categorical. This only holds at $q = 2$:

$$[1 - Q_1(\tilde{\mathbf{D}}, \mathbf{p})]^{-1} = \left[1 - \sum_{i=1}^n \sum_{j=1}^n p_i p_j (1 - \delta_{ij}) \right]^{-1} = \left(\sum_{i=1}^n p_i^2 \right)^{-1} = \Pi_2(\mathbf{p}). \tag{16}$$

Based on this result, Ricotta and Szeidl [21] define the numbers equivalent quadratic entropy \hat{Q}_e as

$$\hat{Q}_e(\tilde{\mathbf{D}}, \mathbf{p}) = (1 - Q_1(\tilde{\mathbf{D}}, \mathbf{p}))^{-1}. \tag{17}$$

This can be interpreted as the inverse Simpson concentration of an idealized categorical reference system whose average pairwise distance between states is equal to $Q_1(\tilde{\mathbf{D}}, \mathbf{p})$.

2.2.2. Functional Hill Numbers

Chiu and Chao [8] derived the Functional Hill Numbers, denoted F_q , based on a similar procedure to that of Ricotta and Szeidl [21]. However, whereas \hat{Q}_e uses a purely categorical system as the idealized reference, F_q requires only that

$$Q_1(\mathbf{D}, \mathbf{p}) = \sum_{i=1}^{n_*} \sum_{j=1}^{n_*} Q_1(\mathbf{D}, \mathbf{p}) p_{*i} p_{*j} = \sum_{i=1}^{n_*} \sum_{j=1}^{n_*} Q_1(\mathbf{D}, \mathbf{p}) n_*^{-2}, \tag{18}$$

which means that the idealized reference system is one for which the between-state distance matrix is set to $Q_1(\mathbf{D}, \mathbf{p})$ everywhere (or to 0 along the leading diagonal and $Q_1(\mathbf{D}, \mathbf{p}) n_* / (n_* - 1)$ on the off diagonals).

Chiu and Chao [8] generalized Rao’s quadratic entropy to include the elasticity parameter $q \geq 0$

$$Q_q(\mathbf{D}, \mathbf{p}) = \sum_{i=1}^n \sum_{j=1}^n D_{ij} (p_i p_j)^q, \tag{19}$$

and sought to find n_* for the idealized reference system satisfying Equation (18) and the following:

$$Q_q(\mathbf{D}, \mathbf{p}) = \sum_{i=1}^{n_*} \sum_{j=1}^{n_*} Q_1(\mathbf{D}, \mathbf{p}) \left(\frac{1}{n_*} \frac{1}{n_*} \right)^q. \tag{20}$$

Solving Equation (20) for n_* yields the functional Hill numbers of order q :

$$F_q(\mathbf{D}, \mathbf{p}) = \left(\frac{Q_q(\mathbf{D}, \mathbf{p})}{Q_1(\mathbf{D}, \mathbf{p})} \right)^{\frac{1}{2(1-q)}} = n_*, \tag{21}$$

which is the effective number of states in an idealized categorical reference system whose distance function is scaled by a factor of $Q_1(\mathbf{D}, \mathbf{p}) n_*/(n_* - 1)$.

2.2.3. Leinster–Cobbold Index

The index derived by Leinster and Cobbold [10], denoted L_q , is distinct from \hat{Q}_e and F_q in two ways. First, for a given system X , the L_q is not derived based on finding an idealized reference system X_* whose average between-state dissimilarity is equal to that of X . Second, it does not use a dissimilarity matrix; rather, it uses a measure of similarity or affinity.

The Leinster–Cobbold index may be derived by simple extension of Equation (3). Assuming X has state space $\mathcal{X} = \{1, 2, \dots, n\}$ with probability distribution $\mathbf{p} = (p_i)_{i=1,2,\dots,n}$, we note that

$$\Pi_q(\mathbf{p}) = \left(\sum_{i=1}^n p_i^q \right)^{\frac{1}{1-q}} = \left[\sum_{i=1}^n p_i (\mathbf{I}\mathbf{p})_i^{q-1} \right]^{\frac{1}{1-q}}. \tag{22}$$

Here, \mathbf{I} is the $n \times n$ identity matrix representing the pairwise similarities between states in X . The Leinster–Cobbold index generalizes \mathbf{I} to be any $n \times n$ similarity matrix \mathbf{S} , yielding the following formula:

$$L_q(\mathbf{S}, \mathbf{p}) = \left[\sum_{i=1}^n p_i \left(\sum_{j=1}^n S_{ij} p_j \right)^{q-1} \right]^{\frac{1}{1-q}}. \tag{23}$$

The similarity matrix can be obtained from a dissimilarity matrix by the transformation $S_{ij} = e^{-uD_{ij}}$, where $u \geq 0$ is a scaling factor. When $u = 0$, then \mathbf{S} is 1 everywhere. Conversely, when $u \rightarrow \infty$, then \mathbf{S} approaches \mathbf{I} . The Leinster–Cobbold index can thus be interpreted as an effective number if the states are in an idealized reference system (i.e., one with uniform probabilities over states) whose topology is also governed by the similarity matrix \mathbf{S} .

2.2.4. Limitations of Existing Non-Categorical Heterogeneity Indices

We illustrate several limitations of the \hat{Q}_e , F_q , and L_q indices using a simple 3-state system X with event space $\mathcal{X} = \{1, 2, 3\}$ over which we specify a probability distribution

$$\mathbf{p}(\kappa) = \begin{cases} (1, 0, 0)^\top & \kappa = 0 \\ \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)^\top & \kappa = 1 \\ (0, 0, 1)^\top & \kappa = \infty \\ \left(\frac{1}{1+\sqrt{\kappa+\kappa'}}, \frac{\sqrt{\kappa}}{1+\sqrt{\kappa+\kappa'}}, \frac{\kappa}{1+\sqrt{\kappa+\kappa'}}\right)^\top & \text{Otherwise} \end{cases} \tag{24}$$

where $0 \leq \kappa$ is a parameter that smoothly varies the level of inequality. When $\kappa = 1$ the distribution is perfectly even (Figure 1A). Since an undirected graph of the system is arranged in a triangle with height h and base b , we also specify the following parametric distance matrix,

$$\mathbf{D}(h, b) = \begin{pmatrix} 0 & b & \sqrt{\frac{b^2}{4} + h^2} \\ b & 0 & \sqrt{\frac{b^2}{4} + h^2} \\ \sqrt{\frac{b^2}{4} + h^2} & \sqrt{\frac{b^2}{4} + h^2} & 0 \end{pmatrix}, \tag{25}$$

which allows us to smoothly vary the level of dissimilarity between states in X . Importantly, Equation (25) allows us to generate distance matrices that are either metric (when $h < b\sqrt{3}/2$; Definition 1) or ultrametric (when $h \geq b\sqrt{3}/2$; Definition 2). This is illustrated in Figure 1B.

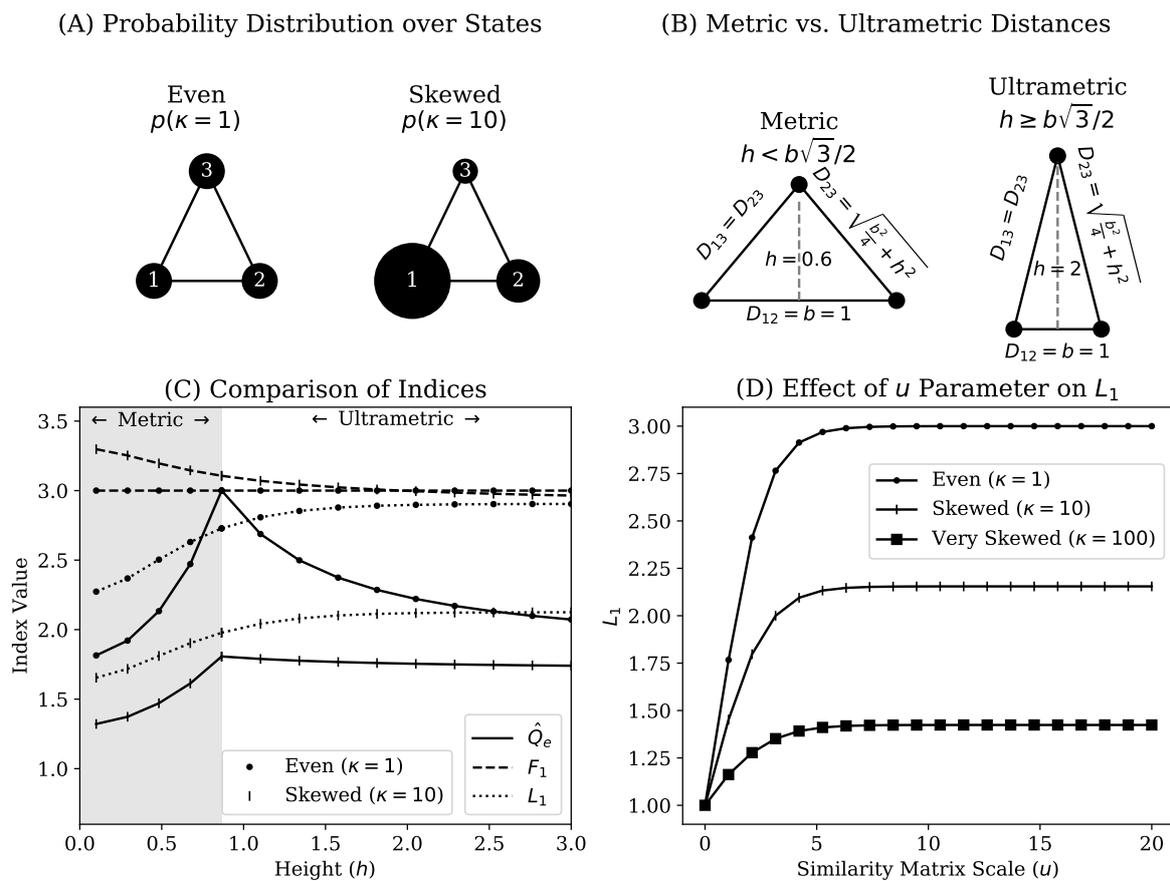


Figure 1. Illustration of simple three-state system under which we compare existing non-categorical heterogeneity indices. **Panel A** depicts a three state system X as an undirected graph, with node sizes corresponding to state probabilities governed by Equation (24). As $0 \leq \kappa$ diverges further from $\kappa = 1$, the probability distribution over states becomes more unequal. **Panel B** visually represents the parametric pairwise distance matrix $\mathbf{D}(h, b)$ of Equation (25) (h is height, b is base length, D_{ij} is distance between states i and j). In the examples shown in Panels B and C, we set $b = 1$. Specifically, we provide visual illustration of settings for which the distance function on X is a metric (Definition 1; when $h < b\sqrt{3}/2$) or ultrametric (Definition 2; when $h \geq b\sqrt{3}/2$). **Panel C** compares the numbers equivalent quadratic entropy (solid lines marked \hat{Q}_e ; Section 2.2.1), functional Hill numbers (at $q = 1$, dashed lines marked F_1 ; Section 2.2.2), and the Leinster–Cobbold Index (at $q = 1$, dotted lines marked L_1 ; Section 2.2.3) for reporting the heterogeneity of X . The y-axis reports the value of respective indices. The x-axis plots the height parameter for the distance matrix $\mathbf{D}(h, 1)$ (Equation (25) and Panel B). The range of h at which $\mathbf{D}(h, 1)$ is only a metric is depicted by the gray shaded background. The range of h shown with a white background is that for which $\mathbf{D}(h, 1)$ is ultrametric. For each index, we plot values for a probability distribution over states that is perfectly even ($\kappa = 1$; dotted markers) or skewed ($\kappa = 10$; vertical line markers). **Panel D** shows the sensitivity of the Leinster–Cobbold index (L_1 ; y-axis) to the scaling parameter $0 \leq u$ (x-axis) used to transform a distance matrix into a similarity matrix ($S_{ij} = e^{-uD_{ij}}$). This is shown for three levels of skewness for the probability distribution over states (no skewness at $\kappa = 1$, dotted markers; significant skewness at $\kappa = 10$, vertical line markers; extreme skewness at $\kappa = 100$, square markers).

Definition 1 (Metric distance). A function $d : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ on a set \mathcal{X} is a metric if and only if all of the following conditions are satisfied for all $(x, y, z) \in \mathcal{X}$:

1. Non-negativity: $d(x, y) \geq 0$
2. Identity of indiscernibles: $d(x, y) = 0 \iff x = y$
3. Symmetry: $d(x, y) = d(y, x)$

4. Triangle inequality: $d(x, z) \leq d(x, y) + d(y, z)$

Definition 2 (Ultrametric distance). A function $d : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ on a set \mathcal{X} is ultrametric if and only if, for all $(x, y, z) \in \mathcal{X}$, criteria 1-3 for a metric are satisfied (Definition 1), in addition to the ultrametric triangle inequality:

$$d(x, z) \leq \max \{d(x, y), d(y, z)\} \quad (26)$$

Figure 1C compares the \hat{Q}_e , F_q , and L_q indices when applied to X across variation in between-state distances (via Equation (25)) and skewness in the probability distribution over states (Equation (24)). With respect to the numbers equivalent quadratic entropy (\hat{Q}_e ; Section 2.2.1), we note that its behavior is categorically different with respect to whether the distance matrix is ultrametric. That is \hat{Q}_e increases with the triangle height parameter h (Equation (25)) until it passes the ultrametric threshold, after which it decreases monotonically with h . The behavior of \hat{Q}_e is sensible in the ultrametric range. When the distance matrix is scaled, as in Equation (12), pulling one of the three states in X further away from the remaining two should function similarly to progressively merging the latter states. Thus, the behavior of \hat{Q}_e is highly sensitive to whether a given distance matrix is ultrametric (which will often not be the case in real-world applications).

With respect to F_q , a notable benefit in comparison to \hat{Q}_e is that F_q behaves consistently regardless of whether distance is ultrametric. However, Figure 1 shows other drawbacks. First, we can see that F_q becomes insensitive to $\mathbf{D}(h, 1)$ when $\mathbf{p}(\kappa)$ is perfectly even (shown analytically in Appendix A). Second, F_q can paradoxically estimate a greater number of states than the theoretical maximum allows. That this occurs when the state probability distribution is more unequal violates the principle of transfers [20,33,39,40] (Section 2.1.1). This is made more problematic since Figure 1C shows it occurs when one state is being pushed closer to the others (i.e., with smaller values of h). To summarize, the functional Hill numbers are estimating more states than are really present despite the reduction in between-state distances and greater inequality in the probability mass function.

Figure 1C shows that the Leinster-Cobbold index compares favorably to F_q because the former does not lose sensitivity to dissimilarity when $\mathbf{p}(\kappa)$ is perfectly even. However, Figure 1D shows that the Leinster-Cobbold index is particularly sensitive to the form of similarity transformation. In the present case, the maximal value of the L_q gradually approaches 3 as u grows (and only when $u \rightarrow \infty$ does it reach 3), while progressively losing sensitivity to distance. As mentioned by Leinster and Cobbold [10], the choice of u or other similarity transformation is dependent on the importance assigned to functional differences between states. However, it is not clear how a given similarity transformation (e.g., u), and therefore the idealized reference system of L_q , should be validated.

Above all of the idiosyncratic limitations of existing numbers equivalent heterogeneity indices, we must highlight two basic assumptions they all share. First, they continue to assume that some valid and reliable categorical partitioning on X is known a priori. Second, they assume that a distance function specified a priori describes semantically relevant geometry of the system in question. These two limitations are not independent, since an unreliable categorical partitioning of the state space will lead to erroneous estimates of the pairwise distances between states. Thus, we seek an approach for measuring heterogeneity that has neither these limitations, nor those shown above to be specific to the other numbers equivalent heterogeneity indices for non-categorical systems.

3. Representational Rényi Heterogeneity

In this section, we propose an alternative approach to the indices of Section 2.2 that we call representational Rényi heterogeneity (RRH). It involves transforming X into a representation Z , defined on an unobservable or latent event space \mathcal{Z} , that satisfies two criteria:

1. The representation Z captures the semantically relevant variation in X
2. Rényi heterogeneity can be directly computed on Z

Satisfaction of the first criterion can only be ascertained in a domain-specific fashion. Since Z is essentially a model of X , investigators must justify that this model is appropriate for the scientific question at hand. For example, an investigator may evaluate the ability of X to be reconstructed from representation Z under cross-validation. The second criterion simply means that the transformation of $X \rightarrow Z$ must specify a probability distribution on Z upon which the Rényi heterogeneity can be directly computed.

Figure 2 illustrates the basic idea of RRH. However, the specifics of this framework differ based on the topology of the representation Z . Thus, the remainder of this section discusses the following approaches:

- A. Application of standard Rényi heterogeneity (Section 2.1) when Z is a categorical representation
- B. Deriving parametric forms for Rényi heterogeneity when Z is a non-categorical representation

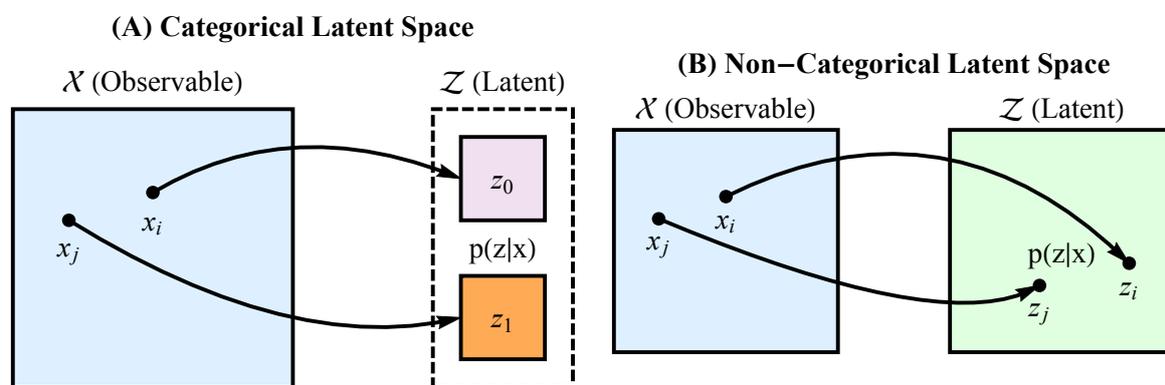


Figure 2. Graphical illustration of the two main approaches for computing representational Rényi heterogeneity. In both cases, we map sampled points on an observable space \mathcal{X} onto a latent space \mathcal{Z} , upon which we apply the Rényi heterogeneity measure. The mapping is illustrated by the curved arrows, and should yield a posterior distribution over the latent space. **Panel A** shows the case in which the latent space is categorical (for example, discrete components of a mixture distribution on a continuous space). **Panel B** illustrates the case in which the latent space has non-categorical topology. A special case of the latter mapping may include probabilistic principal components analysis. When the latent space is continuous, we must derive a parametric form for the Rényi heterogeneity.

3.1. Rényi Heterogeneity on Categorical Representations

Let X be a system defined on an observable space \mathcal{X} that is non-categorical and n_x -dimensional. Consider the scenario in which the semantically relevant variation in X is categorical: for instance, images of different object categories stored in raw form as real-valued vectors. An investigator may be interested in measuring the effective number of states in X with respect to this categorical variation. This requires transforming X into a semantically relevant categorical representation Z upon which Equation (3) can be applied.

Assume we have a large random sample of N points $\mathbf{X} = (\mathbf{x}_i)_{i=1,2,\dots,N}$ from system X . We can conceptualize each discrete observation \mathbf{x}_i in this sample as the single point in the event space of a perfectly homogeneous subsystem X_i . When pooled, the subsystems $\{X_i\}_{i=1,2,\dots,N}$ constitute X . The contribution weights of each subsystem to X as a whole are denoted $\mathbf{w} = (w_i)_{i=1,2,\dots,N}$, where $\sum_{i=1}^N w_i = 1$ and $w_i \geq 0$.

We now specify a vector-valued function $\mathbf{f}: \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Z})$ such that $\mathbf{x} \mapsto \mathbf{f}(\mathbf{x}) = [f_j(\mathbf{x})]_{j=1,2,\dots,n_z}$ is a mapping from n_x -dimensional coordinates on the observable space, $\mathbf{x} \in \mathcal{X}$, onto an n_z -dimensional discrete probability distribution over $\mathcal{Z} = \{1, 2, \dots, n_z\}$. Thus, $\mathbf{f}(\mathbf{x}_i)$ can be conceptualized as mapping

subsystem X_i onto its categorical representation Z_i . After defining \mathbf{f} , the effective number of states in the latent representation of X_i can be computed as

$$\Pi_q(\mathbf{x}_i) = \left(\sum_{j=1}^{n_z} f_j^q(\mathbf{x}_i) \right)^{\frac{1}{1-q}}. \quad (27)$$

When $\Pi_q(\mathbf{x}_i) = 1$, then \mathbf{f} assigns \mathbf{x} to a single category with perfect certainty. Conversely, when $\Pi_q(\mathbf{x}_i) = n_z$, then either \mathbf{x}_i belongs to all categorical states with equal probability, or \mathbf{f} is maximally uncertain about the mapping of point \mathbf{x}_i .

Mapping all points \mathbf{X} onto the categorical latent space yields a collection of subsystems $\{Z_i\}_{i=1,2,\dots,N}$, which generate Z when pooled. Using Equation (6), we can compute the effective number of total states in Z as the pooled heterogeneity:

$$\Pi_q^P(\mathbf{X}, \mathbf{w}) = \left[\sum_{j=1}^{n_z} \left(\sum_{i=1}^N w_i f_j^q(\mathbf{x}_i) \right)^q \right]^{\frac{1}{1-q}}, \quad (28)$$

Unfortunately, $\Pi_q^P(\mathbf{X}, \mathbf{w})$ counts some heterogeneity that is due to uncertainty in the model (i.e., that quantified by Equation (27)). We, therefore, compute the effective number of states in Z per point $\mathbf{x} \in \mathcal{X}$ using the within-group heterogeneity formula (Equation (7)):

$$\Pi_q^W(\mathbf{X}, \mathbf{w}) = \left[\frac{\sum_{i=1}^N w_i^q \left(\sum_{j=1}^{n_z} f_j^q(\mathbf{x}_i) \right)}{\sum_{k=1}^N w_k^q} \right]^{\frac{1}{1-q}}. \quad (29)$$

Finally, the effective number of states (points) in X —with respect to the categorical variation modeled by Z —can then be computed using the between-group heterogeneity formula (Equation (8)):

$$\Pi_q^B(\mathbf{X}, \mathbf{w}) = \frac{\Pi_q^P(\mathbf{X}, \mathbf{w})}{\Pi_q^W(\mathbf{X}, \mathbf{w})}. \quad (30)$$

Example 1 demonstrates that current methods of measuring biodiversity and wealth concentration can be viewed as special cases of categorical RRH.

Example 1 (Classical measurement of biodiversity and economic equality as categorical RRH). *Definitions necessary for this example are shown in Table 2. The traditional analysis of species diversity and economic equality can be recovered from an RRH-based formulation when \mathbf{f} is assumed to be deterministic and $\mathbf{w} = (N^{-1})_{i=1,2,\dots,N}$. In this case within-group heterogeneity can be shown to reduce to 1:*

$$\begin{aligned} \Pi_q^W(\mathbf{X}, \mathbf{w}) &= \left[\sum_{i=1}^N \frac{N^{-q}}{\sum_{k=1}^N N^{-q}} \left(\sum_{j=1}^{n_z} f_j^q(\mathbf{x}_i) \right) \right]^{\frac{1}{1-q}} \\ &= \left[\sum_{i=1}^N N^{-1} (1) \right]^{\frac{1}{1-q}} \\ &= 1. \end{aligned} \quad (31)$$

Thus, we have

$$\begin{aligned}
 \Pi_q^B(\mathbf{X}, \mathbf{w}) &= \Pi_q^P(\mathbf{X}, \mathbf{w}) \\
 &= \left[\sum_{j=1}^{n_z} \left(\sum_{i=1}^N N^{-1} f_j(\mathbf{x}_i) \right)^q \right]^{\frac{1}{1-q}} \\
 &= \left[\sum_{j=1}^{n_z} \left(\frac{N_j}{N} \right)^q \right]^{\frac{1}{1-q}},
 \end{aligned}
 \tag{32}$$

which yields the categorical Rényi heterogeneity (Hill numbers for biodiversity analysis and Hannah–Kay indices in the economic setting [19,20]), and by extension many diversity indices to which it is connected (Table 1). Thus, traditional analysis of species biodiversity and economic equality are special cases of representational Rényi heterogeneity where the representation is specified by a mapping onto degenerate distributions over categorical labels. The only differences lie in the definition of observable and latent spaces, and the representational models.

In the case of biodiversity analysis, the model \mathbf{f} in real-world practice may simply be a human expert assigning species labels to a sample of organisms from a field study. In the economic setting, one may speculate that \mathbf{f} would essentially reduce to contracts specifying ownership of assets, whose value is deemed by market forces.

Table 2. Definitions in formulation of classical biodiversity and economic equality analysis as categorical representational Rényi heterogeneity. Superscripted indexing on $\mathbf{x} = (x_i)^{i=1, \dots, n_x}$ denotes that this is a row vector.

Symbol	Analytical Context	
	Biodiversity	Economic Equality
X	Ecosystem, whose observation yields an organism denoted by vector $\mathbf{x} = (x_i)^{i=1, \dots, n_x} \in \mathcal{X}$	A system of resources, whose observation yields an asset denoted by vector $\mathbf{x} = (x_i)^{i=1, \dots, n_x} \in \mathcal{X}$
$\mathcal{X} \subseteq \mathbb{R}^{n_x}$	n_x -dimensional feature space of organisms in the ecosystem	n_x -dimensional feature space of assets in the economy, whose topology is such that the “economic” or monetary value is equal at each coordinate $\mathbf{x} \in \mathcal{X}$
$\mathcal{Z} = \{\mathbf{z} \in \{0, 1\}^{n_z} : \sum_{i=1}^{n_z} z_i = 1\}$	n_z -dimensional space of one-hot species labels	n_z -dimensional space of one-hot labels over wealth-owning agents
$\mathbf{f} : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Z})$	A model that performs the mapping $\mathbf{x} \mapsto \mathbf{f}(\mathbf{x})$ of organisms to discrete probability distributions over \mathcal{Z}	A model that performs the mapping $\mathbf{x} \mapsto \mathbf{f}(\mathbf{x})$ of assets to discrete probability distributions over \mathcal{Z}
$N_i \in \mathbb{N}_+$	The number of organisms observed belonging to species $i \in \{1, \dots, n_z\}$	The number of equal valued assets belonging to agent $i \in \{1, \dots, n_z\}$
$N = \sum_{i=1}^{n_z} N_i$	The total number of organisms observed	The total quantity of assets observed
$\mathbf{X} = \left(x_{ij} \right)_{i=1, \dots, N}^{j=1, \dots, n_x}$	A sample of N organisms	A sample of N assets
$\mathbf{w} = (w_i)_{i=1, \dots, N}$	Sample weights, such that $w_i \geq 0$ and $\sum_{i=1}^N w_i = 1$	

3.2. Rényi Heterogeneity on Non-Categorical Representations

In Section 3.1, we dealt with instances in which semantically relevant variation in X is categorical, such as when object categories are embedded in images stored as real-valued vectors. Here, we consider scenarios in which the semantically relevant information in an observable system X is non-categorical: for instance, where a piece of text contains information about semantic concepts best represented as real-valued “word vectors” [44,45]. Measuring the effective number of distinct states in X with respect to this continuous variation requires transforming X into a semantically relevant continuous representation Z upon which procedures analogous to those of Section 3.1 may be undertaken.

Let Z be defined on an n_z -dimensional event space $\mathcal{Z} \subseteq \mathbb{R}^{n_z}$ over which there exists a family of parametric probability distributions $\mathcal{P}(\mathcal{Z})$ of a form chosen by the experimenter. Let $f : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Z})$ be a model that performs the mapping $\mathbf{x} \mapsto f(\cdot|\mathbf{x})$ from a point $\mathbf{x} \in \mathcal{X}$ on the observable space to a probability density on \mathcal{Z} . For example, if $\mathcal{P}(\mathcal{Z})$ is the family of multivariate Gaussians, then $f(\mathbf{z}|\mathbf{x}_i) = \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$, where $\boldsymbol{\mu}_i$ and $\boldsymbol{\Sigma}_i$ are the Gaussian mean and covariance functions at \mathbf{x}_i , respectively. Given a sample $\mathbf{X} = (\mathbf{x}_i)_{i=1,2,\dots,N}$, as in Section 3.1, we compute the continuous analogue of Equation (27) as follows

$$\Pi_q(\mathbf{x}_i) = \left(\int_{\mathcal{Z}} f^q(\mathbf{z}|\mathbf{x}_i) \, d\mathbf{z} \right)^{\frac{1}{1-q}}. \quad (33)$$

This formula yields the effective size of the domain of a uniform distribution on \mathbb{R}^{n_z} whose Rényi heterogeneity is equal to $\Pi_q(\mathbf{x}_i)$ (proof is given in Appendix A). Thus, it is possible for $\Pi_q(\mathbf{x}_i)$ to be less than 1, though it will remain non-negative.

Similar to the procedure in Section 3.1, we now define a continuous version of the within-observation heterogeneity

$$\Pi_q^W(\mathbf{X}, \mathbf{w}) = \left[\sum_{i=1}^N \frac{w_i^q}{\sum_{j=1}^N w_j^q} \int_{\mathcal{Z}} f^q(\mathbf{z}|\mathbf{x}_i) \, d\mathbf{z} \right]^{\frac{1}{1-q}}, \quad (34)$$

which estimates the effective size of the latent space occupied per observable point $\mathbf{x} \in \mathcal{X}$.

In order to compute the pooled heterogeneity $\Pi_q^P(\mathbf{X}, \mathbf{w})$, the experimenter must specify the form of the pooled distribution, here denoted $\bar{f}_{\mathbf{w}}$. The conceptually most simple approach is non-parametric, using a model average,

$$\bar{f}_{\mathbf{w}}(\mathbf{z}|\mathbf{X}) = \sum_{i=1}^N w_i f(\mathbf{z}|\mathbf{x}_i), \quad (35)$$

whereby the pooled heterogeneity would be

$$\Pi_q^P(\mathbf{X}, \mathbf{w}) = \left[\int_{\mathcal{Z}} \left(\sum_{i=1}^N w_i f(\mathbf{z}|\mathbf{x}_i) \right)^q \, d\mathbf{z} \right]^{\frac{1}{1-q}}. \quad (36)$$

The integral in Equation (36) may often be analytically intractable and potentially difficult to solve accurately in high dimensions with numerical methods. Furthermore, some areas of \mathcal{Z} may be assigned low probability by $f(\mathbf{z}|\mathbf{x}_i)$ for all $i \in \{1, 2, \dots, N\}$. This is not a problem as the sample \mathbf{X} becomes infinitely large. However, with finite samples, it may be the case that some representational states in \mathcal{Z} are unlikely simply because we have not sampled from the corresponding regions of \mathcal{X} . An alternative to Equation (35) is therefore to specify a parametric pooled distribution

$$\bar{f}_{\mathbf{w}}(\cdot|\mathbf{X}) = \Xi_f(\mathbf{X}, \mathbf{w}), \quad (37)$$

where Ξ_f is a deterministic function that combines $f(\cdot|\mathbf{x}_i)$ for $i \in \{1, 2, \dots, N\}$ into a valid probability density on \mathcal{Z} . In this case, the pooled Rényi heterogeneity is simply

$$\Pi_q^P(\mathbf{X}, \mathbf{w}) = \left(\int_{\mathcal{Z}} \bar{f}_{\mathbf{w}}^q(\mathbf{z}|\mathbf{X}) \, d\mathbf{z} \right)^{\frac{1}{1-q}}. \tag{38}$$

Using either Equation (36) or (38) as the pooled heterogeneity and Equation (34) as the within-group heterogeneity, the effective number of distinct states in X —with respect to the non-categorical representation Z —can then be computed using Equation (30).

Figure 3 demonstrates the difference between the parametric and non-parametric approaches to pooling for non-categorical RRH, and Example 2 demonstrates one approach to parametric pooling for a mixture of multivariate Gaussians.

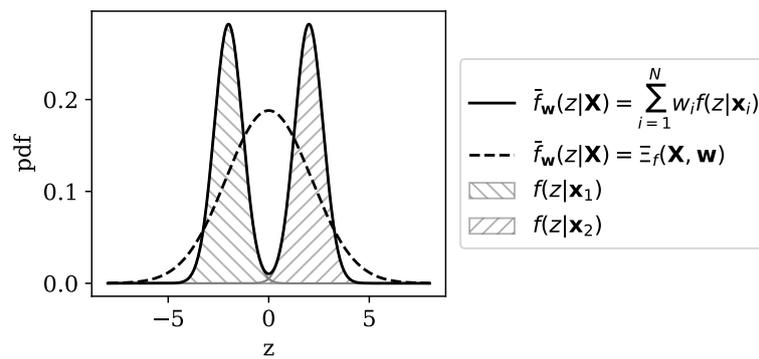


Figure 3. Illustration of approaches to computing the pooled distribution on a simple representational space $\mathcal{Z} = \mathbb{R}$. In this example, two points on the observable space, $(\mathbf{x}_1, \mathbf{x}_2) \in \mathcal{X}$, are mapped onto the latent space via model $f(\cdot|\mathbf{x}_i)$ for $i \in \{1, 2\}$, which indexes univariate Gaussians over \mathcal{Z} (depicted as hatched patterns for \mathbf{x}_1 and \mathbf{x}_2 , respectively). A pooled distribution computed non-parametrically by model-averaging (Equation (35)) is depicted as the solid black line. The parametrically pooled distribution (see Example 2) is depicted as the dashed black line. The parametric approach implies the assumption that further samples from \mathcal{X} would yield latent space projections in some regions assigned low probability by $f(z|\mathbf{x}_1)$ and $f(z|\mathbf{x}_2)$.

Example 2 (Parametric pooling of multivariate Gaussian distributions). Let $\mathbf{X} = (\mathbf{x}_i)_{i=1,2,\dots,N}$ be a sample of n_x -dimensional vectors from a system X with event space $\mathcal{X} \subseteq \mathbb{R}^{n_x}$. Let Z be a latent representation of X with n_z -dimensional event space $\mathcal{Z} = \mathbb{R}^{n_z}$. Let

$$f(\mathbf{z}|\mathbf{x}_i) = \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i) \tag{39}$$

be a model that returns a multivariate Gaussian density with mean $\boldsymbol{\mu}_i$ and covariance $\boldsymbol{\Sigma}_i$ given point $\mathbf{x}_i \in \mathcal{X}$. Finally, let $\mathbf{w} = (w_i)_{i=1,2,\dots,N}$ be weights assigned to each sample in \mathbf{X} such that $w_i \geq 0$ and $\sum_{i=1}^N w_i = 1$.

If one assumes that the pooled distribution over \mathcal{Z} given the set of components $f(\mathbf{z}|\mathbf{x}_1), f(\mathbf{z}|\mathbf{x}_2), \dots, f(\mathbf{z}|\mathbf{x}_N)$ is itself a multivariate Gaussian,

$$\bar{f}_{\mathbf{w}}(\mathbf{z}|\mathbf{X}) = \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}_*, \boldsymbol{\Sigma}_*) \tag{40}$$

with $n_z \times 1$ pooled mean,

$$\boldsymbol{\mu}_* = \sum_{i=1}^N w_i \boldsymbol{\mu}_i \tag{41}$$

and $n_z \times n_z$ pooled covariance matrix

$$\Sigma_* = -\boldsymbol{\mu}_* \boldsymbol{\mu}_*^\top + \sum_{i=1}^N w_i [\Sigma_i + \boldsymbol{\mu}_i \boldsymbol{\mu}_i^\top], \tag{42}$$

then the pooled heterogeneity Π_q^P is therefore simply the Rényi heterogeneity of a multivariate Gaussian,

$$\Pi_q(\Sigma) = \begin{cases} \text{Undefined} & q = 0 \\ (2\pi e)^{\frac{n_z}{2}} \sqrt{|\Sigma|} & q = 1 \\ (2\pi)^{\frac{n_z}{2}} \sqrt{|\Sigma|} & q = \infty \\ (2\pi)^{\frac{n_z}{2}} q^{\frac{n_z}{2(q-1)}} \sqrt{|\Sigma|} & \text{Otherwise} \end{cases} \tag{43}$$

evaluated at Σ_* . The derivation is provided in Appendix A [46]. Equation (43) at Σ_* is interpreted as the effective size of space \mathcal{Z} occupied by the complete latent representation of X under model f .

The within-group heterogeneity can be obtained for the set of components $[f(\mathbf{z}|\mathbf{x}_i)]_{i=1,2,\dots,N}$ by solving Equation (34) for the Gaussian densities, yielding:

$$\Pi_q^W(\Sigma_{1:N}, \mathbf{w}) = \begin{cases} \text{Undefined} & q = 0 \\ \exp\left\{\frac{1}{2}\left(n_z + \sum_{i=1}^N w_i \log |2\pi \Sigma_i|\right)\right\} & q = 1 \\ 0 & q = \infty \\ (2\pi)^{\frac{n_z}{2}} \left(\sum_{i=1}^N \frac{\bar{w}_i^q |\Sigma_i|^{\frac{1}{2}}}{q^{\frac{n_z}{2}}}\right)^{\frac{1}{1-q}} & \text{Otherwise} \end{cases}, \tag{44}$$

where we denote $\Sigma_{1:N} = \{\Sigma_i\}_{i=1,2,\dots,N}$ for parsimony, and $\bar{w}_i = w_i \left(\sum_{j=1}^N w_j^q\right)^{-1/q}$. Equation (44) estimates the effective size of the n_z -dimensional representational space occupied per state $\mathbf{x} \in \mathcal{X}$.

The effective number of states in X with respect to the continuous representation \mathcal{Z} is thus the between-group heterogeneity Π_q^B which can be computed as the ratio $\Pi_q(\Sigma_*) / \Pi_q^W(\Sigma_{1:N}, \mathbf{w})$. The properties of this decomposition—specifically the conditions under which $\Pi_q^B \geq 1$ (Lande’s requirement [28,42])—are discussed further elsewhere [46].

4. Empirical Applications of Representational Rényi Heterogeneity

In this section, we demonstrate two applications of RRH under assumptions of categorical (Section 4.1) and continuous (Section 4.2) latent spaces. First, Section 4.1, uses a simple closed-form system consisting of a mixture of two beta distributions on the (0,1) interval to give exact comparisons of the behavior of RRH against that of existing non-categorical heterogeneity indices (Section 2.2). This experiment provides evidence that existing non-categorical heterogeneity indices can demonstrate counterintuitive behavior under various circumstances. Second, Section 4.2 demonstrates that RRH can yield heterogeneity measurements that are sensible and tractably computed, even for highly complex mappings $f : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Z})$. There, we use a deep neural network to compute the effective number of observations in a database of handwritten images with respect to compressed latent representations on a continuous space.

4.1. Comparison of Heterogeneity Indices Under a Mixture of Beta Distributions

Consider a system X with event space \mathcal{X} on the open interval (0, 1), containing an embedded, unobservable, categorical structure represented by the latent system Z with event space $\mathcal{Z} = \{1, 2\}$. The systems’ collective behavior is governed by the joint distribution of a beta mixture model (BMM),

$$p(x, z) = \mathbb{1}[z = 1](1 - \theta_1)\text{Beta}_{\theta_2, \theta_3}(x) + \mathbb{1}[z = 2]\theta_1\text{Beta}_{\theta_3, \theta_2}(x), \tag{45}$$

where $\text{Beta}_{\alpha,\beta}(x)$ is the probability density function for a beta distribution with shape parameters α, β , and $\theta = (\theta_1, \theta_2, \theta_3)$ are parameters. The indicator function $\mathbb{1}[\cdot]$ evaluates to 1 if its argument is true, and to 0 otherwise. The prior distribution is

$$p(z) = \mathbb{1}[z = 1](1 - \theta_1) + \mathbb{1}[z = 2]\theta_1, \tag{46}$$

and marginal probability of observable data is as follows (see Figure 4 for illustrations):

$$p(x) = (1 - \theta_1)\text{Beta}_{\theta_2,\theta_3}(x) + \theta_1\text{Beta}_{\theta_3,\theta_2}(x). \tag{47}$$

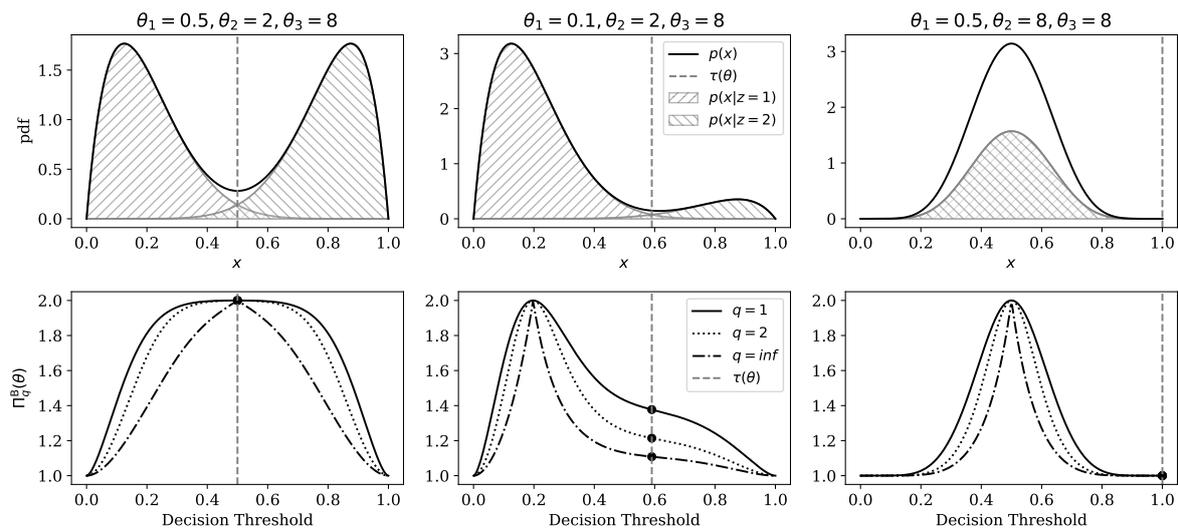


Figure 4. Demonstration of data-generating distribution (top row; Equations (45)–(47)), and relationship between the representational model’s decision threshold (Equations (48) and (50)) and categorical representational Rényi heterogeneity (bottom row). The optimal decision boundary (Equation (50)) is shown as a gray vertical dashed line in all plots. Each column depicts a specific parameterization of the data-generating system (parameters are stated above the top row). **Top Row:** Probability density functions for data-generating distributions. Shaded regions correspond to the two mixture components. Solid black lines denote the marginal distribution (Equation (47)). The x-axis represents the observable domain, which is the (0,1) interval. **Bottom Row:** Effect of varying categorical representational Rényi heterogeneity (RRH) for $q \in \{1, 2, \infty\}$ across different category assignment thresholds for the beta-mixture models shown in the top row. Varying levels of decision boundary are plotted on the x-axis. The y-axis shows the resulting between-observation RRH. Black dots highlight the RRH computed at the optimal decision boundary.

To facilitate exact comparisons between heterogeneity indices, below, let us assume we have a model $f : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Z})$ that maps an observation $x \in \mathcal{X}$ onto a degenerate distribution over \mathcal{Z} :

$$f_{\theta}(z|x) = \mathbb{1}[z = 1]\mathbb{1}[x \leq \tau(\theta)] + \mathbb{1}[z = 2]\mathbb{1}[x > \tau(\theta)]. \tag{48}$$

The subscripting of f_{θ} denotes that the model is optimized such that the threshold $0 \leq \tau(\theta) \leq 1$ is the solution to

$$p(z = 1|x = \tau(\theta)) = p(z = 2|x = \tau(\theta)), \tag{49}$$

which is

$$\tau(\theta) = \begin{cases} \left[\left(\theta_1^{-1} - 1 \right)^{\frac{1}{2(\theta_2 - \theta_3)}} (1 - \theta_1)^{\frac{1}{2(\theta_2 - \theta_3)}} \theta_1^{-\frac{1}{2(\theta_2 - \theta_3)}} + 1 \right]^{-1} & \theta_2 - \theta_3 \neq 0 \\ 0 & ((\theta_2 = \theta_3) \wedge (\theta_1 > \frac{1}{2})) \\ 1 & \text{Otherwise} \end{cases} \quad (50)$$

Under this model, the categorical RRH at point $x \in \mathcal{X}$ is

$$\Pi_q(x) = \left(\sum_{i=1}^2 f_{\theta}^q(z = i|x) \right)^{\frac{1}{1-q}} = (\mathbb{1}^q[x \leq \tau(\theta)] + \mathbb{1}^q[x > \tau(\theta)])^{\frac{1}{1-q}} = 1. \quad (51)$$

The expected value of $f_{\theta}(z = 2|x)$ with respect to the data generating distribution (Equation (47)) is

$$\begin{aligned} \bar{f}_{\theta}(z = 2) &= \mathbb{E}_{x \sim p(x)} [f_{\theta}(z = 2|x)] \\ &= \int_0^1 p(x) \mathbb{1}[x > \tau(\theta)] dx \\ &= \int_{\tau(\theta)}^1 p(x) dx \\ &= (1 - \theta_1) I_x^1(\theta_2, \theta_3) + \theta_1 I_x^1(\theta_3, \theta_2), \end{aligned} \quad (52)$$

where $I_{x_0}^{x_1}(a, b)$ is the generalized regularized incomplete beta function (BetaRegularized[x_0, x_1, a, b] command in the Wolfram language and `betainc(a, b, x0, x1, regularized=True)` in Python's `mpmath` package). Equation (52) implies that $\bar{f}_{\theta}(z = 1) = 1 - \bar{f}_{\theta}(z = 2)$. The pooled heterogeneity is thus expressed as a function of θ as follows:

$$\Pi_q^P(\theta) = \begin{cases} \sum_{i=1}^2 \mathbb{1}[\bar{f}_{\theta}(z = i) > 0] & q = 0 \\ \exp \left\{ -\sum_{i=1}^2 \bar{f}_{\theta}(z = i) \log \bar{f}_{\theta}(z = i) \right\} & q = 1 \\ (\max_i \bar{f}_{\theta}(z = i))^{-1} & q = \infty \\ \left(\sum_{i=1}^2 \bar{f}_{\theta}^q(z = i) \right)^{\frac{1}{1-q}} & \text{Otherwise} \end{cases} \quad (53)$$

As a function of θ , the within-group heterogeneity is

$$\begin{aligned} \Pi_q^W(\theta) &= \left[\int_0^1 \frac{p^q(x)}{\int_0^1 p^q(u) du} \left(\sum_{i=1}^2 f_{\theta}(z = i|x) \right)^q dx \right]^{\frac{1}{1-q}} \\ &= \left[\int_0^1 \frac{p^q(x)}{\int_0^1 p^q(u) du} (1) dx \right]^{\frac{1}{1-q}} \\ &= 1, \end{aligned} \quad (54)$$

and therefore the between-group heterogeneity is $\Pi_q^B(\theta) = \Pi_q^P(\theta)$.

Analytic expressions for the existing non-categorical heterogeneity indices \hat{Q}_e (Equation (17)), F_q (Equation (21)), and L_q (Equation (23)) were computed as “best-case” scenarios, as follows. First, the probability distributions over states for all expressions was the true prior distribution (Equation (46)). Distance matrices—and by extension, the similarity matrix for L_q —were computed using the closed-form expectation of the absolute distance between two beta-distributed random variables (see Appendix B and the Supplementary Materials).

Figure 5 compares the categorical RRH against \hat{Q}_e , F_q , and L_q for BMM distributions of varying degrees of separation, and across different mixture component weights ($0.5 \leq \theta_1 < 1$). Without significant loss of generality, we show only those comparisons at $q = 1$ (which excludes the numbers equivalent quadratic entropy), and $q = 2$.

The most salient differences between these indices occur when the BMM mixture components completely overlap (i.e., at $\theta_2 = \theta_3$). The RRH correctly identifies that there is effectively only one component, regardless of mixture weights. Only the Leinster–Cobbold index showed invariance to the mixture weights when $\theta_2 = \theta_3$, but it could not correctly identify that data were effectively unimodal.

The other stark difference arose when the mixture components were furthest apart (here when $\theta_2 = 5$ and $\theta_3 = 20$). At this setting, the functional Hill numbers showed a paradoxical increase in the heterogeneity estimate as the prior distribution on components was skewed. The Leinster–Cobbold index was appropriately concave throughout the range of prior weights, but it never reached a value of 2 at its peak (as expected based on the predictions outlined in Section 2.2.3). Conversely, the RRH was always concave and reached a peak of 2 when both mixture components were equally probable.

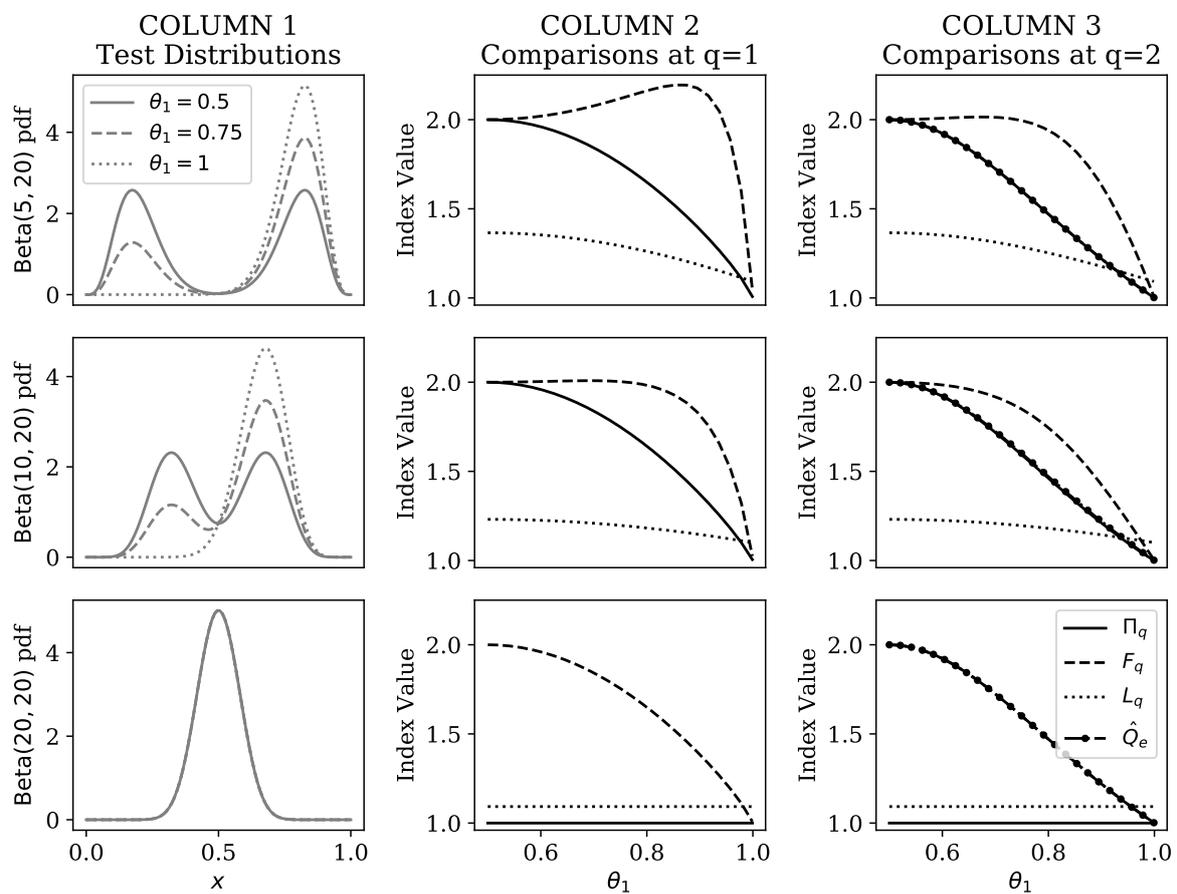


Figure 5. Comparison of categorical representational Rényi heterogeneity (Π_q), the functional Hill numbers (F_q), the numbers equivalent quadratic entropy (\hat{Q}_e), and the Leinster–Cobbold index (L_q) within the beta mixture model. Each row of plots corresponds to a given separation between the beta mixture components. **Column 1** illustrates the beta mixture distributions upon which indices were compared. The x-axis plots the domain of the distribution (open interval between 0 and 1). The y-axis shows the corresponding probability density. Different line styles in Column 1 provides visual examples of the effect of changing the θ_1 parameter over the range $[0.5, 1]$. **Column 2** compares Π_q (solid line), F_q (dashed line), and L_q (dotted line), each at elasticity $q = 1$. The x-axis shows the value of the $0.5 \leq \theta_1 < 1$ parameter at which the indices were compared. Index values are plotted along the y-axis. **Column 3** compares the indices shown in Column 2, as well as \hat{Q}_e (dot-dashed line).

4.2. Representational Rényi Heterogeneity is Scalable to Deep Learning Models

In this example, the observable system X is that of images of handwritten digits defined on an event space $\mathcal{X} = [0, 1]^{784}$ of dimension $n_x = 784$ (the black and white images are flattened from 28×28 pixel matrices into 784-dimensional vectors). Our sample $\mathbf{X} = (x_{ij})_{i=1,2,\dots,N}^{j=1,2,\dots,784}$ from this space is the familiar MNIST training dataset [22] (Figure 6), which consists of $N = 60,000$ images roughly evenly distributed across digits $\{0, 1, \dots, 9\}$, and where approximately 10% of all images come from each class. We assume each image carries equal importance, given by a weight vector $\mathbf{w} = (N^{-1})_{i=1,2,\dots,N}$. We are interested in measuring the heterogeneity of X with respect to a continuous latent representation Z defined on event space $\mathcal{Z} = \mathbb{R}^2$. In the present example, this space is simply the continuous 2-dimensional compression of an image that best facilitates its reconstruction. We choose a dimensionality of $n_z = 2$ for the latent space in order to facilitate a pedagogically useful visualization of the latent feature representation, below. Unlike Section 4.1, in the present case we have no explicit representation of the true marginal distribution over the data, $p(\mathbf{x})$.

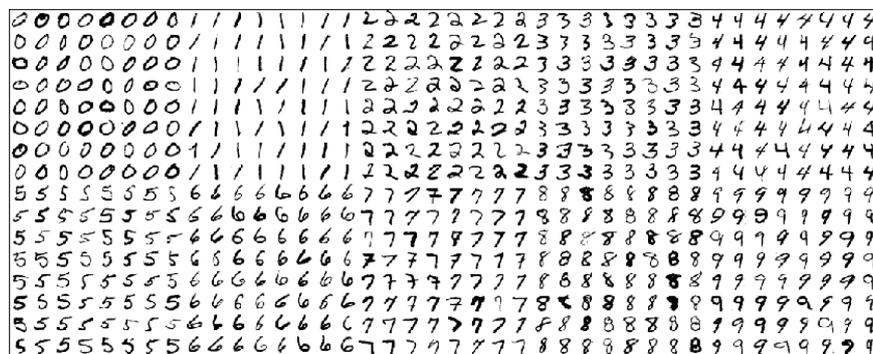
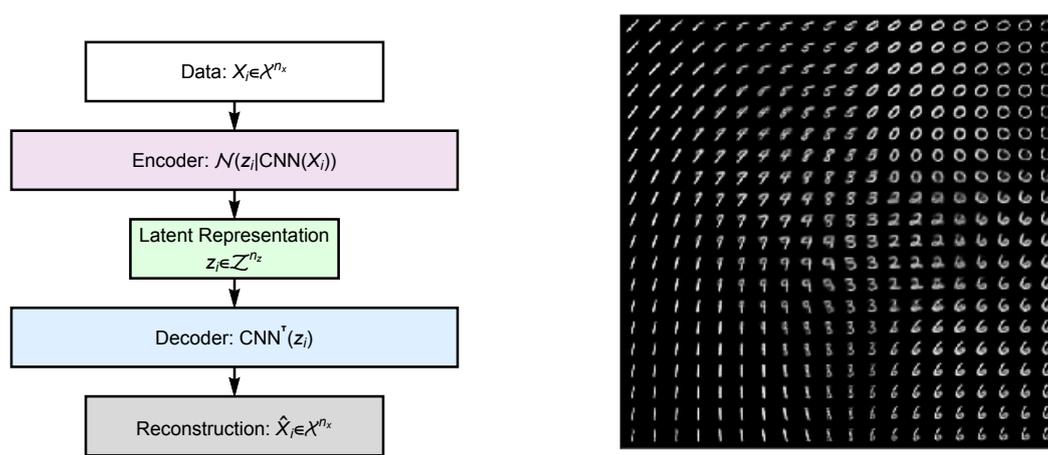


Figure 6. Sample images from the MNIST dataset [22].

Having defined the observable and latent spaces, measuring RRH now requires defining a model $f : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Z})$ that maps a (flattened) image vector $x_i \in \mathcal{X}$ onto a probability distribution over the latent space. Our chosen model is the encoder module of a pre-trained convolutional variational autoencoder (cVAE) provided by the (https://colab.research.google.com/github/smartgeometry-ucl/dl4g/blob/master/variational_autoencoder.ipynb, Smart Geometry Processing Group at University College London) (Figure 7) [23,24]:



(a) Schematic of the model architecture. (b) Visualization of the two-dimensional latent space.

Figure 7. Panel A: Illustration of the convolutional variational autoencoder (cVAE) [23]. The computational graph is depicted from top to bottom. An n_x -dimensional input data X_i (white rectangle) is passed through an encoder (in our experiment this is a convolutional neural network, CNN) which parameterizes an n_z -dimensional multivariate Gaussian over the coordinates z_i for the image’s embedding on the latent space $\mathcal{Z} = \mathbb{R}^{n_z}$. The latent embedding can then be passed through a decoder (blue rectangle) which is a neural network employing transposed convolutions (here denoted CNN^\top) to yield a reconstruction \hat{X}_i of the original input data. The objective function for this network is a variational lower bound on the model evidence of the input data (see Kingma and Welling [23] for details). **Panel B:** Depiction of the latent space learned by the cVAE. This model was a pre-trained model from the (https://colab.research.google.com/github/smartgeometry-ucl/dl4g/blob/master/variational_autoencoder.ipynb, Smart Geometry Processing Group at University College London).

$$f_\phi(\mathbf{z} | \mathbf{x}_i) = \mathcal{N}(\mathbf{z} | \mathbf{m}(\mathbf{x}_i), \mathbf{C}(\mathbf{x}_i)) \quad (55)$$

where ϕ are the encoder’s parameters, which specify a convolutional neural network (CNN) whose output layer returns a 2×1 mean vector $\mathbf{m}(\mathbf{x}_i)$ and a 2×1 log-variance vector $\mathbf{s}(\mathbf{x}_i)$ given \mathbf{x}_i . For simplicity, we denote the latter as the 2×2 diagonal covariance matrix $\mathbf{C}(\mathbf{x}_i) = \left(e^{s_j(\mathbf{x}_i)} \delta_{jk} \right)_{j=1,2}^{k=1,2}$. Further details of the cVAE and its training can be found in Kingma and Welling [23,24], although the specific implementation in this paper was a pre-trained implementation by the (https://colab.research.google.com/github/smartgeometry-ucl/dl4g/blob/master/variational_autoencoder.ipynb, Smart Geometry Processing Group at University College London). Briefly, the cVAE learns to generate a compressed latent representation (via encoder f_ϕ , which is an approximate posterior distribution) that contains enough information about the input \mathbf{x}_i to facilitate its reconstruction by a “decoder” module. The objective function is a lower bound on the model evidence $p(\mathbf{x})$, which if maximized is equivalent to minimizing the Kullback–Leibler divergence between the approximate and true (but unknown) posteriors f_ϕ and $p(\mathbf{z} | \mathbf{x})$, respectively.

The continuous RRH under the model in Equation (55) for a single example $\mathbf{x}_i \in \mathcal{X}$ can be computed by merely evaluating the Rényi heterogeneity of a multivariate Gaussian (Equation (43) in Example 2) for the covariance matrix given by $\mathbf{C}(\mathbf{x}_i)$. This is interpreted as the effective area of the 2-dimensional latent space consumed by representation of \mathbf{x}_i .

Since the handwritten digit images belong to groups of “Zeros, Ones, Twos, ..., Nines,” this section will call the quantity Π_q^W the within-observation heterogeneity (rather than the “within-group” heterogeneity) in order to avoid its interpretation as measuring the heterogeneity of a group of digits. Rather, it is interpreted as the effective area of latent space consumed by representation of a single observation $\mathbf{x} \in \mathcal{X}$ on average. It is computed by evaluation of Equation (44) at $\mathbf{C}(\mathbf{X}) = \{\mathbf{C}(\mathbf{x}_i)\}_{i=1,2,\dots,N}$, given uniform weights on samples.

Finally, to compute the pooled heterogeneity Π_q^P , we use the parametric pooling approach detailed in Example 2, wherein the pooled distribution is a multivariate Gaussian with mean and covariance given by Equations (41) and (42), respectively. The pooled heterogeneity is then merely Equation (43) evaluated at $\mathbf{C}_*(\mathbf{X})$, and represents the total amount of area in the latent space consumed by the representation of X under f_ϕ . The effective number of observations in X with respect to the continuous latent representation Z is, therefore, given by the between-observation heterogeneity:

$$\Pi_q^B(\mathbf{C}(\mathbf{X}), \mathbf{w}) = \frac{\Pi_q^P(\mathbf{C}_*(\mathbf{X}))}{\Pi_q^W(\mathbf{C}(\mathbf{X}), \mathbf{w})}. \quad (56)$$

Equation (56) gives the effective number of observations in X because it uses the entire sample \mathbf{X} (of course, assuming \mathbf{X} provides adequate coverage of the observable event space). However, one could compute the effective number of observations in a subset of \mathbf{X} , if necessary. Let $\mathbf{X}^{(j)} = (\mathbf{x}_k)_{k=1,2,\dots,N_j}$ be the subset of N_j points in \mathbf{X} found in the observable subspace $\mathcal{X}_j \subset \mathcal{X}$ (such as the subspace of MNIST digits corresponding to a given digit class). Given corresponding weights $\mathbf{w}^{(j)} = (N_j^{-1})_{k=1,2,\dots,N_j}$, Equation (56) is then simply

$$\Pi_q^B(\mathbf{C}(\mathbf{X}^{(j)}), \mathbf{w}^{(j)}) = \frac{\Pi_q^P(\mathbf{C}_*(\mathbf{X}^{(j)}))}{\Pi_q^W(\mathbf{C}(\mathbf{X}), \mathbf{w}^{(j)})}. \quad (57)$$

Figure 8 shows the effective number of observations in the subsets of MNIST images belonging to each image class, under the continuous representation learned by the cVAE. One can appreciate that the MNIST class of “Ones” (in the training set) has the smallest effective number of observations. Subjective visual inspection of the MNIST samples in Figure 6 may suggest that the Ones are indeed relatively more homogeneous as a group than the other digits (this claim is given further objective support in Appendix C based on deep similarity metric learning [47,48]).

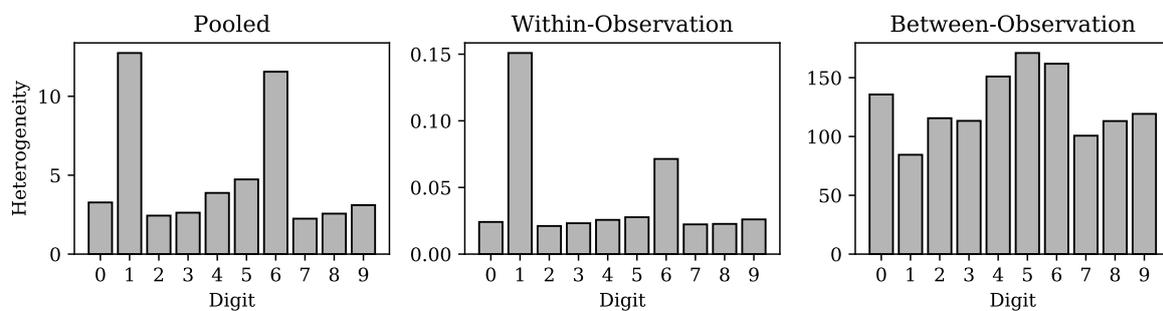
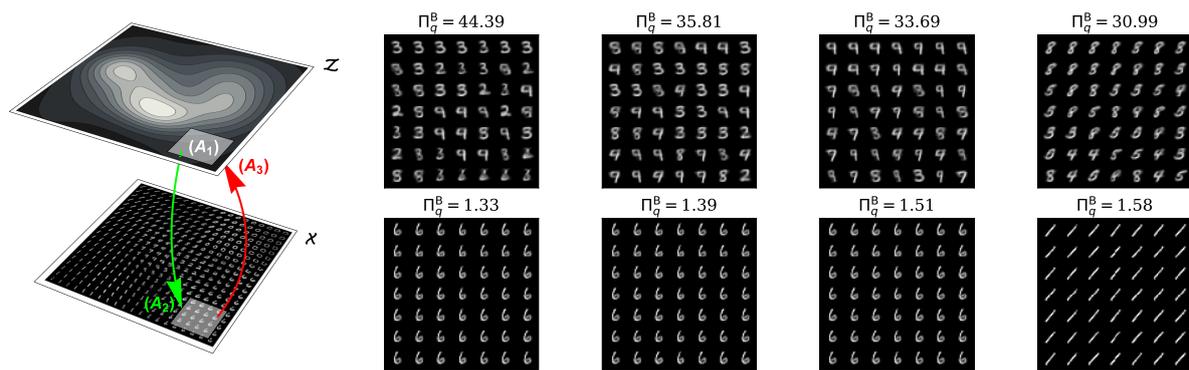


Figure 8. Heterogeneity for the subset of MNIST training data belonging to each digit class respectively projected onto the latent space of the convolutional variational autoencoder (cVAE). The leftmost plot shows the pooled heterogeneity for each digit class (the effective total area of latent space occupied by encoding each digit class). The middle plot shows the within-observation heterogeneity (the effective total area of latent space per encoded observation of each digit class, respectively). The rightmost plot shows the between-observation heterogeneity (the effective number of observations per digit class). Recall that Rényi heterogeneity on a continuous distribution gives the effective size of the domain of an equally heterogeneous uniform distribution on the same space, which explains why the within-observation heterogeneity values here are less than 1.

Figure 9 demonstrates the correspondence of between-observation heterogeneity (i.e., the effective number of observations) and the visual diversity of different samples from the latent space of our cVAE model. For each image in the MNIST training dataset, we computed the effective location of its latent representation: $\mathbf{m}(\mathbf{x}_i)$ for $i \in \{1, 2, \dots, N\}$. For each of these image representations,

we defined a “neighborhood” including the 49 other images whose latent coordinates were closest in Euclidean distance (which is sensible on the latent space given the Gaussian prior). For all such neighbourhoods defined, we then reconstructed the corresponding images on \mathcal{X} , whose between-observation heterogeneity was then computed using Equation (57). Figure 9b shows the estimated effective number of observations for the latent neighborhoods with the greatest and least heterogeneity. One can appreciate that neighborhoods with Π_q^B close to 1 include images with considerably less diversity than neighborhoods with Π_q^B closer to the upper limit of 49. These data suggest that the between-observation heterogeneity—which is the effective number of observations in X with respect to the latent features learned by a cVAE—can indeed correspond to visually appreciable sample diversity.



(a) Illustration of analysis.

(b) Heterogeneity of patches in the latent space.

Figure 9. Visual illustration of MNIST image samples corresponding to different levels of representational Rényi heterogeneity under the convolutional variational autoencoder (cVAE). **Panel (a)** illustrates the approach to this analysis. Here, the surface Z shows hypothetical contours of a probability distribution over the 2-dimensional latent feature space. The surface \mathcal{X} represents the observable space, upon which we have projected an “image” of the latent space Z for illustrative purposes. We first compute the expected latent locations $\mathbf{m}(x_i)$ for each image $x_i \in \mathcal{X}$. **(A₁)** We then define the latent neighbourhood of image x_i as the 49 images whose latent locations are closest to $\mathbf{m}(x_i)$ in Euclidean distance. **(A₂)** Each coordinate in the neighbourhood of $\mathbf{m}(x_i)$ is then projected onto a corresponding patch on the observable space of images. **(A₃)** These images are then projected as a group back onto the latent space, where Equation (57) can be applied, given equal weights over images, to compute the effective number of observations in the neighbourhood of x_i . **Panel (b)** plots the most and least heterogeneous neighbourhoods so that we may compare the estimated effective number of observations with the visually appreciable sample diversity.

5. Discussion

This paper introduced representational Rényi heterogeneity, a measurement approach that satisfies the replication principle [1,38,41] and is decomposable [28] while requiring neither a priori (A) categorical partitioning nor (B) specification of a distance function on the input space. Rather, the experimenter is free to define a model that maps observable data onto a semantically relevant domain upon which Rényi heterogeneity may be tractably computed, and where a distance function need not be explicitly manipulated. These properties facilitate heterogeneity measurements for several new applications. Compared to state-of-the-art comparator indices under a beta mixture distribution, RRH more reliably quantified the number of unique mixture components (Section 4.1), and under a deep generative model of image data, RRH was able to measure the effective number of distinct images with respect to latent continuous representations (Section 4.2). In this section, we further synthesize our conclusions, discuss their implications, and highlight open questions for future research.

The main problem we set out to address was that all state of the art numbers equivalent heterogeneity measures (Section 2.2) require a priori specification of a distance function and categorical partitioning on the observable space. To this end, we showed that RRH does not require categorical

partitioning of the input space (Section 3). Although our analysis under the two-component BMM assumed that the number of components was known, RRH was the only index able to accurately identify an effectively singular cluster (i.e., where mixture components overlapped; Figure 5). We also showed that the categorical RRH did not violate the principle of transfers [39,40] (i.e., it was strictly concave with respect to mixture component weights), unlike the functional Hill numbers (Figure 5). Future studies should extend this evaluation to mixtures of other distributional forms in order to better characterize the generalizability of our conclusions.

Sections 3.1 and 3.2 both showed that RRH does not require specification of a distance function on the observable space. Instead, one must specify a model that maps the observable space onto a probability distribution over the latent representation. This is beneficial since input space distances are often irrelevant or misleading. For example, latent representations of image data learned by a convolutional neural network will be robust to translations of the inputs since convolution is translation invariant. However, pairwise distances on the observable space will be exquisitely sensitive to semantically irrelevant translations of input data. Furthermore, semantically relevant information must often be learned from raw data using hierarchical abstraction. Ultimately, when (A) pre-defined distance metrics are sensitive to noisy perturbations of the input space, or (B) the relevant semantic content of some input data is best captured by a latent abstraction, the RRH measure will be particularly useful.

The requirement of specifying a representational model $f : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Z})$ implies the additional problem of model selection. In Section 3, we noted that the determination of whether a model is appropriate must be made in a domain-specific fashion. For instance, the method by which ecologists assign species labels prior to measurement of species diversity implies the use of a mapping from the observable space of organisms to a degenerate distribution over species labels (Example 1). In Section 4.2, we used the encoder module of a cVAE (a generative model based on a convolutional neural network architecture [23,24]) to represent images as 2-dimensional real-valued vectors in order to demonstrate our ability to capture variation in digits' written forms (see Figures 7B and 9). Someone concerned with measuring heterogeneity of image batches in terms of the digit-class distribution could choose a categorical latent representation corresponding to the digit classes (this would return the effective number of digit classes per sample). Regardless, the model used to map between observations and the latent space should be validated using either explanatory power (e.g., maximization of a lower bound on the model evidence), generalizability (e.g., out of sample predictive power), or another approach that is justifiable within the investigator's scientific domain of interest.

In addition to the results of empirical applications of RRH in Section 4, we were also able to show that RRH generalizes the process by which species diversity and indices of economic equality are computed (Example 1). In doing so, we are able to clarify some of the assumptions inherent in those indices. Specifically, that assignment of species or ownership labels (in ecological and economic settings, respectively) corresponds to mapping from an observable space, such as the space of organisms' identifiable features or the space of economic resources, onto a degenerate distribution over the categorical labels (Table 2). It is possible that altering the form of that mapping may yield new insights about ecological and economic diversity.

In conclusion, we have introduced an approach for measuring heterogeneity that requires neither (A) categorical partitioning nor (B) distance measure on the observable space. Our RRH method enables measurement of heterogeneity in disciplines where categorical entities are unreliably defined, or where relevant semantic content of some data is best captured by a hierarchical abstraction. Furthermore, our approach includes many existing heterogeneity indices as special cases, while facilitating clarification of many of their assumptions. Future work should evaluate the RRH in practice and under a broader array of models.

Supplementary Materials: The following are available online at <http://www.mdpi.com/1099-4300/22/4/417/s1>, Supplementary materials include code for Sections 2–4 and Appendix B (RRH_Supplement_3State_BMM_CVAE.ipynb), and Appendix C (RRH_Supplement_Siamese.ipynb).

Author Contributions: Conceptualization, A.N.; methodology, A.N.; validation, A.N.; formal analysis, A.N.; investigation, A.N.; resources, T.T.; data curation, A.N.; writing—original draft preparation, A.N.; writing—review and editing, A.N., M.A., T.B., T.T.; visualization, A.N.; supervision, M.A., T.B., T.T.; project administration, A.N.; funding acquisition, A.N., M.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Genome Canada (A.N., M.A.), the Nova Scotia Health Research Foundation (A.N.), the Killam Trusts (A.N.), and the Ruth Wagner Memorial Fund (A.N.).

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Appendix A. Mathematical Appendix

Proposition A1. Rényi heterogeneity (Equation (3)) obeys the replication principle.

Proof. The Rényi heterogeneity for a single distribution $\mathbf{p}_i = (p_{ij})_{j=1,2,\dots,n_i}$, where $n_i \in \mathbb{N}_+$ is the size of the state space in system i , is

$$\Pi_q(\mathbf{p}_i) = \left(\sum_{j=1}^{n_i} p_{ij}^q \right)^{\frac{1}{1-q}} \tag{A1}$$

and for the aggregation of N subsystems is

$$\Pi_q(\bar{\mathbf{p}}_i) = \left(\sum_{i=1}^N \sum_{j=1}^{n_i} \left(\frac{p_{ij}}{N} \right)^q \right)^{\frac{1}{1-q}}. \tag{A2}$$

The replication principle asserts that

$$\Pi_q(\bar{\mathbf{p}}) = N\Pi_q(\mathbf{p}_i). \tag{A3}$$

Let $\lambda_i = \sum_{j=1}^{n_i} p_{ij}^q$ and recall that $\lambda_i = \lambda_k$ for all $(i, k) \in \{1, 2, \dots, N\}$. Then,

$$\begin{aligned} \left(N^{-q} \sum_{i=1}^N \sum_{j=1}^{n_i} p_{ij}^q \right)^{\frac{1}{1-q}} &= N \left(\sum_{j=1}^{n_i} p_{ij}^q \right)^{\frac{1}{1-q}} \\ \left(N^{-q} \sum_{i=1}^N \lambda_i \right)^{\frac{1}{1-q}} &= N\lambda_i^{\frac{1}{1-q}} \\ \left(N^{1-q} \lambda_i \right)^{\frac{1}{1-q}} &= N\lambda_i^{\frac{1}{1-q}} \\ N\lambda_i^{\frac{1}{1-q}} &= N\lambda_i^{\frac{1}{1-q}}. \end{aligned} \tag{A4}$$

Since $\lim_{q \rightarrow 1} \lambda_i^{\frac{1}{1-q}}$ exists (it is the perplexity index), the result also holds at $q = 1$. \square

Proposition A2. For a system X with probability mass function represented by the vector $\mathbf{p} = (p_i)_{i=1,2,\dots,n}$ on event space $\mathcal{X} = \{1, 2, \dots, n\}$, with distance function $d_X : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ represented by the $n \times n$ matrix $\mathbf{D} = [d_X(i, j)]_{i=1,2,\dots,n}^{j=1,2,\dots,n}$, the functional Hill numbers family of indices

$$F_q(\mathbf{D}, \mathbf{p}) = \left(\frac{Q_q(\mathbf{D}, \mathbf{p})}{Q_1(\mathbf{D}, \mathbf{p})} \right)^{\frac{1}{2(1-q)}} \tag{A5}$$

is insensitive to $d_X(i, j)$ for all $(i, j) \in \mathcal{X}$ when \mathbf{p} is uniform.

Proof. The proof is direct given substitution of $\mathbf{p} = (n^{-1})_{i=1,2,\dots,n}$ into Equation (A5).

$$F_q(\mathbf{D}, \mathbf{p}) = \left(\frac{Q_q(\mathbf{D}, \mathbf{p})}{Q_1(\mathbf{D}, \mathbf{p})} \right)^{\frac{1}{2(1-q)}} = \left(\frac{n^{-2q} \sum_{i=1}^n \sum_{j=1}^n d_X(i, j)}{n^{-2} \sum_{i=1}^n \sum_{j=1}^n d_X(i, j)} \right)^{\frac{1}{2(1-q)}} = n \tag{A6}$$

□

Proposition A3 (Rényi Heterogeneity of a Continuous System). *The Rényi heterogeneity of a system X with event space $\mathcal{X} \subseteq \mathbb{R}^n$ and pdf $f \in \mathcal{P}(\mathcal{X})$ is equal to the magnitude of the volume of an n -cube over which there is a uniform probability density with the same Rényi heterogeneity as that given by f .*

Proof. Let the basic integral of X be defined as $\int_{\mathcal{X}} f^q(\mathbf{x}) \, d\mathbf{x}$. Furthermore, let X_* be an idealized reference system with a uniform probability density f_* on \mathcal{X} with lower bounds $\mathbf{0} = (0)_{i=1, \dots, n}$ and upper bounds $\mathbf{u} = (u_*)_{i=1, \dots, n}$ where $u_* \geq 0$ is the side length of an n -cube. We assume that X_* has basic integral $\int_{\mathcal{X}} f_*^q(\mathbf{x}) \, d\mathbf{x}$ such that

$$\begin{aligned} \int_{\mathcal{X}} f^q(\mathbf{x}) \, d\mathbf{x} &= \int_{\mathcal{X}} f_*^q(\mathbf{x}) \, d\mathbf{x} \\ &= \prod_{i=1}^n u_*^{1-q} \\ &= u_*^{n(1-q)}. \end{aligned} \tag{A7}$$

Solving Equation (A7) for u_*^n gives the Rényi heterogeneity of order q . At $q \neq 1$,

$$u_*^n = \left(\int_{\mathcal{X}} f^q(\mathbf{x}) \, d\mathbf{x} \right)^{\frac{1}{1-q}} \tag{A8}$$

and in the limit of $q \rightarrow 1$, Equation (A8) becomes the exponential of the Shannon (differential) entropy. Thus, Π_q is interpreted as the volume of an n -cube of side length u_* , over which there is a uniform distribution giving the same heterogeneity as X . □

Proposition A4 (Rényi heterogeneity of a multivariate Gaussian). *The Rényi heterogeneity of an n -dimensional multivariate Gaussian with probability density function (pdf)*

$$f(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma}) = (2\pi)^{-\frac{n}{2}} |\boldsymbol{\Sigma}|^{-\frac{1}{2}} e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})}, \tag{A9}$$

with mean $\boldsymbol{\mu} = (\mu_i)_{i=1,2, \dots, n}$ and covariance matrix $\boldsymbol{\Sigma} = (\Sigma_{ij})_{i=1,2, \dots, n}^{j=1,2, \dots, n}$ is

$$\Pi_q(\boldsymbol{\Sigma}) = \begin{cases} \text{Undefined} & q = 0 \\ (2\pi e)^{\frac{n}{2}} \sqrt{|\boldsymbol{\Sigma}|} & q = 1 \\ (2\pi)^{\frac{n}{2}} \sqrt{|\boldsymbol{\Sigma}|} & q = \infty \\ (2\pi)^{\frac{n}{2}} q^{\frac{n}{2(q-1)}} \sqrt{|\boldsymbol{\Sigma}|} & \text{Otherwise} \end{cases}. \tag{A10}$$

Proof. Let $\boldsymbol{\Sigma}^{-1} = \mathbf{U}\boldsymbol{\Lambda}\mathbf{U}^{-1}$ be the eigendecomposition of the inverse covariance matrix into an orthonormal matrix of eigenvectors \mathbf{U} and $n \times n$ diagonal matrix $\boldsymbol{\Lambda}$ with eigenvalues $(\lambda_i)_{i=1,2, \dots, n}$

down the leading diagonal. Furthermore, let $\frac{dx_i}{dy_j} = U_{ij}$ and use the substitution $\mathbf{y} = \mathbf{U}^{-1}(\mathbf{x} - \boldsymbol{\mu})$ to proceed as follows:

$$\begin{aligned} \Pi_q(\boldsymbol{\Sigma}) &= \left[(2\pi)^{-\frac{qn}{2}} |\boldsymbol{\Sigma}|^{-\frac{q}{2}} \int e^{-\frac{q}{2}(\mathbf{x}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})} d\mathbf{x} \right]^{\frac{1}{1-q}} \\ &= \left((2\pi)^{-\frac{qn}{2}} |\boldsymbol{\Sigma}|^{-\frac{q}{2}} \int e^{-\frac{q}{2}\mathbf{y}^\top \boldsymbol{\Lambda} \mathbf{y}} d\mathbf{y} \right)^{\frac{1}{1-q}} \\ &= \left((2\pi)^{-\frac{qn}{2}} |\boldsymbol{\Sigma}|^{-\frac{q}{2}} \left(\frac{(2\pi)^n}{q^n \prod_{i=1}^n \lambda_i} \right)^{\frac{1}{2}} \right)^{\frac{1}{1-q}} \\ &= \left((2\pi)^{-\frac{qn}{2}} |\boldsymbol{\Sigma}|^{-\frac{q}{2}} \left(\frac{(2\pi)^n}{q^n |\boldsymbol{\Lambda}|} \right)^{\frac{1}{2}} \right)^{\frac{1}{1-q}} \\ &= q^{\frac{n}{2(q-1)}} (2\pi)^{\frac{n}{2}} \sqrt{|\boldsymbol{\Sigma}|} \end{aligned} \tag{A11}$$

which holds only at $q \notin \{0, 1, \infty\}$. At $q = 1$, we have

$$\begin{aligned} \lim_{q \rightarrow 1} \log \Pi_q(\boldsymbol{\Sigma}) &= \lim_{q \rightarrow 1} \left(\frac{n}{2(q-1)} \log q \right) + \frac{n}{2} \log(2\pi) + \frac{1}{2} \log |\boldsymbol{\Sigma}| \\ &= \frac{n}{2} + \frac{n}{2} \log(2\pi) + \frac{1}{2} \log |\boldsymbol{\Sigma}|, \end{aligned} \tag{A12}$$

and therefore,

$$\Pi_1(\boldsymbol{\Sigma}) = (2\pi e)^{\frac{n}{2}} \sqrt{|\boldsymbol{\Sigma}|}. \tag{A13}$$

One can then easily show that $\Pi_0(\boldsymbol{\Sigma})$ is undefined and that as $q \rightarrow \infty$,

$$\Pi_\infty(\boldsymbol{\Sigma}) = (2\pi)^{\frac{n}{2}} \sqrt{|\boldsymbol{\Sigma}|}. \tag{A14}$$

□

Appendix B. Expected Distance Between two Beta-Distributed Random Variables

To compute the numbers equivalent RQE \hat{Q}_e , the functional Hill numbers F_q , and the Leinster-Cobbold index L_q under the beta mixture model, we must derive an analytical expression for the distance matrix. This involves the following integral:

$$d(x, y) = \int_0^1 \int_0^1 |x - y| f(x) g(y) dx dy, \tag{A15}$$

where $f(x) = \text{Beta}_{\alpha_1, \beta_1}(x)$ and $g(y) = \text{Beta}_{\alpha_2, \beta_2}(y)$. By exploiting the identity

$$|x - y| = x + y - 2 \min\{x, y\}, \tag{A16}$$

and expanding, the integral is greatly simplified and gives the following closed-form solution:

$$d(x, y) = \langle x \rangle - \langle y \rangle + \eta (\Phi_a - \alpha_1 \Phi_b), \tag{A17}$$

where

$$\eta = \frac{2\Gamma(\alpha_1)\Gamma(\beta_2)\Gamma(\alpha_1 + \alpha_2 + 1)}{B(\alpha_1, \beta_1)B(\alpha_2, \beta_2)}, \tag{A18}$$

and where $\langle y \rangle = \frac{\alpha_2}{\alpha_2 + \beta_2}$, $\langle x \rangle = \frac{\alpha_1}{\alpha_1 + \beta_1}$, and the Φ 's are regularized hypergeometric functions:

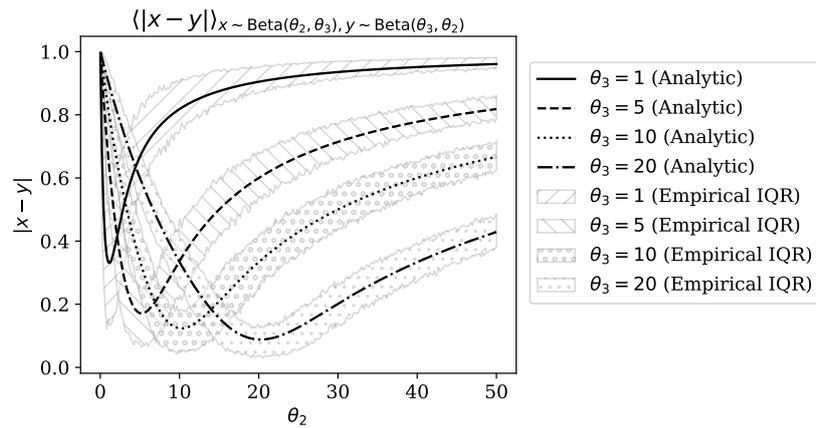


Figure A1. Numerical verification of the analytical expression for the expected absolute distance between two Beta-distributed random variables. Solid lines are the theoretical predictions. Ribbons show the bounds between 25th–75th percentiles (the interquartile range, IQR) of the simulated values.

$$\Phi_a = {}_3\tilde{F}_2 \left[\begin{matrix} \alpha_1, \alpha_1 + \alpha_2 + 1, 1 - \beta_1 \\ \alpha_1 + 1, \alpha_1 + \alpha_2 + \beta_2 + 1 \end{matrix} , 1 \right] \tag{A19}$$

$$\Phi_b = {}_3\tilde{F}_2 \left[\begin{matrix} \alpha_1 + 1, \alpha_1 + \alpha_2 + 1, 1 - \beta_1 \\ \alpha_1 + 2, \alpha_1 + \alpha_2 + \beta_2 + 1 \end{matrix} , 1 \right] \tag{A20}$$

Figure A1 provides numerical verification of this result. One simply uses Equation (A17) to compute the analytic distance matrix

$$\mathbf{D}(\alpha_1, \beta_1, \alpha_2, \beta_2) = \begin{pmatrix} d(x, x) & d(x, y) \\ d(y, x) & d(y, y) \end{pmatrix}, \tag{A21}$$

which, with the component probabilities (Equation (46)), can be used to compute $\hat{Q}_e, F_q,$ and L_q using the formulas shown in the main body.

Appendix C. Evidence Supporting Relative Homogeneity of MNIST “Ones”

In our evaluation of non-categorical RRH using the MNIST data, we asserted that the class of handwritten Ones were relatively more homogeneous than other digits. Our initial statement was based simply on visual inspection of samples from the dataset, wherein the Ones ostensibly demonstrate fewer relevant feature variations than other classes. However, to test this hypothesis more objectively, we conducted an empirical evaluation using similarity metric learning.

We implemented a deep neural network architecture known as a “siamese network” [47] to learn a latent distance metric on the MNIST classes. Our siamese network architecture is depicted in Figure A2a. Training is conducted by sampling batches of 10,000 image pairs from the MNIST test set, where 5000 pairs are drawn from the same class (i.e., a pair of Fives or a pair of Threes), and 5000 pairs are drawn from different classes (i.e., the pairs [2,3] or [1,7]). The siamese network is then optimized using gradient-based methods over 100 epochs using the contrastive loss function [48] (Figure A2a). This analysis may be reproduced in the Supplementary Materials.

After training, we sampled same-class pairs ($n = 25,000$) and different-class pairs ($n = 25,000$) from the MNIST training set (which contains 60,000 images). Pairwise distances for each sample were computed using the trained siamese network. If the “ones” are indeed the most homogeneous class, they should demonstrate a generally smaller pairwise distance than other digit class pairs. We evaluated this hypothesis by comparing empirical cumulative distribution functions (CDF) on

the class-pair distances (Figure A2b). Our results show that the empirical CDF for “1–1” image pairs dominate that of all other class pairs (where the distance between pairs of “ones” is lower).

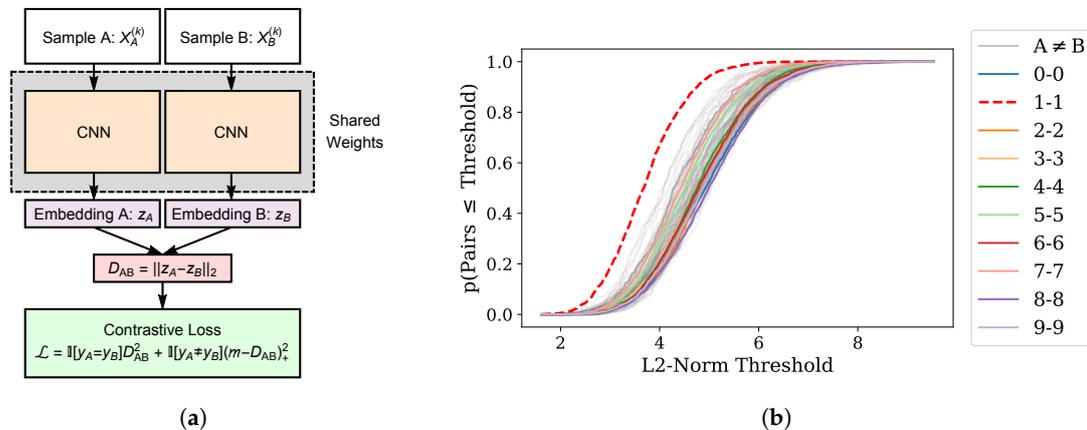


Figure A2. Depiction of the siamese network architecture and the empirical cumulative distribution function for pairwise distances between digit classes. **(a)** Depiction of a siamese network architecture. At iteration k , each of two samples, $X_A^{(k)}$ and $X_B^{(k)}$, are passed through a convolutional neural network to yield embeddings z_A and z_B , respectively. The class label for samples A and B are denoted y_A and y_B , respectively. The L2-norm of these embeddings is computed as D_{AB} . The network is optimized on the contrastive loss [48] \mathcal{L} . Here $\mathbb{I}[\cdot]$ is an indicator function, **(b)** Empirical cumulative distribution functions (CDF) for pairwise distances between images of the listed classes under the siamese network model. The x-axis plots the L2-norm between embedding vectors produced by the siamese network. The y-axis shows the proportion of samples in the respective group (by line color) whose embedded L2 norms were less than the specified threshold on the x-axis. Class groups are denoted by different line colors. For instance, “0-0” refers to pairs where each image is a “zero.” We combine all disjoint class pairs, for example “0–8” or “3–4,” into a single empirical CDF denoted as “A≠B.”

References

- Jost, L. Entropy and diversity. *Oikos* **2006**, *113*, 363–375. [\[CrossRef\]](#)
- Prehn-Kristensen, A.; Zimmermann, A.; Tittmann, L.; Lieb, W.; Schreiber, S.; Baving, L.; Fischer, A. Reduced microbiome alpha diversity in young patients with ADHD. *PLoS ONE* **2018**, *13*, e0200728. [\[CrossRef\]](#) [\[PubMed\]](#)
- Cowell, F. *Measuring Inequality*, 2nd ed.; Oxford University Press: Oxford, UK, 2011.
- Higgins, J.P.T.; Thompson, S.G.; Deeks, J.J.; Altman, D.G. Measuring inconsistency in meta-analyses. *BMJ Br. Med. J.* **2003**, *327*, 557–560. [\[CrossRef\]](#)
- Hooper, D.U.; Chapin, F.S.; Ewel, J.J.; Hector, A.; Inchausti, P.; Lavorel, S.; Lawton, J.H.; Lodge, D.M.; Loreau, M.; Naeem, S.; et al. Effects of biodiversity on ecosystem functioning: A consensus of current knowledge. *Ecol. Monogr.* **2005**, *75*, 3–35. [\[CrossRef\]](#)
- Botta-Dukát, Z. The generalized replication principle and the partitioning of functional diversity into independent alpha and beta components. *Ecography* **2018**, *41*, 40–50. [\[CrossRef\]](#)
- Mouchet, M.A.; Villéger, S.; Mason, N.W.; Mouillot, D. Functional diversity measures: An overview of their redundancy and their ability to discriminate community assembly rules. *Funct. Ecol.* **2010**, *24*, 867–876. [\[CrossRef\]](#)
- Chiu, C.H.; Chao, A. Distance-based functional diversity measures and their decomposition: A framework based on hill numbers. *PLoS ONE* **2014**, *9*, e113561. [\[CrossRef\]](#) [\[PubMed\]](#)
- Petchey, O.L.; Gaston, K.J. Functional diversity (FD), species richness and community composition. *Ecol. Lett.* **2002**. [\[CrossRef\]](#)
- Leinster, T.; Cobbold, C.A. Measuring diversity: The importance of species similarity. *Ecology* **2012**, *93*, 477–489. [\[CrossRef\]](#)

11. Chao, A.; Chiu, C.H.; Jost, L. Unifying Species Diversity, Phylogenetic Diversity, Functional Diversity, and Related Similarity and Differentiation Measures Through Hill Numbers. *Annu. Rev. Ecol. Evol. Syst.* **2014**, *45*, 297–324. [[CrossRef](#)]
12. American Psychiatric Association. *Diagnostic and Statistical Manual of Mental Disorders*, 5th ed.; American Psychiatric Publishing: Washington, DC, USA, 2013.
13. Regier, D.A.; Narrow, W.E.; Clarke, D.E.; Kraemer, H.C.; Kuramoto, S.J.; Kuhl, E.A.; Kupfer, D.J. DSM-5 field trials in the United States and Canada, part II: Test-retest reliability of selected categorical diagnoses. *Am. J. Psychiatr.* **2013**, *170*, 59–70. [[CrossRef](#)] [[PubMed](#)]
14. Bengio, Y.; Courville, A.; Vincent, P. Representation learning: A review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.* **2013**, *35*, 1798–1828. [[CrossRef](#)] [[PubMed](#)]
15. Arvanitidis, G.; Hansen, L.K.; Hauberg, S. Latent Space Oddity: On the Curvature of Deep Generative Models. In Proceedings of the International Conference on Learning Representations, Vancouver, BC, Canada, 30 April–3 May 2018; pp. 1–15, [[1710.11379](#)].
16. Shao, H.; Kumar, A.; Thomas Fletcher, P. The Riemannian geometry of deep generative models. In Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, Salt Lake City, UT, USA, 18–22 June 2018.
17. Nickel, M.; Kiela, D. Poincaré embeddings for learning hierarchical representations. In Proceedings of the Advances in Neural Information Processing Systems, Long Beach, CA, USA, 4–9 December 2017; pp. 6339–6348.
18. Rényi, A. On measures of information and entropy. *Proc. Fourth Berkeley Symp. Math. Stat. Probab.* **1961**, *114*, 547–561.
19. Hill, M.O. Diversity and Evenness: A Unifying Notation and Its Consequences. *Ecology* **1973**, *54*, 427–432. [[CrossRef](#)]
20. Hannah, L.; Kay, J.A. *Concentration in Modern Industry: Theory, Measurement and The U.K. Experience*; The MacMillan Press, Ltd.: London, UK, 1977.
21. Ricotta, C.; Szeidl, L. Diversity partitioning of Rao’s quadratic entropy. *Theor. Popul. Biol.* **2009**, *76*, 299–302. [[CrossRef](#)] [[PubMed](#)]
22. LeCun, Y.; Bottou, L.; Bengio, Y.; Haffner, P. Gradient-based learning applied to document recognition. *Proc. IEEE* **1998**, *86*, 2278–2324. [[CrossRef](#)]
23. Kingma, D.P.; Welling, M. Auto-Encoding Variational Bayes. *ICLR 2014* **2014**, arXiv:1312.6114v10.
24. Kingma, D.P.; Welling, M. An Introduction to Variational Autoencoders. *Found. Trend. Mach. Learn.* **2019**, *12*, 307–392. [[CrossRef](#)]
25. Eliazar, I.I.; Sokolov, I.M. Measuring statistical evenness: A panoramic overview. *Phys. A Stat. Mech. Its Appl.* **2012**, *391*, 1323–1353. [[CrossRef](#)]
26. Patil, A.G.P.; Taillie, C. Diversity as a Concept and its Measurement. *J. Am. Stat. Assoc.* **1982**, *77*, 548–561. [[CrossRef](#)]
27. Adelman, M.A. Comment on the “H” Concentration Measure as a Numbers-Equivalent. *Rev. Econ. Stat.* **1969**, *51*, 99–101. [[CrossRef](#)]
28. Jost, L. Partitioning Diversity into Independent Alpha and Beta Components. *Ecology* **2007**, *88*, 2427–2439. [[CrossRef](#)] [[PubMed](#)]
29. Shannon, C.E. A mathematical theory of communication. *Bell Syst. Tech. J.* **1948**, *27*, 379–423. [[CrossRef](#)]
30. Eliazar, I. How random is a random vector? *Ann. Phys.* **2015**, *363*, 164–184. [[CrossRef](#)]
31. Gotelli, N.J.; Chao, A. Measuring and Estimating Species Richness, Species Diversity, and Biotic Similarity from Sampling Data. In *Encyclopedia of Biodiversity*, 2nd ed.; Levin, S.A., Ed.; Academic Press: Waltham, MA, USA, 2013; pp. 195–211.
32. Berger, W.H.; Parker, F.L. Diversity of planktonic foraminifera in deep-sea sediments. *Science* **1970**, *168*, 1345–1347. [[CrossRef](#)] [[PubMed](#)]
33. Daly, A.; Baetens, J.; De Baets, B. Ecological Diversity: Measuring the Unmeasurable. *Mathematics* **2018**, *6*, 119. [[CrossRef](#)]
34. Tsallis, C. Possible generalization of Boltzmann-Gibbs statistics. *J. Stat. Phys.* **1988**, *52*, 479–487. [[CrossRef](#)]
35. Simpson, E.H. Measurement of Diversity. *Nature* **1949**, *163*, 688. [[CrossRef](#)]
36. Gini, C. *Variabilità e mutabilità. Contributo allo Studio delle Distribuzioni e delle Relazioni Statistiche*; C. Cuppini: Bologna, Italy, 1912.

37. Shorrocks, A.F. The Class of Additively Decomposable Inequality Measures. *Econometrica* **1980**, *48*, 613–625. [[CrossRef](#)]
38. Jost, L. Mismeasuring biological diversity: Response to Hoffmann and Hoffmann (2008). *Ecol. Econ.* **2009**, *68*, 925–928. [[CrossRef](#)]
39. Pigou, A.C. *Wealth and Welfare*; MacMillan and Co., Ltd: London, England, 1912.
40. Dalton, H. The Measurement of the Inequality of Incomes. *Econ. J.* **1920**, *30*, 348. [[CrossRef](#)]
41. Macarthur, R.H. Patterns of species diversity. *Biol. Rev.* **1965**, *40*, 510–533. [[CrossRef](#)]
42. Lande, R. Statistics and partitioning of species diversity and similarity among multiple communities. *Oikos* **1996**, *76*, 5–13. [[CrossRef](#)]
43. Rao, C.R. Diversity and dissimilarity coefficients: A unified approach. *Theor. Popul. Biol.* **1982**, *21*, 24–43. [[CrossRef](#)]
44. Mikolov, T.; Chen, K.; Corrado, G.; Dean, J. Distributed representations of words and phrases and their compositionality. In Proceedings of the NIPS 2013, Lake Tahoe, NV, USA, 5–10 December 2013; pp. 1–9.
45. Pennington, J.; Socher, R.; Manning, C. Glove: Global Vectors for Word Representation. In Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), Doha, Qatar, 25–29 October 2014; pp. 1532–1543. doi:10.3115/v1/D14-1162. [[CrossRef](#)]
46. Nunes, A.; Alda, M.; Trappenberg, T. On the Multiplicative Decomposition of Heterogeneity in Continuous Assemblages. *arXiv* **2020**, arXiv:2002.09734.
47. Bromley, J.; Guyon, I.; LeCun, Y.; Säckinger, E.; Shah, R. Signature verification using a “siamese” time delay neural network. In Proceedings of the Advances in Neural Information Processing Systems 6, Denver, CO, USA, 29 November–2 December 1993; pp. 737–744.
48. Hadsell, R.; Chopra, S.; LeCun, Y. Dimensionality Reduction by Learning an Invariant Mapping. In Proceedings of the CVPR 2006, New York, NY, USA, 17–22 June 2006; pp. 1735–1742.



© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).