Clone-Resistant Weights in Metric Spaces: A Framework for Handling Redundancy Bias

Damien Berriaud and **Roger Wattenhofer**
ETH Zürich
{dberriaud, wattenhofer}@ethz.ch

Abstract

We are given a set of elements in a metric space. The distribution of the elements is arbitrary, possibly adversarial. Can we weigh the elements in a way that is resistant to such (adversarial) manipulations? This problem arises in various contexts. For instance, the elements could represent data points, requiring robust domain adaptation. Alternatively, they might represent tasks to be aggregated into a benchmark; or questions about personal political opinions in voting advice applications. This article introduces a theoretical framework for dealing with such problems. We propose *clone-proof representation functions* as a solution concept. These functions distribute importance across elements of a set such that similar objects ("clones") share (some of) their weights, thus avoiding a potential bias introduced by their multiplicity. Our framework extends the maximum uncertainty principle to accommodate general metric spaces and includes a set of axioms — symmetry, continuity, and clone-proofness — that guide the construction of representation functions. Finally, we address the existence of representation functions satisfying our axioms in the significant case of Euclidean spaces and propose a general method for their construction.

1 Introduction

Morpheus: "You take the **blue** pill and the story ends. You wake up in your bed and believe whatever you want to believe. You take the **red** pill, you stay in Wonderland and I show you how deep the rabbit hole goes." Before Neo can make his choice, Morpheus continues: "Or you can take this **indigo** pill, and wake up with $100 in your pocket. Or this **navy** pill – with a different hair color." Why would Morpheus present these insignificant shades of blue? Neo already feels manipulated, but Morpheus continues excitedly and adds more pills colored in **bordeaux**, **cyan**, and **green**?!

Is there an objective way to make a choice without being bamboozled (see Figure 1)? This problem is at the heart of machine learning as well as various other areas, e.g., distributed systems or social choice. Let us first discuss some application examples:

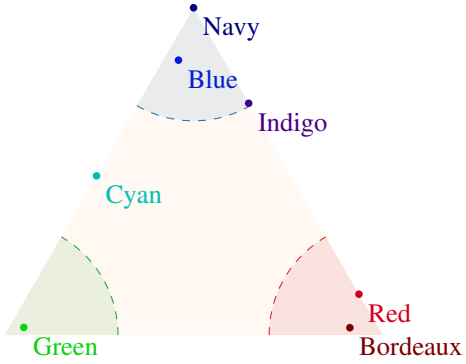

Figure 1: Which weight should we give to each individual point? By symmetry, one would expect the areas in blue, red and green to sum up to the same value, even though they contain different numbers of points. How to deal with the addition of cyan though?

(i) In machine learning, one may want to tackle class imbalance in multi-label classification problems, e.g., by giving weights to the individual contribution of each sample to the loss. More generally, reweighing a loss function according to a clone-proof representation can be thought of as a distribution-agnostic importance sampling technique. As such, it could be beneficial whenever the most informative samples become increasingly difficult to obtain, either because of the *computational cost* associated with computing the associated true label, e.g., for protein folding [Jumper *et al.*, 2021], Graph Neural Networks-based SAT solvers [Wang *et al.*, 2024], climate modeling [Eyring *et al.*, 2024]; or simply because of their *relative rarity*, e.g., for rare disease diagnosis in medical image analysis [Banerjee *et al.*, 2023], fraud detection in financial systems [Motie and Raahemi, 2024], low-resources languages in Natural Language Processing [Hedderich *et al.*, 2021].

(ii) Benchmarking theory presents another natural context. Consider a composite benchmark consisting of different tasks; however, it quickly becomes apparent that many tasks are quite similar. Perhaps they should share some of their weight? Recent research has proposed to revisit benchmarking practices through the lens of social choice theory [Colombo *et al.*, 2022; Rofin *et al.*, 2023; Zhang and Hardt, 2024]. While this approach enables a more accurate aggregation of individual tasks' performance into a final score, we

argue that benchmarking diverges from traditional voting scenarios in a key way: anonymity, or the equal treatment of votes, is not inherently required. Instead, the benchmark's outcome should remain unaffected by the inclusion of numerous highly similar tasks, as this could unfairly favor models that perform well on the original task over those excelling in other areas. Weighing tasks with a clone-proof representation offers a solution to this problem.

(iii) Finally, clone-proof representations offer a novel set of tools to mitigate Sybil attacks in distributed systems. While reputation mechanisms are designed to promote cooperation, significant research has focused on preventing adversaries from exploiting these systems by creating fake identities that mutually reinforce one another [Resnick and Sami, 2009; Seuken and Parkes, 2014; Stannat *et al.*, 2021]. Similarly, researchers have sought to design mechanisms that incentivize information diffusion within social networks without encouraging the creation of fraudulent identities [Babaioff *et al.*, 2016; Chen *et al.*, 2022; Zhang *et al.*, 2020]. In both scenarios, clone-proof representations could help regulate the influence of Sybils once detected, effectively shifting the challenge to identifying these artificially generated identities.

Contribution & Outline This work proposes a mathematical framework for handling redundancy in a metric space. Specifically, we tackle the problem of determining the relative importance of elements in a finite set such that close-by elements, or "clones," share some of their weight. To extend the well-understood case of discrete metrics, where elements are either similar or equally dissimilar, we introduce in Section 3 the concept of representation function and propose a set of axioms that such functions should satisfy in general metric spaces. These properties can be broadly categorized into three key principles: symmetry, continuity, and clone-proofness. Building on these foundations, we address in Section 4 the challenges of constructing functions that adhere to these axioms. For the specific case of Euclidean spaces, we resolve the question of existence and construct, in Theorems 1 and 2, a family of desirable representation functions. We finally explore in Section 5 possible extensions of our construction to more general spaces and discuss the computational hurdles associated with its evaluation.

2 Related Work

Our work intersects with several bodies of literature; we focus hereafter on three of the most relevant topics.

Benchmark Aggregation and Measuring Task Similarity. In a recent line of work [Colombo *et al.*, 2022; Himmi *et al.*, 2023; Rofin *et al.*, 2023; Tatiana and Valentin, 2021; Zhang and Hardt, 2024], multi-tasks benchmarking practices have been scrutinized through the lenses of social choice theory. In particular, these works question the usage of the arithmetical mean to aggregate scores of different tasks in popular benchmarks [Koh *et al.*, 2021; Wang *et al.*, 2020] and investigate different aggregates such as the Pythagorean means [Tatiana and Valentin, 2021], the Bradley-Terry model [Peyrard *et al.*, 2021], or other classical voting rules [Colombo *et al.*, 2022; Himmi *et al.*, 2023; Rofin *et al.*, 2023].

Contrary to usual voting scenarios where the equal treatment of voters is of utmost importance, there is apriori no requirement to treat each task equally in benchmark aggregation scenarios, and we may want to consider voting schemes with different weights, e.g., chosen arbitrarily by the benchmark creator. These weights may however carry more information than arbitrary preference. In [Balduzzi *et al.*, 2018], researchers proposed to model the evaluation of agents on different tasks through a zero-sum meta-game played between an "agent" and a "task" player, each choosing a probability distribution over the corresponding set. Scores on different tasks are then aggregated with a weighted average, where the weights correspond to the probability of playing each task in the entropy-maximizing Nash Equilibrium. One of the desirable properties of this technique is that it is invariant under the addition of *exact* copies of agents, a property which has been studied under the appellation *false-name-proofness* in social choice theory [Conitzer and Yokoo, 2010; Nehama *et al.*, 2022; Todo *et al.*, 2009]. Note however that this property is very brittle and only applies to purely adversarial scenarios, as we lose all guarantees whenever a vanishingly small amount of noise is added to one of the copies. Our work aims exactly at extending this property to near-copies, thus ensuring it remains usable in real-world scenarios. Calculating the weights for different tasks using the representation function introduced in this work ensures an automatic scaling of the benchmark, i.e., the evaluation always benefits from adding new although slightly redundant tasks.

Importantly, our framework assumes that practitioners provide a suitable distance metric, as its selection lies beyond the scope of this work. Identifying an appropriate distance metric between tasks or datasets is a critical prerequisite for applying our approach effectively. Fortunately, this challenge has been extensively explored [Alvarez-Melis and Fusi, 2020; Gretton *et al.*, 2012; Liu *et al.*, 2022], particularly within the transfer learning literature [Achille *et al.*, 2019; Peng *et al.*, 2020]. This existing body of research complements our work and provides valuable guidance for practitioners seeking to use our framework for benchmark aggregation.

Domain Adaptation and Samples Reweighting. In traditional learning setups, training and testing data are assumed to follow the same distribution. *Domain adaptation* [Wang and Deng, 2018], however, addresses scenarios where this assumption is violated, such as in the presence of class imbalance or label noise [Torralba and Efros, 2011].

One approach to handle biases in training datasets involves assigning weights to individual samples and minimizing a weighted loss. Classical algorithms, such as AdaBoost [Freund and Schapire, 1997], hard-negative mining [Chang *et al.*, 2018], self-paced learning [Jiang *et al.*, 2015], adapt these weights dynamically during training based on the observed training loss. In contrast, the *meta-learning* framework [Jamal *et al.*, 2020; Ren *et al.*, 2019; Shu *et al.*, 2019] iteratively optimizes the weighting to minimize loss on a small, unbiased validation dataset. Diverging from these methods, we consider a one-shot scenario, where the sample weighting is determined a priori and remains fixed.

In the context of imbalanced classification and long-tailed

datasets, reweighting techniques have been explored extensively [Cao *et al.*, 2019; Dong *et al.*, 2017; Gebru *et al.*, 2017]. These methods typically assign weights inversely proportional to the number of instances in each class. Recent approaches, such as those proposed in [Cui *et al.*, 2019], go further by accounting for data overlap. They suggest weighting samples based on the effective number of samples, under the intuition that the marginal benefit of adding a new sample diminishes as the sample count grows. Specifically, they expand each data point to include its surrounding neighborhood and define the informativeness of a sample as the additional coverage it contributes compared to the scenario where the sample is excluded. Our theoretical framework, in particular the locality axiom (see Axiom 6) and the volume-based construction (see Section 4), draws a close parallel to their total volume of sampled data, but extends this intuition beyond simple classification problems.

More generally, reweighting is a key technique in addressing *covariate shift* within domain adaptation. Covariate shift occurs when the input distribution differs between training and evaluation datasets, i.e., $P_{train}(x) \neq P_{test}(x)$, but the conditional distribution $P_{train}(y|x) = P_{test}(y|x)$ remains consistent. Originating in importance sampling – a technique commonly used to reduce variance in Monte Carlo estimation – different methods [Y *et al.*, 2019] tackle covariate shift by reweighting samples with the ratio $P_{test}(x)/P_{train}(x)$. Approaches such as *Kernel Density Estimation* [Härdle *et al.*, 2004] approximate these distributions using Gaussian kernels but suffer from the curse of dimensionality. Alternatively, *Kernel Mean Matching* [Gretton *et al.*, 2009] minimizes the discrepancy between training and test distributions by aligning their means in a reproducing kernel Hilbert space, effectively estimating $P_{test}(x)/P_{train}(x)$ directly. In contrast to these statistical methods, our approach does not rely on assumptions about the stochasticity of the sampling process. Instead, we adopt an axiomatic framework that provides robustness guarantees even when samples are adversarially selected. This makes our method more resilient to challenges like dataset poisoning [Carlini *et al.*, 2024].

Metric Learning and Hierarchical clustering. Metric learning and clustering techniques adopt fundamentally opposing philosophies. On the one hand, metric learning generally assumes access to ground-truth labels of similarity, e.g., whether points belong to the same class, and seeks to derive a distance metric, often within the class of generalized Mahalanobis metrics, that best separates dissimilar points while bringing similar ones closer [Ghojogh *et al.*, 2022]. On the other hand, clustering assumes some ground truth distance metric and aims to recover a notion of class by grouping similar points into clusters. Our framework aligns more closely with clustering techniques through its shared starting assumption – the availability of an informative distance metric. However, it diverges in its objective, focusing instead on the implications of similarity for unbiased weighting of data points.

Still, clustering techniques may represent a useful step toward this goal: one could first group points into clusters, assign clusters equal weights, and then share these weights uniformly within each cluster. Such "hard" clusters may however lack the smooth properties we aim for. These limitations can be mitigated by adopting hierarchical clustering techniques, where points can belong to multiple clusters arranged in a tree structure (dendrogram) [Murtagh and Contreras, 2017; Ran *et al.*, 2023]. This approach enables contributions across different scales, akin to the role of the probability distribution ν in Theorem 2.

One may even consider more flexible structures, such as *Fuzzy Hierarchical Clustering (FHC)* or *Overlapping Hierarchical Clustering (OHC)*. FHC [Varshney *et al.*, 2022] allows points to have partial membership in several clusters at once, with the sum of memberships normalized across clusters. OHC [Jeantet *et al.*, 2020], on the other hand, constructs directed acyclic graphs of clusters (quasi-dendrograms) instead of traditional trees, enabling a soft merging process. This approach allows points to have full membership in multiple clusters simultaneously, letting clusters overlap without the need for fuzzy memberships.

In any case, transitioning from clusters of points to individual weights becomes a non-trivial task for "soft" clusters. In this work, we move away from the concept of clusters and allow for non-transitive similarity relations.

3 Representation Functions and Desirable Axioms

In this section, we formally introduce representation functions and propose a set of axioms that we consider essential for generalizing the well-understood case of discrete metrics. Notations are introduced as they appear, a summary is however provided in Appendix A.

Consider a metric space (E, d), that is a set E equipped with a notion of distance in the form of an operator $d : E \times E \mapsto \mathbb{R}_{\geq 0}$ satisfying *separability*, *symmetry* and *triangular inequality*. We now formally define the object of interest of this work, called *representation functions of* (E, d).

Definition 1 (Representation functions of (E, d))**.** *A representation function of (E, d) is a function f that maps finite sets of E to probability distributions over their elements, i.e.,*

$$f : S \in \mathcal{P}(E) \mapsto p_S \in \Delta(S),$$

where $\mathcal{P}(E)$ denotes the set containing all finite subsets of E (outside the empty set), and $\Delta(S) = \left\{ p_S : S \mapsto [0, 1] \mid \sum_{x \in S} p_S(x) = 1 \right\}$ denotes the simplex over the elements of S. We moreover refer to the probability distribution $f(S) : S \mapsto [0, 1]$ as the representation of S.

Note that this definition encompasses the *uniform distribution* as a particular case of representation function. Indeed, consider the discrete metric space (E, ρ), where $\rho(x, y)$ is equal to one if $x \neq y$ and zero otherwise. Then the maximum entropy principle compels us to use the *uniform representation function* $\mathcal{U} : S \in \mathcal{P}(E) \mapsto \mathbb{1}_S(\cdot)/|S| \in \Delta(S)$.

Drawing inspiration from the properties of this particular representation function, we next introduce a few axioms that we argue are desirable for a general metric space (E, d) and representation function f thereof. The first desirable property that the uniform representation $\mathcal{U}$ verifies is rather simple: it ensures that all elements of a finite set are represented with

positive probability. This means that it never hurts to add new elements to a set as the support of the probability distribution given by the representation function only increases.

Axiom 1 (Positivity)**.** *Every element of a finite set is represented with positive probability, i.e., for all finite subset $S \in \mathcal{P}(E)$ and element x in S, we have $f(S)(x) > 0$.*

The second property of $\mathcal{U}$ that we would want to extend to a generic f is that of symmetry: when the distance is uninformative and some elements are isomorphic, they receive similar weights. In particular, if all elements of a finite subset S are equidistant, then $f(S)$ should be uniform over S.

Axiom 2 (Symmetry)**.** *Elements of a set that are symmetric with respect to the metric are equally represented, i.e., for all finite subset $S \in \mathcal{P}(E)$ and self-isometry $\sigma_S : S \mapsto S$, it holds for all $x \in S$ that $f(S)(x) = f(S)(\sigma_S(x))$.*

Importantly, we consider in the above definition that the permutation σ_S preserves the distance on S, but need not be extendable to a full isometry on E. Moreover, determining the automorphism group of a set S is an instance of the *graph automorphism problem*, which is known to be solvable in quasi-polynomial time [Helfgott *et al.*, 2017], but is neither known to be in P nor to be NP-complete. Luckily, two symmetric elements x and $\sigma_S(x)$ possess the same multi-set of distances $\{\{d(x,y)\}\}_{y\in S}$ and we only need to make sure that similar multi-sets lead to similar representations.

Finally, the third property that $\mathcal{U}$ trivially satisfies is that of continuity, since the topology induced by the metric ρ is the discrete one. In order to define a notion of continuity for general metric spaces (E,d), we let $d(y,X) = \min_{x\in X} d(x,y)$ denote the distance from an individual element y in E to a finite subset X in $\mathcal{P}(E)$. We then equip the domain $\mathcal{P}(E)$ of a representation function f with the natural metric associated with d, that is the *Hausdorff distance* $d_H(X,Y) = \max\big\{\max_{x\in X} d(x,Y), \max_{y\in Y} d(X,y)\big\}$ where X, Y are two finite subsets of E. Finally, we denote by π_X the canonical projection on a finite subset X, that is the map $\pi_X : y \in E \mapsto \operatorname{argmin}_{x\in X} d(x,y)$, and let $\underline{d}(X) = \min_{x\neq x'\in X} d(x,x')$ be the minimal distance between two distinct elements of X. The following result characterizes neighborhoods in the metric space $(\mathcal{P}(E), d_H)$.

Lemma 1. *Let X be a finite subset of E and let δ satisfy $\underline{d}(X)/2 > \delta > 0$. A finite subset $Y \in \mathcal{P}(E)$ is at distance $d_H(X,Y) \leq \delta$ if and only if the canonical projection $\pi_X|_Y$ is the unique surjective map $\pi : Y \mapsto X$ such that $\max_{y\in Y} d(y,\pi(y)) \leq \delta$.*

The proof of Lemma 1 is included in Appendix B. Interestingly, the surjectivity means that, in small enough neighborhoods of a fixed finite subset X, there are only sets of cardinality greater or equal to that of X. Moreover, the projection $\pi = \pi_X|_Y$ offers a neat way to identify which elements of a neighboring set Y correspond to some element x of X: we then refer to each $\pi^{-1}(x) = \{y \in Y \mid \pi(y) = x\}$ as a *class of clones* since all y, y' in $\pi^{-1}(x)$ are at distance at most 2δ by the triangle inequality. Summing representations over each class of clones we are then able to collapse some of the dimensions of the codomain $\Delta(Y)$ and identify it with $\Delta(X)$. Note that such an homeomorphism was apriori not clearly identifiable since the two simplices would be of different dimensions whenever $|Y| > |X|$. With this intuition, we then define *class continuity* as follows.

Axiom 3 (Class Continuity)**.** *Representation is class-wise continuous, i.e., for a finite subset $X \in \mathcal{P}(E)$ and $\epsilon > 0$, there exists $\delta > 0$ such that, for each finite subset $Y \in \mathcal{P}(E)$ satisfying $d_H(X,Y) \leq \delta$, we have $\max_{x\in X} \big|f(X)(x) - \sum_{y\in\pi^{-1}(x)} f(Y)(y)\big| \leq \epsilon$, where $\pi = \pi_X|_Y$ denotes the canonical projection on X.*

We show in Appendix B that this formulation corresponds to the classical definition of continuity between two carefully chosen metric spaces. Note moreover that Axiom 3 ensures a form of *cloneproofness*, i.e., robustness of representation under the addition of clones. Intuitively, when a set Y contains many δ-clones, we can find a set X of smaller size in its δ-Hausdorff neighborhood: Axiom 3 then ensures that X and Y have "similar representations" when summing probabilities locally over the redundancies in Y.

However, we argue that this definition of continuity, although intuitive, is perhaps not desirable from an application standpoint. Indeed, we show hereafter that a representation function f satisfying Axioms 2 and 3 may give very different individual representations to points in nearby sets.

Example 1 (Diverging individual representations.)**.** *Let f be a representation function on the three-dimensional Euclidean space $(\mathbb{R}^3, d_2)$ satisfying both Axioms 2 and 3, and define the parametric family $S_{\alpha,\beta,\gamma} = \{o, u_\alpha, v^+_{\beta,\gamma}, v^-_{\beta,\gamma}\}$, where $o = (0,0,0)$, $u_\alpha = (1,\alpha,0)$, $v^+_{\beta,\gamma} = (1,-\beta,\gamma)$ and $v^-_{\beta,\gamma} = v^+_{\beta,-\gamma}$. Figure 2 summarizes our construction.*

On one hand, consider the set $S_{\alpha,\alpha,\gamma}$, where $\alpha > 0$ is fixed and $\gamma > 0$ is much smaller than α. Since $S_{\alpha,\alpha,\gamma}$ converges to the set $S_\alpha = \{o, u_\alpha, u_{-\alpha}\}$ in the Hausdorff sense when γ goes to zero, Axiom 3 implies $\lim_{\gamma\to 0} f(S_{\alpha,\alpha,\gamma})(u_\alpha) = f(S_\alpha)(u_\alpha)$. Moreover, S_α converges in turn to the symmetric set $S_0 = \{o, u_0\}$, and Axioms 2 and 3 together imply that $\lim_{\alpha\to 0} f(S_\alpha)(u_\alpha) = f(S_0)(u_0)/2 = 1/4$. Combining these two results, we get

$$\lim_{\alpha\to 0}\lim_{\gamma\to 0} f(S_{\alpha,\alpha,\gamma})(u_\alpha) = 1/4.$$

On the other hand, consider the set $S_{\alpha,\alpha/2,\sqrt{3}\alpha/2}$. Note that the points u_α, $v^+_{\alpha/2,\sqrt{3}\alpha/2}$ and $v^-_{\alpha/2,\sqrt{3}\alpha/2}$ form an equilateral triangle centered in u_0 and orthogonal to the origin, hence by Axiom 2 they must receive similar representation. As $S_{\alpha,\alpha/2,\sqrt{3}\alpha/2}$ also converges to the symmetric S_0 when α goes to zero, Axioms 2 and 3 finally imply

$$\lim_{\alpha\to 0} f\big(S_{\alpha,\alpha/2,\sqrt{3}\alpha/2}\big)(u_\alpha) = 1/6.$$

We hence constructed two arbitrarily close sets of the same cardinality whose individual representations differ.

The consequences of this interaction between Axioms 2 and 3 are perhaps better understood from a *fairness* perspective. Imagine for example an online lottery where participants can register for free by providing personal parameters that are difficult to forge, i.e., whose values may only be be slightly

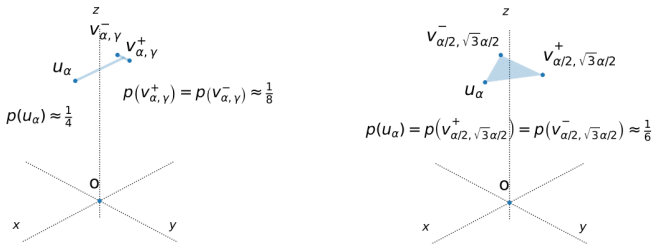

(a) $S_{\alpha,\alpha,\gamma}$ for $1 \gg \alpha \gg \gamma > 0$.

(b) $S_{\alpha,\alpha/2,\sqrt{3}\alpha/2}$ for $1 \gg \alpha > 0$

Figure 2: Visualization of the divergence of individual representations under Axioms 2 and 3. The edges in 2a highlight the symmetries of $S_{\alpha,\alpha,\gamma}$ in the limit $\gamma \to 0$; the equilateral triangle in 2b displays the symmetry of $S_{\alpha,\alpha/2,\sqrt{3}\alpha/2}$.

modified. In this scenario, *cloneproofness* means ensuring that a malevolent attacker cannot gain from registering to the lottery many times, i.e., keeping the lottery fair for players without the capacity or the envy to forge many fake identities. However, it may still happen that two individuals, e.g., two siblings, have similar personal parameters and cannot be distinguished from a behavioral perspective from two malevolent clones. In their misfortune, the two siblings would at least hope that they share the price equitably: the above example shows that using a representation function satisfying Axioms 2 and 3 would not ensure such *clone fairness*, a property formalized in the following axiom. This can be thought of as a continuous extension of Axiom 2 since perfect clones at distance precisely zero would always be isomorphic with one another, and thus have equal representations.

Axiom 4 (Uniform Clone Fairness)**.** *Representation is fair among δ-clones, i.e., for all $\epsilon > 0$, there exists $\delta > 0$ such that, for all finite subset $S \in \mathcal{P}(E)$ and x, y in S satisfying $d(x, y) \leq \delta$, it holds that $|f(S)(x) - f(S)(y)| \leq \epsilon$.*

Given this incompatibility, we propose to replace Axiom 3 with two slightly weaker axioms that are however compatible with the clone fairness in Axiom 4. Because Axiom 3 considers sets of different cardinality, it collapses high-dimensional simplices into smaller ones, and only looks at the continuity of representation between different classes of clones. In so doing, it loses track of individual representation which end up locally diverging. A simple solution to this issue is then to restrict the requirements of Axiom 3 to sets of similar cardinality. Note that any surjective map $\pi : Y \mapsto X$ becomes bijective in such a case, and this applies in particular to $\pi_X|_Y$ when δ is small enough in Lemma 1.

Axiom 5 (Individual Continuity)**.** *Representation is elementwise continuous, i.e., for a finite subset $X \in \mathcal{P}(E)$ and $\epsilon > 0$, there exists $\delta > 0$ such that, for all subset $Y \in \mathcal{P}(E)$ with $|Y| = |X|$ and $d_H(X, Y) \leq \delta$, we have $\max_{x \in X} |f(X)(x) - f(Y)(\pi^{-1}(x))| \leq \epsilon$, where $\pi = \pi_X|_Y$ denotes the canonical projection on X.*

Note that Axiom 5, unlike Axiom 3, does not address the addition of clones. To account for this, we introduce a second weakening of Axiom 3, essentially requiring continuity of representation everywhere except in the vicinity of the newly added clone. This relaxation allows for greater flexibility in how the mass is redistributed locally.

Axiom 6 (α-Locality under Addition of Clones)**.** *The addition of a clone only changes the representation of points in the α-neighborhood of the clone, i.e., for a finite subset $S \in \mathcal{P}(E)$ and $\epsilon > 0$, there exists $\delta > 0$ such that for each element $x \in S$ and δ-clone x' satisfying $d(x, x') \leq \delta$, we have for all $z \in S$ such that $d(x, z) \geq \alpha$ that $|f(S)(z) - f(S \cup \{x'\})(z)| \leq \epsilon$.*

Though similar in formulation, note that Axiom 4 cannot be derived from Axiom 5 by simply plugging in $Y = X$, since the restricted projection $\pi_X|_X$ is always the identity. Moreover, Axioms 4 and 6 provide orthogonal restrictions in the presence of clones: the former dictates how to shift weights around the recently introduced clone, while the latter ensures weights do not change away from it. We further discuss the relationship between these axioms in Appendix B.

We finally denote by $\mathcal{R}_\alpha(E, d)$ the set of representation functions on (E, d) satisfying Axioms 1, 2, 4, 5 and 6 with parameter $\alpha > 0$. The burning question is now this: does the set $\mathcal{R}_\alpha(E, d)$ contain any elements at all?

4 Local Voting Approach

To construct representation functions that satisfy our axioms, the first step is to identify invariant objects under clones' addition: we argue that the open balls of the topology, that is the $B_r(x) := \{y \in E \mid d(x, y) < r\}$ for some element $x \in E$ and radius $r > 0$, are natural invariants for our problem. Indeed, they are stable under the addition of clones, in the sense that for some *δ-clones* x, y in E satisfying $d(x, y) \leq \delta$, the triangle inequality ensures that $B_r(x) \subseteq B_r(x) \cup B_r(y) \subseteq B_{r+\delta}(x)$. If we then equip our space with a measure μ defined on the open balls of the space[1] and associate with each finite subset $X \subseteq E$ its neighborhood $B_r(X) := \bigcup_{x \in X} B_r(x)$, we obtain a map invariant under clone addition. Indeed, for each neighboring finite set $Y \subseteq B_\delta(X)$ with $r > \delta > 0$, we have $\mu\big(B_{r-\delta}(X)\big) \leq \mu\big(B_r(Y)\big) \leq \mu\big(B_{r+\delta}(X)\big)$ and the map $X \in \mathcal{P}(E) \mapsto \mu\big(B_r(X)\big)$ is continuous with respect to the Hausdorff distance.[2]

Note however that further requirements are needed to satisfy the symmetry in Axiom 2, essentially regarding the homogeneity and the isotropy of the underlying measure space. For this reason, we focus on Euclidean spaces $(\mathbb{R}^n, d_2)$ for the remainder of the section, where $d_2(x, y) = \sum_{i=1}^n (x_i - y_i)^2$ for all $x = (x_i)_{1 \leq i \leq n}$ and $y = (y_i)_{1 \leq i \leq n}$ in $\mathbb{R}^n$. We will discuss in Section 5 how to adapt our approach to more general metric spaces.

Based on the above invariant, we construct a representation function as a local voting scheme. For a fixed $r > 0$ and finite subset $S \subseteq \mathbb{R}^n$, we consider each element of $B_r(S)$ as a voter that approves only of the candidates in S close to him, and as such spreads his voting power equally among them. Formally, we define the grade that each voter z in $B_r(S)$ attributes to a

[1] I.e., on the Borel σ-algebra.

[2] At least when μ is locally finite.

candidate x in S as follows

$$g_{r,S,x}(z) = \frac{\mathbb{1}_{B_r(x)}(z)}{\sum_{y \in S} \mathbb{1}_{B_r(y)}(z)}.$$

We then aggregate the ballots with the Lebesgue measure μ and finally define the representation function g_r, for each finite subset $S \subseteq \mathbb{R}^n$, as follows

$$g_r(S) : x \in S \mapsto \int_{B_r(S)} \frac{g_{r,S,x}(z)}{\mu\big(B_r(S)\big)}\, d\mu(z).$$

As illustrated in Figure 3, the representation function g_r computes a weighted average of the inverse depth of each cell, with the depth defined as the number of intersecting balls forming the cell and the weights based on the cell's size.

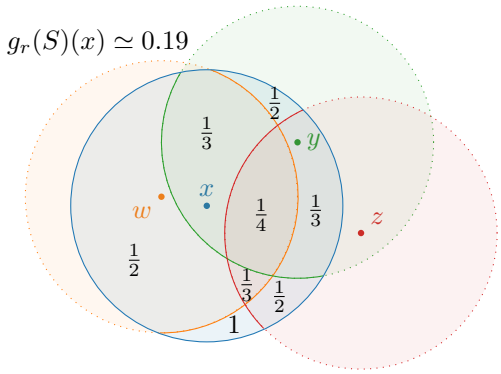

Figure 3: Computation of $g_r(S)(x)$ in the two-dimensional Euclidean space $(\mathbb{R}^2, d_2)$, where the set $S = \{w, x, y, z\}$ contains four elements. A cell $A_r(U)$ is uniquely defined by the subset $U \subseteq S$ as the possibly empty intersection of the balls around each element in U and the complement of the balls of each element of S absent from U (c.f. Appendix C). For each subset U containing x, the grading function $g_{r,S,x}$ is constant on the cell $A_r(U)$ and equal to the inverse depth of the cell, i.e., $g_{r,S,x}(z) = 1/U$ for all z in $A_r(U)$. The weight of x in S is then equal to the weighted average of $g_{r,S,x}$ on the ball centered in x, where the weight of each cell corresponds to its area normalized by the total area of the balls' union. We estimated the value $g_r(S)(x) \simeq 0.19$ via Monte Carlo sampling.

We next show that this class of representation functions satisfies the desirable axioms we introduced in Section 3.

Theorem 1. *For $r > 0$, the representation function g_r is well-defined and belongs in $\mathcal{R}_{2r}(\mathbb{R}^n, d_2)$.*

The detailed proof of Theorem 1 is included in Appendix C, and we provide here a sketch of the proof. As depicted in Figure 3, the weight $g_r(S)(x)$ is in fact a weighted average of positive elements, hence it is positive and Axiom 1 trivially holds. Showing that g_r is symmetric (Axiom 2) is also relatively straightforward after observing the following two properties of Euclidean spaces: first, the fact that one can uplift any self-isometry σ_S on a finite subset S to an isometry on the entire space $\mathbb{R}^n$, c.f. Appendix D; second, the fact that the Lebesgue measure is invariant under *translations*, *rotations* and *reflections*, which generate the group of Euclidean isometries [Gallian, 2020]. The most challenging aspect of the proof is verifying that g_r satisfies Axioms 4, 5 and 6. While these proofs are technically intricate, they fortunately follow a similar structure. We illustrate our approach by focusing on the simpler case of Axiom 4 below. In order to bound the difference of representation function $|g_r(S)(x) - g_r(S)(y)|$, we first show that the grading functions $g_{r,S,x}$ and $g_{r,S,y}$ are equal outside of a thin spherical shell parametrized by δ, the distance between the neighboring elements x and y (c.f. Figure 4a). This allows us to obtain a difference of Lebesgue measure $\mu(B_r(x)) - \mu(B_{r-\delta}(x))$, which we then bound in terms of δ by taking the limit of this difference as δ approaches zero (c.f. Figure 4b). The formalization of this argument relies on tools from *geometric measure theory*, particularly the $n-1$-dimensional Minkowski content. These arguments are illustrated in Figure 4.

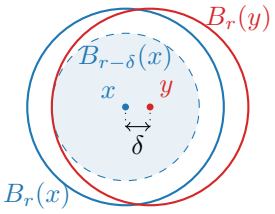

(a) The ball $B_{r-\delta}(x)$ belongs in the intersection of $B_r(x)$ and $B_r(y)$, hence x and y receive the same grade from every voter on $B_{r-\delta}(x)$.

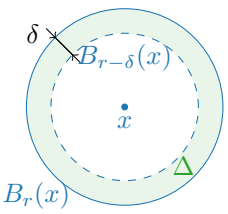

(b) The n^{th}-dimensional volume of the green set $\mu(\Delta)$ can be approximated as δ times S_r^{n-1}, the $n-1^{\text{th}}$-dimensional surface of a ball of radius r.

Figure 4: Key steps in demonstrating that g_r satisfies Axiom 4.

Since $\mathcal{R}_\alpha(E, d)$ is monotonically increasing in positive α, Theorem 1 actually ensures that the whole collection $\{g_r\}_{\alpha/2 \geq r > 0}$ belongs in $\mathcal{R}_\alpha(\mathbb{R}^n, d_2)$. Moreover, it is relatively straightforward to see that $\mathcal{R}_\alpha(\mathbb{R}^n, d_2)$ is a convex set and, as such, contains all finite convex combinations of $\{g_r\}_{\alpha/2 \geq r > 0}$. Since the representation functions g_r are well-behaved, we generalize this result as follows.

Theorem 2. *Let ν be a probability density function over $[0, \alpha/2]$. Then the representation function $f_\nu : S \in \mathcal{P}\big(\mathbb{R}^n\big) \mapsto \int_0^{\alpha/2} \nu(r) g_r(S)\, dr$ belongs in $\mathcal{R}_\alpha(\mathbb{R}^n, d_2)$.*

The detailed proof of Theorem 2, provided in Appendix C, relies on inequalities derived for the proof of Theorem 1.

5 Discussion

We gather in this section different remarks on our results as well as possible extensions of our work.

Extension to Perfect Clones. The framework we considered until now only allows for δ-clones with $\delta > 0$, but not perfect clones, i.e., with $\delta = 0$. The appropriate analytical tool to handle this is to consider a *pseudo-metric space* (E, d) instead of a metric one, where the pseudo-metric d verifies *non-negativity*, *symmetry*, *triangle inequality* like a full-fledged metric, but only verifies *identity* instead of *separability*. This exactly means that two different elements $x \neq y$ in E may be perfect clones, i.e., $d(x, y) = 0$.

Most of the axioms we introduced directly extend to a pseudo-metric space (E, d); however, the surjectivity in

Lemma 1 breaks, and with it Axioms 3 and 5. To circumvent this issue, we rely on the equivalence relationship induced by the vanishing of the pseudo-metric, i.e., we define the *metric identification* as $x \sim y$ if and only if $d(x, y) = 0$, and write $[x]$ for the equivalence class of x in E. Letting $E^* = E/\sim$ denote the quotient space of E by the equivalence relation $\sim$, and defining $d^* : ([x], [y]) \in E^* \times E^* \mapsto d(x, y)$, we then refer to (E^*, d^*) as the metric space induced by the pseudo-metric space (E, d). In order to adapt Axiom 5 to a pseudo-metric space, we simply ask that its requirements be verified after we collapse all perfect clones. In other words, there must exists a representation function f^* of (E^*, d^*) verifying Axiom 5 such that $f^*(X/\sim)([x]) = \sum_{x' \in [x] \cap X} f(X)(x')$ holds for all choices of X and $x \in X$. We refer the interested reader to Appendix E for more details.

Note that a similar extension is possible for Axiom 3. As argued in Section 3 however, representation functions satisfying both this extension and Axiom 2 would treat δ-clones and perfect clones differently since the latter are always isomorphic but the former not, thus violating Axiom 4.

Extension beyond Euclidean spaces. While the solution proposed in Section 4 is restricted to Euclidean spaces, similar ideas could be applied in more general metric spaces. Using a Radon measure μ, one could define the representation functions g_r in full generality and show similarly as in Theorem 1 that Axioms 1, 4, 5 and 6 hold. The real challenge however is to satisfy Axiom 2.

Indeed, our proof relies on two convenient properties of Euclidean spaces: first, the uplifting of self-isometry σ_S to the entire space; second the invariance of the Lebesgue measure under translations, rotations and reflections. What can be done without these properties? The first problem could be entirely shunned by arguing that only full-fledged isometries should be considered in the definition of Axiom 2. The second issue is however tougher to ward off. To extend invariance by translation beyond vector spaces, one should consider *uniformly distributed measures*, i.e., measures that give the same weight to all balls of the same radius. However, such measures turn out to be very rigid objects and are uniquely defined up to a multiplicative constant in most metric spaces.

Lemma 2 (From [Christensen, 1970])**.** *Let (E, d) be a locally compact metric space. There exists a Radon measure μ defined on the Borel σ-algebra of E that is uniformly distributed, i.e., it verifies $0 < \mu(B_r(x)) = \mu(B_r(y)) < \infty$ for all $r > 0$ and x, y in E. Moreover, this measure is unique up to a multiplicative constant if E is separable.*

As a particular example, this essentially implies that the Lebesgue measure is the only Borel measure invariant by translation on $\mathbb{R}^n$. This comes as an awful news since this indicates that our approach is doomed even in the simple case of $\mathbb{R}^n$ endowed with the L^1 distance $d_1(x, y) = \sum_{i=1}^n |x_i - y_i|$, as explained in Figure 5.

Such metric spaces thus require developing techniques different from the one introduced in this work. *Topologically independent* representation functions would provide an elegant solution to this issue, i.e., functions that do not rely on the topological properties of (E, d), but rather solely depend on the distance matrices associated with each finite set.

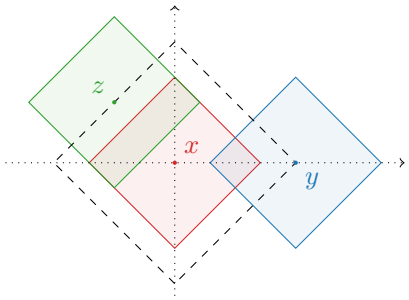

Figure 5: The representation function g_r does not satisfy Axiom 2 in $(\mathbb{R}^2, d_1)$. As illustrated by the dashed L^1 ball centered in x, points y and z are indeed at the same distance of x, thus belong in a common isometry class in $S = \{x, y, z\}$ and should receive similar representation under Axiom 2. Note however that the Lebesgue measure, i.e., the area, of the intersection between the red and the green ball differs from that of the intersection between the red and the blue ball, hence $g_r(S)(y) \neq g_r(S)(z)$.

Axiom 7 (Topological Invariance)**.** *Representation only depends on the distance matrix associated with each finite set, i.e., there exists a family $(h_n)_{n \geq 1}$ with $h_n : \mathbb{R}^{n \times n} \mapsto \Delta(n)$ such that, for all $S \in \mathcal{P}(E)$ of cardinality $|S| = n$, we have $f(S) = h_n(M(S))$, where $M = (d(x, y))_{x,y \in S} \in \mathbb{R}^{n \times n}$ denotes the distance matrix associated to S and d, unique up to permutations.*

Identifying representation functions within $\mathcal{R}_\alpha(E, d)$ that adhere to Axiom 7 is a promising direction for future works.

Computability Our focus thus far has been on identifying representation functions with theoretically desirable properties. However, from a practical standpoint, such tools are of little utility if they cannot be computed efficiently. This concern is encapsulated in the following principle.

Axiom 8 (Exact Computability)**.** *The representation of any given subset is efficiently computable, i.e., for any subset $S \in \mathcal{P}(E)$, the probability distribution $f(S)$ can be exactly computed in time polynomially bounded by the cardinality $|S|$ of the subset, and the dimension n of the space if $E = \mathbb{R}^n$.*

It is worth noting that the representation functions introduced in Section 4 are unlikely to meet this criterion. For example, computing $g_r(S)(x)$ would a priori involve averaging $g_{r,S,x}$ over as many as $O\big(2^{|S|}\big)$ disjoint cells, making the approach computationally infeasible. Additionally, even the simpler task of calculating the volume $\mu\big(B_r(S)\big)$ of the union of Euclidean balls becomes increasingly challenging in higher dimensions (see [Cazals *et al.*, 2011] for the case $n = 3$). While hardness results for this specific problem are not readily available, related problems, such as computing the exact volume of the union of general axis-aligned boxes, are known to be #P-hard [Bringmann and Friedrich, 2010].

In practice, as often occurs, we may want to relax Axiom 8 and settle for approximate evaluations of $g_r(S)$. Monte-Carlo-based methods [Bringmann and Friedrich, 2010; Mitchell *et al.*, 2018] could potentially be adapted to our framework and enable us to efficiently estimate $g_r(S)$ within an ϵ factor with high probability, where $\epsilon \propto 1/\sqrt{k}$ typically decreases quadratically with the number of samples k.

References

[Achille *et al.*, 2019] Alessandro Achille, Michael Lam, Rahul Tewari, Avinash Ravichandran, and al. Task2Vec: Task Embedding for Meta-Learning, February 2019.

[Alvarez-Melis and Fusi, 2020] David Alvarez-Melis and Nicolò Fusi. Geometric Dataset Distances via Optimal Transport, February 2020.

[Babaioff *et al.*, 2016] Moshe Babaioff, Shahar Dobzinski, Sigal Oren, and Aviv Zohar. On Bitcoin and Red Balloons, June 2016.

[Balduzzi *et al.*, 2018] David Balduzzi, Karl Tuyls, Julien Perolat, and Thore Graepel. Re-evaluating Evaluation, October 2018.

[Banerjee *et al.*, 2023] Jineta Banerjee, Jaclyn N. Taroni, Robert J. Allaway, Deepashree Venkatesh Prasad, and al. Machine learning in rare disease. *Nature Methods*, 20(6), June 2023.

[Bringmann and Friedrich, 2010] Karl Bringmann and Tobias Friedrich. Approximating the volume of unions and intersections of high-dimensional geometric objects. *Computational Geometry*, 43(6), August 2010.

[Cao *et al.*, 2019] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning Imbalanced Datasets with Label-Distribution-Aware Margin Loss, October 2019.

[Carlini *et al.*, 2024] Nicholas Carlini, Matthew Jagielski, Christopher A. Choquette-Choo, Daniel Paleka, and al. Poisoning Web-Scale Training Datasets is Practical, May 2024.

[Cazals *et al.*, 2011] Frederic Cazals, Harshad Kanhere, and Sébastien Loriot. Computing the volume of a union of balls: A certified algorithm. *ACM Trans. Math. Softw.*, 38(1), December 2011.

[Chang *et al.*, 2018] Haw-Shiuan Chang, Erik Learned-Miller, and Andrew McCallum. Active Bias: Training More Accurate Neural Networks by Emphasizing High Variance Samples, January 2018.

[Chen *et al.*, 2022] Hongyin Chen, Xiaotie Deng, Ying Wang, Yue Wu, and Dengji Zhao. Sybil-Proof Diffusion Auction in Social Networks, November 2022.

[Christensen, 1970] Jens Peter Reus Christensen. On Some Measures Analogous to Haar Measure. *Mathematica Scandinavica*, 26, June 1970.

[Colombo *et al.*, 2022] Pierre Colombo, Nathan Noiry, Ekhine Irurozki, and Stephan Clemencon. What are the best systems? New perspectives on NLP Benchmarking, October 2022.

[Conitzer and Yokoo, 2010] Vincent Conitzer and Makoto Yokoo. Using Mechanism Design to Prevent False-Name Manipulations. *AI Magazine*, 31(4), December 2010.

[Cui *et al.*, 2019] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-Balanced Loss Based on Effective Number of Samples, January 2019.

[Dong *et al.*, 2017] Qi Dong, Shaogang Gong, and Xiatian Zhu. Class Rectification Hard Mining for Imbalanced Deep Learning, December 2017.

[Eyring *et al.*, 2024] Veronika Eyring, William D. Collins, Pierre Gentine, Elizabeth A. Barnes, and al. Pushing the frontiers in climate modelling and analysis with machine learning. *Nature Climate Change*, 14(9), September 2024.

[Federer, 1996] Herbert Federer. *Geometric Measure Theory*. Classics in Mathematics. Springer, 1996.

[Freund and Schapire, 1997] Yoav Freund and Robert E Schapire. A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting. *Journal of Computer and System Sciences*, 55(1), August 1997.

[Gallian, 2020] Joseph Gallian. *Contemporary Abstract Algebra*. CRC Press, 10 edition, December 2020.

[Gebru *et al.*, 2017] Timnit Gebru, Judy Hoffman, and Li Fei-Fei. Fine-Grained Recognition in the Wild: A Multi-task Domain Adaptation Approach. In *IEEE International Conference on Computer Vision*, October 2017.

[Ghojogh *et al.*, 2022] Benyamin Ghojogh, Ali Ghodsi, Fakhri Karray, and Mark Crowley. Spectral, Probabilistic, and Deep Metric Learning: Tutorial and Survey, January 2022.

[Gretton *et al.*, 2009] A. Gretton, A.J. Smola, J. Huang, Marcel Schmittfull, and al. Covariate Shift by Kernel Mean Matching. *Dataset Shift in Machine Learning*, January 2009.

[Gretton *et al.*, 2012] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *J. Mach. Learn. Res.*, 13, March 2012.

[Hedderich *et al.*, 2021] Michael A. Hedderich, Lukas Lange, Heike Adel, Jannik Strötgen, and Dietrich Klakow. A Survey on Recent Approaches for Natural Language Processing in Low-Resource Scenarios, April 2021.

[Helfgott *et al.*, 2017] Harald Andrés Helfgott, Jitendra Bajpai, and Daniele Dona. Graph isomorphisms in quasi-polynomial time, October 2017.

[Himmi *et al.*, 2023] Anas Himmi, Ekhine Irurozki, Nathan Noiry, Stephan Clemencon, and Pierre Colombo. Towards More Robust NLP System Evaluation: Handling Missing Scores in Benchmarks, May 2023.

[Horn and Johnson, 2012] Roger A. Horn and Charles R. Johnson. *Matrix analysis*. Cambridge University Press, 2nd edition, 2012.

[Härdle *et al.*, 2004] Wolfgang Härdle, Axel Werwatz, Marlene Müller, and Stefan Sperlich. *Nonparametric and Semiparametric Models*. Springer Series in Statistics. 2004.

[Jamal *et al.*, 2020] Muhammad Abdullah Jamal, Matthew Brown, Ming-Hsuan Yang, Liqiang Wang, and Boqing Gong. Rethinking Class-Balanced Methods for Long-Tailed Visual Recognition from a Domain Adaptation Perspective, March 2020.

[Jeantet *et al.*, 2020] Ian Jeantet, Zoltan Miklos, and David Gross-Amblard. Overlapping Hierarchical Clustering (OHC). In *Inteligent Data Analysis*, April 2020.

[Jiang *et al.*, 2015] Lu Jiang, Deyu Meng, Qian Zhao, Shiguang Shan, and Alexander G. Hauptmann. Self-paced curriculum learning. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, January 2015.

[Jumper *et al.*, 2021] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, and al. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), August 2021.

[Koh *et al.*, 2021] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, and al. WILDS: A Benchmark of in-the-Wild Distribution Shifts, July 2021.

[Liu *et al.*, 2022] Xinran Liu, Yikun Bai, Yuzhe Lu, Andrea Soltoggio, and Soheil Kolouri. Wasserstein Task Embedding for Measuring Task Similarities, August 2022.

[Mitchell *et al.*, 2018] Scott A. Mitchell, Muhammad A. Awad, Mohamed S. Ebeida, and Laura P. Swiler. Fast Approximate Union Volume in High Dimensions with Line Samples. Technical Report SAND–2018-8684, 1464880, Sandia National Laboratories, August 2018.

[Motie and Raahemi, 2024] Soroor Motie and Bijan Raahemi. Financial fraud detection using graph neural networks: A systematic review. *Expert Systems with Applications*, 240, April 2024.

[Murtagh and Contreras, 2017] Fionn Murtagh and Pedro Contreras. Algorithms for hierarchical clustering: an overview, II. *WIREs Data Mining and Knowledge Discovery*, 7(6), 2017.

[Nehama *et al.*, 2022] Ilan Nehama, Taiki Todo, and Makoto Yokoo. Manipulation-resistant false-name-proof facility location mechanisms for complex graphs. *Autonomous Agents and Multi-Agent Systems*, 36(1), January 2022.

[Peng *et al.*, 2020] Xingchao Peng, Yichen Li, and Kate Saenko. Domain2Vec: Domain Embedding for Unsupervised Domain Adaptation, July 2020.

[Peyrard *et al.*, 2021] Maxime Peyrard, Wei Zhao, Steffen Eger, and Robert West. Better than Average: Paired Evaluation of NLP Systems. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics*, 2021.

[Ran *et al.*, 2023] Xingcheng Ran, Yue Xi, Yonggang Lu, Xiangwen Wang, and Zhenyu Lu. Comprehensive survey on hierarchical clustering algorithms and the recent developments. *Artificial Intelligence Review*, 56(8), August 2023.

[Ren *et al.*, 2019] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to Reweight Examples for Robust Deep Learning, May 2019.

[Resnick and Sami, 2009] Paul Resnick and Rahul Sami. Sybilproof transitive trust protocols. In *Proceedings of the 10th ACM conference on Electronic commerce*, July 2009.

[Rofin *et al.*, 2023] Mark Rofin, Vladislav Mikhailov, Mikhail Florinsky, Andrey Kravchenko, and al. Vote'n'Rank: Revision of Benchmarking with Social Choice Theory. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, May 2023.

[Seuken and Parkes, 2014] Sven Seuken and David C. Parkes. Sybil-proof accounting mechanisms with transitive trust. In *Proceedings of The 15th International Conference on Autonomous Agents and Multi-Agent Systems*, May 2014.

[Shu *et al.*, 2019] Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, and al. Meta-Weight-Net: Learning an Explicit Mapping For Sample Weighting, September 2019.

[Stannat *et al.*, 2021] Alexander Stannat, Can Umut Ileri, Dion Gijswijt, and Johan Pouwelse. Achieving Sybil-Proofness in Distributed Work Systems. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, May 2021.

[Tatiana and Valentin, 2021] Shavrina Tatiana and Malykh Valentin. How not to Lie with a Benchmark: Rearranging NLP Leaderboards, December 2021.

[Todo *et al.*, 2009] Taiki Todo, Atsushi Iwasaki, Makoto Yokoo, and Yuko Sakurai. Characterizing false-name-proof allocation rules in combinatorial auctions. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems*, May 2009.

[Torralba and Efros, 2011] Antonio Torralba and Alexei A. Efros. Unbiased look at dataset bias. In *CVPR*, June 2011.

[Varshney *et al.*, 2022] Ayush K. Varshney, Pranab K. Muhuri, and Q. M. Danish Lohani. PIFHC: The Probabilistic Intuitionistic Fuzzy Hierarchical Clustering Algorithm. *Applied Soft Computing*, 120, May 2022.

[Wang and Deng, 2018] Mei Wang and Weihong Deng. Deep Visual Domain Adaptation: A Survey. *Neurocomputing*, 312, October 2018. arXiv:1802.03601.

[Wang *et al.*, 2020] Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, and al. SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems, February 2020.

[Wang *et al.*, 2024] Wenxi Wang, Yang Hu, Mohit Tiwari, Sarfraz Khurshid, and al. NeuroBack: Improving CDCL SAT Solving using Graph Neural Networks, May 2024.

[Y *et al.*, 2019] Geeta Dharani. Y, Nimisha G Nair, Pallavi Satpathy, and Jabez Christopher. Covariate Shift: A Review and Analysis on Classifiers. In *Global Conference for Advancement in Technology*, October 2019.

[Zhang and Hardt, 2024] Guanhua Zhang and Moritz Hardt. Inherent Trade-Offs between Diversity and Stability in Multi-Task Benchmarks, May 2024.

[Zhang *et al.*, 2020] Yao Zhang, Xiuzhen Zhang, and Dengji Zhao. Sybil-proof Answer Querying Mechanism, May 2020.

A Preliminaries & Notations

In this section, we provide an overview of the notations used throughout the paper and introduce the key tools necessary for our demonstrations.

Metric spaces and Hausdorff norm. Let E be a set and let $d : E \times E \mapsto \mathbb{R}_{\geq 0}$ be a metric on E, that is an operator satisfying for all $x, y, z \in E$

1. (*Non-negativity*) $d(x, y) \geq 0$;
2. (*Symmetry*) $d(x, y) = d(y, x)$;
3. (*Triangle inequality*) $d(x, z) \leq d(x, y) + d(y, z)$;
4. (*Separability*) $d(x, y) = 0 \iff x = y$.

For $n \in \mathbb{N}$, let $\mathcal{P}_n(E) := \{S \subseteq E \mid n = |S|\}$ denote the powerset of subsets of E of cardinality n; we further denote by $\mathcal{P}(E) := \bigcup_{n\geq 1} \mathcal{P}_n(E)$ the set of finite subsets of E. In particular, $\mathcal{P}(E)$ does not contain the empty-set.

We equip $\mathcal{P}(E)$ with the *Hausdorff distance* d_H, defined for two finite subsets X, Y in $\mathcal{P}(E)$ as

$$d_H(X, Y) := \max\big\{ \max_{x\in X} d(x, Y), \max_{b\in Y} d(X, y)\big\}.$$

A metric space is finally defined as an ordered pair where the first element is a set and the second is a metric on this set; (E, d) and $(\mathcal{P}(E), d_H)$ constitute the two most prominent examples of metric spaces in this work.

Isometries and self-isometries. In the metric space (E, d), an isometry is defined as a map $\sigma : E \mapsto E$ that preserves distances, i.e., that verifies for all x, y in E,

$$d(\sigma(x), \sigma(y)) = d(x, y).$$

For $X \in \mathcal{P}(E)$ a finite subset of E, we moreover refer to an isometry in the subspace induced by X as a *self-isometry on* X, i.e., a permutation $\sigma_X : X \mapsto X$ such that $d(\sigma_X(x), \sigma_X(y)) = d(x, y)$ for all x, y in X.

While restricting an isometry σ to a subset $X \subseteq E$ gives a self-isometry on X, note that the converse may not hold in general. We however show in Appendix D that it holds in the particular case of Euclidean spaces and that a self-isometry σ_X on $X \subseteq E$ may then be extended to a full-fledged isometry σ on the entire space E.

Projection. For X a finite subset of E, we let $\underline{d}(X)$ denote the *inner diameter of* X, defined as

$$\underline{d}(X) := \min_{x\neq x'\in X} d(x, x').$$

Moreover, we define the *canonical projection on* X as the operator $\pi_X : E \mapsto \mathcal{P}(X)$ verifying $\pi_X(y) = \operatorname{argmin}_{x\in X} d(x, y)$ for all y in E.

In particular, note that the codomain of π_X is $\mathcal{P}(X)$, since the minimum may be achieved in multiple points simultaneously. Lemma 1 shows however that this does not happen when the domain of π_X is restricted to finite sets Y close enough to X; we denote the restricted operator as $\pi_X|_Y$. In such cases, we may associate each singleton $\{x\}$ in $\mathcal{P}(X)$ to the corresponding element $x \in X$ and recover the formulation in the main body.

Topology. In a metric space (E, d), we denote by $B_r(x)$ the open ball of radius $r > 0$ centered in $x \in E$, that is the set $B_r(x) := \{y \in E \mid d(x, y) < r\}$. When X is a set of points in E, we extend this definition and write $B_r(X) := \bigcup_{x\in X} B_r(x)$ for the union of the open balls of radius $r > 0$ centered at each element of X.

A set $X \subseteq E$ is then said to be *open* if it contains a ball of positive size centered in each of its elements, i.e., for all $x \in X$, there exists $r > 0$ such that $B_r(x) \subseteq X$. On the contrary, we say that a set X is *closed* when its complement $X^c = E \setminus X$ is open. We next define the *closure* $\overline{X}$ of a set X as the smallest closed set that contains X, and its *interior* $\mathring{X}$ as the largest open set contained within X.

Hausdorff measure. A *σ-algebra on E* is a non-empty collection of subsets of E closed under complement, countable union and countable intersections. In particular, the *Borel σ-algebra* Σ is the smallest σ-algebra by set inclusion containing all open sets of E.

We refer to the ordered pair (E, Σ) as a *measurable space*, and define a *measure* as a function $\mu : \Sigma \mapsto \mathbb{R} \cup \{\pm\infty\}$ that verifies

1. *Non-negativity:* $\mu(X) \geq 0$ for all $X \in \Sigma$;
2. $\mu(\emptyset) = 0$;
3. *Countable additivity:* $\mu\big(\bigcup_{k\in\mathbb{N}} X_k\big) = \sum_{k\in\mathbb{N}} \mu(X_k)$ for all countable collection $\{X_k\}_{k\in\mathbb{N}}$ of pairwise disjoint sets in Σ.

The *Hausdorff measure* is particularly important, we recall its definition hereafter. For U a subset of E, we first define the *diameter of* U as

$$\operatorname{diam}(U) = \sup\{d(x, y) \mid x, y, \in U\}.$$

We moreover adopt the convention $\operatorname{diam}(\emptyset) = 0$. For two positive real number δ and m, as well as a subset $X \in \Sigma$, we further write

$$\mathcal{H}^m_\delta(X) = \inf\left\{ \sum_{k\in\mathbb{N}} \operatorname{diam}(U_k)^m \mid X \subseteq \bigcup_{k\in\mathbb{N}} U_k, \operatorname{diam}(U_k) < \delta \right\},$$

where the infimum is taken over countable collections $\{U_k\}_{k\in\mathbb{N}}$ that cover X with sets of diameter smaller than δ. The *m-dimensional Hausdorff measure* is then finally defined as

$$\mathcal{H}^m(X) = 2^{-m}\alpha_m \cdot \lim_{\delta\to 0} \mathcal{H}^m_\delta(X),$$

where $\alpha(m) = \frac{\pi^{m/2}}{\Gamma(m/2+1)}$ denotes the volume of the unit m-ball, and Γ represents Euler's gamma function.

When m is an integer, the scaling $2^{-m}\alpha_m$ ensures that the m-dimensional Hausdorff measure coincides with the classical Lebesgue measure on the Borel sets of an m-dimensional Euclidean space.

For an Euclidean space $(E, d) = (\mathbb{R}^n, d_2)$, we can therefore write the $n-1$-dimensional surface of a ball of radius $r > 0$ as

$$S^{n-1}_r = \mathcal{H}^{n-1}\big(B_r(x)\big),$$

where x may be any point in E since the Lebesgue measure, hence also the Hausdorff measure, are invariant by translation. Denoting by V_r^n the n-dimensional volume of a ball of radius $r > 0$, we moreover have $V_r^n = \alpha(n) r^n$ as well as the relation $S_r^{n-1} = \frac{n}{r} V_r^n$.

Geometric measure theory. For two subsets X, Y of $E = \mathbb{R}^n$, we define the *Minkowski sum* $X + Y$ as follows, i.e., $X + Y = \{x + y \mid x \in X, y \in Y\}$.

Let $o = (0, \dots, 0)$ denote the origin and let m be an integer verifying $0 \leq m \leq n$. We next define the *m-dimensional Minkowski content of X* as

$$\mathcal{M}^{*m}(X) = \limsup_{\delta \to 0} \frac{\mu\big(X + B_\delta(o)\big)}{\alpha(n-m)\delta^{n-m}},$$

where μ denotes the m-dimensional Lebesgue measure, and $X + B_\delta(o))$ is a Minkowski sum. The *m-dimensional lower Minkowski content* $\mathcal{M}_*^m(X)$ is similarly defined by replacing the $\sup$ by an $\inf$ in the definition of $\mathcal{M}^{*m}(X)$.

When the upper and lower m-dimensional Minkowski contents are equal, we call their common value $\mathcal{M}^m(X)$ the *m-dimensional Minkowski content of X*.

In particular for $m = n - 1$, we have $\alpha(1) = 2$ and we get

$$\mathcal{M}^{n-1}(X) = \lim_{\delta \to 0} \frac{\mu\big(X + B_\delta(o)\big)}{2\delta}.$$

A set X is said to be *m- rectifiable* if and only if there exists a Lipschitz function mapping some bounded subset of $\mathbb{R}^m$ onto X.

With these definitions in place, we can now present a fundamental result in geometric measure theory that establishes a connection between the Minkowski content of well-behaved sets and their Hausdorff measure.

Theorem 3 (From [Federer, 1996], Thm 3.2.39)**.** *If X is a closed m-rectifiable set of $\mathbb{R}^n$, then $\mathcal{M}^m(X) = \mathcal{H}^m(X)$.*

B Definition of Continuity in Terms of Metric and Discussion of the Axioms

In this section, we prove the Lemmas discussed in Section 3, and provide a more thorough discussion of the relationships between the different Axioms.

We open this section with a short proof of Lemma 1 characterizing neighborhoods in $(\mathcal{P}(E), d_H)$.

Lemma 1. *Let X be a finite subset of E and δ satisfy $\underline{d}(X)/2 > \delta > 0$. A finite subset $Y \in \mathcal{P}(E)$ is at distance $d_H(X, Y) \leq \delta$ if and only if the canonical projection $\pi_X|_Y$ is the unique surjective map $\pi : Y \mapsto X$ such that $\max_{y \in Y} d(y, \pi(y)) \leq \delta$ holds.*

Proof. Let X be a finite subset in $\mathcal{P}(E)$ and let δ satisfy $\max_{x,x' \in X} d(x, x')/2 > \delta > 0$.

We first show the converse $\Longleftarrow$. Consider a finite subset Y such that $\pi = \pi_X|_Y$ is surjective and verifies $d(y, \pi(y)) \leq \delta$. We directly have $d(X, y) \leq d(\pi(y), y) \leq \delta$ for all $y \in Y$. Moreover for $x \in X$, the set $\pi^{-1}(x)$ is non-empty by surjectivity of π, and we similarly get $d(x, Y) \leq \min_{y \in \pi^{-1}(x)} d(\pi(y), y) \leq \delta$. Taking the maximum over Y and X respectively, we finally establish that $\max_{y \in Y} d(X, y) \leq \delta$ and $\max_{x \in X} d(x, Y) \leq \delta$, which, combined with the definition of the Hausdorff norm, gives the desired result $d_H(X, Y) \leq \delta$.

We now turn to the direction $\Longrightarrow$. Let Y be a finite subset in $\mathcal{P}(E)$ that verifies $d_H(X, Y) \leq \delta$. We first show that $\pi = \pi_X|_Y$ is well-defined and surjective. Indeed, consider x in X and let y be an element of Y closest to x. We then directly get $d(x, y) = d(x, Y) \leq d_H(X, Y) \leq \delta < \max_{z,z' \in X} d(z, z')/2$. Moreover, let x' be an element of X different from x: by triangle inequality, we get $d(y, x') \geq d(x, x') - d(x, y) \geq \max_{z,z' \in X} d(z, z') - \delta > \delta$. This implies that x is the only element in X at distance at most δ of y, hence we have $\pi(y) = x$ and π is well-defined and surjective.

Moreover, any other choice of $\pi(y) = x' \neq x$ will break the property $\max_{y \in Y} d(y, \pi(y))$, and $\pi_X|_Y$ is the only choice.

□

We next show that Axiom 3 is simply an instance of the definition of continuity between two metric spaces. For S a finite subset of E, recall that we denote by $\Delta(S) := \big\{p_S : S \mapsto [0,1] \mid \sum_{x \in S} p_S(x) = 1\big\}$ the set of probability distributions over the elements of S. We further define $\Delta_{\mathcal{P}}(E) := \bigcup_{S \in \mathcal{P}(E)} \Delta(S)$ to be the set of probability distributions over all finite subsets of E. For two finite subsets X, Y of E with $|X| \leq |Y|$ and respective probability distributions p_X in $\Delta(X)$ and p_Y in $\Delta(Y)$, we then define the map $d_\Delta : \Delta_{\mathcal{P}}(E) \times \Delta_{\mathcal{P}}(E) \mapsto \mathbb{R}_{\geq 0}$ as follows:

$$\begin{aligned} d_\Delta(p_X, p_Y) = \min_{\substack{\pi : Y \mapsto X \\ \pi \text{ surjective}}} \max \Bigg\{ & \max_{y \in Y} d(y, \pi(y)), \\ & \sum_{x \in X} \Big| p_X(x) - \sum_{y \in \pi^{-1}(x)} p_Y(y) \Big| \Bigg\} \\ = d_\Delta(p_Y, p_X). & \end{aligned}$$

Lemma 3. *The map d_Δ constitutes a metric on $\Delta_{\mathcal{P}}(E)$.*

Proof. Note that *symmetry* and *non-negativity* clearly hold by definition. We hence focus on *separability* and *triangle inequality.*

- *Separability.* Let X, Y be two finite subsets of E satisfying $|X| \leq |Y|$, and let p_X, p_Y be probability distribution on the respective sets.

  On one hand, we verify that $d_\Delta(p_X, p_X) = 0$: indeed, the choice $\pi = \mathrm{Id}$ renders both terms of $d_\Delta(p_X, p_X)$ null since d and $\|\cdot\|_1$ both satisfy distinguishability.

  On the other hand, suppose that $d_\Delta(p_X, p_Y) = 0$, and let $\pi : Y \mapsto X$ be the surjective map achieving the minimum. Since the first term is null and d verifies distinguishability, we directly get that $X = Y$. Since the second term is moreover null, we moreover obtain $p_X = p_Y$, and d_Δ also verifies distinguishability.

- *Triangle inequality.* Let X, Y, Z be three sets in $\mathcal{P}(E)$ satisfying $|X| \leq |Y| \leq |Z|$, and p_X, p_Y, p_Z be probability distribution on the respective sets. Let $\pi_Y : Z \mapsto Y$ and $\pi_X : Y \mapsto X$ be the two surjective maps that achieve the minimum in $d_\Delta(p_Z, p_Y)$ and in $d_\Delta(p_Y, p_X)$ respectively; we then denote by π the surjective map $\pi_X \circ \pi_Y : Z \mapsto X$.

$$
\begin{aligned}
& d_\Delta(p_X, p_Z) \\
& \overset{(a)}{\leq} \max \Bigg\{ \max_{z \in Z} d(z, \pi(z)), \\
& \sum_{x \in X} \Big| p_X(x) - \sum_{z \in \pi^{-1}(x)} p_Z(z) \Big| \Bigg\}, \\
& \overset{(b)}{\leq} \max \Bigg\{ \max_{z \in Z} d(z, \pi_Y(z)) + d(\pi_Y(z), \pi(z)), \\
& \sum_{x \in X} \Big| p_X(x) - \sum_{y \in \pi_X^{-1}(x)} p_Y(y) \Big| \\
& + \sum_{y \in \pi_X^{-1}(x)} \Big| p_Y(y) - \sum_{z \in \pi_Y^{-1}(y)} p_Z(z) \Big| \Bigg\}, \\
& \overset{(c)}{\leq} \max \Bigg\{ \max_{z \in Z} d(z, \pi_Y(z)) + \max_{y \in Y} d(y, \pi_X(y)), \\
& \sum_{x \in X} \Big| p_X(x) - \sum_{y \in \pi_X^{-1}(x)} p_Y(y) \Big| \\
& + \sum_{y \in Y} \Big| p_Y(y) - \sum_{z \in \pi_Y^{-1}(y)} p_Z(z) \Big| \Bigg\}, \\
& \overset{(d)}{\leq} d_\Delta(p_X, p_Y) + d_\Delta(p_Y, p_Z).
\end{aligned}
$$

Inequality (a) follows from the minimum in the definition of $d_\Delta(p_X, p_Z)$ being smaller than with the particular choice $\pi = \pi_X \circ \pi_Y$; inequality (b) use the triangular inequality, as well as the partition $\pi^{-1}(x) = \bigcup_{y \in \pi_X^{-1}(x)} \pi_Y^{-1}(y)$. Inequality (c) holds by taking the maximum over Y in the first term and summing over the whole set Y instead of simply $\pi^{-1}(x)$ in the second term; inequality (d) finally uses the inequality $\max\{a+b, u+v\} \leq \max\{a,u\} + \max\{b,v\}$, as well as the definitions of π_Y and π_X.

□

Combining these results, we finally characterize the representation functions that satisfy Axiom 3.

Lemma 4. *A representation function f satisfying Axiom 3 is precisely a continuous map from $(\mathcal{P}(E), d_H)$ to $(\Delta_\mathcal{P}(E), d_\Delta)$.*

Proof. Note that a representation function f is indeed a map from $\mathcal{P}(E)$ to $\Delta_\mathcal{P}(E)$, we hereafter focus on the relationship between Axiom 3 and continuity between metric spaces.

We first prove the direction $\Longrightarrow$. Consider an arbitrary small $\epsilon > 0$ and a finite subset X in $\mathcal{P}(E)$. By Axiom 3 and Lemma 1, there exists $\delta > 0$ small enough such that $\max_{x \in X} \left| f(X)(x) - \sum_{y \in \pi^{-1}(x)} f(Y)(y) \right| \leq \epsilon / |X|$ holds for all finite subset Y satisfying $d_H(X, Y) \leq \delta$, where the canonical projection $\pi = \pi_X|_Y$ is surjective. For all such subset Y, we then get

$$
\begin{aligned}
d_\Delta(f(X), f(Y)) \leq \max \Bigg\{ & \max_{y \in Y} d(y, \pi(y)), \\
& \sum_{x \in X} \Big| f(X)(x) - \sum_{y \in \pi^{-1}(x)} f(Y)(y) \Big| \Bigg\}.
\end{aligned}
$$

Note that the first term is bounded by $\delta \leq \epsilon$ by definition of Y and π. Moreover, as shown above, each of the $|X|$ terms of the sum is bounded by $\epsilon/|X|$. We hence conclude that $d_\Delta(f(X), f(Y)) \leq \epsilon$, and f is indeed a continuous map between the two metric spaces $(\mathcal{P}(E), d_H)$ and $(\Delta_\mathcal{P}(E), d_\Delta)$.

We next turn to the direction $\Longleftarrow$. Consider a finite subset X in $\mathcal{P}(E)$ and an arbitrary ϵ such that $\underline{d}(X)/2 > \epsilon > 0$. By definition of continuity, there exists $\delta > 0$ such that for all subsets Y in $\mathcal{P}(E)$ satisfying $d_H(X, Y) \leq \delta$, it holds that $d_\Delta(f(X), f(Y)) \leq \epsilon$. Let π be the minimizer in $d_\Delta(f(X), f(Y))$; since $\pi : Y \mapsto X$ is a surjective map satisfying $d(y, \pi(y)) \leq \epsilon$ for all $y \in Y$, Lemma 1 implies that $\pi = \pi_X|_Y$. We then directly get Axiom 3, i.e.,

$$
\max_{x \in X} \Big| f(X)(x) - \sum_{y \in \pi^{-1}(x)} f(Y)(y) \Big| \leq d_\Delta(f(X), f(Y)) \leq \epsilon.
$$

□

We conclude this section with a discussion of Axioms 3, 4, 5 and 6. First, note that the formulation of Axiom 4 stands out from the three others because it inverses the quantifiers for X and δ, in the spirit of uniform continuity. Note that this detail is of great importance as the non-uniform equivalent of Axiom 4 would always be trivially satisfied for a given finite subset X by choosing δ strictly smaller than $\underline{d}(X)$. Even in the case of perfect clones in a pseudo-metric space (c.f. Section 5), this weaker form of Axiom 4 would not be interesting since perfect clones, being in the same isomorphism class, would already obtain the exact same representation under Axiom 2.

One could however wonder whether stronger uniform versions of the remaining axioms could be considered. We argue this is not easily done for Axioms 3 and 5. Indeed, both are defined in terms of the canonical projection $\pi_X|_Y$, which is only necessarily surjective for a value of δ smaller than $\underline{d}(X)/2$, i.e., that depends on the choice of the underlying finite subset X (c.f. Lemma 1). Note moreover that surjectivity is vital: for each element x in X and $\delta > \underline{d}(X)/2$, we can find a choice of Y such that x is not in the codomain of $\pi_X|_Y$, and x should receive representation smaller than ϵ. This is of course a contradiction. Furthermore, switching to different surjective maps is not an alternative: what about symmetric cases, e.g., $X = \{(0,0), (1,1)\}$ and $Y = \{(0,1), (1,0)\}$ in $\mathbb{R}^2$? One would then have to consider multiple maps simultaneously, each with contradictory constraints.

The same issue does not occur for Axiom 6 however, since it compares the representation of the same element z in two

different nearby sets. One could then strengthen Axiom 6 by asking for δ to be independent of the underlying subset X.

Axiom 9 (Uniform α-Locality under Addition of Clones)**.** *The addition of a clone only changes the representation locally, i.e., for each $\epsilon > 0$, there exists $\delta > 0$ such that for all finite subset $S \in \mathcal{P}(E)$, element $x \in S$ and δ-clone x' satisfying $d(x, x') \leq \delta$, we have for all $z \in S$ such that $d(x, z) \geq \alpha$ that $|f(S)(z) - f(S \cup \{x'\})(z)| \leq \epsilon$.*

Another possible direction to strengthen Axiom 6 would be to require the same to hold also for $\epsilon = 0$.

Axiom 10 (Strict α-Locality under Addition of Clones)**.** *The addition of a clone only changes the representation locally, i.e., for a finite subset $S \in \mathcal{P}(E)$, there exists $\delta > 0$ such that, for each element $x \in S$ and δ-clone x' satisfying $d(x, x') \leq \delta$, we have for all $z \in S$ such that $d(x, z) \geq \alpha$ that $f(S)(z) = f(S \cup \{x'\})(z)$.*

Note that the family of representation function g_r we introduced in Section 4 verifies neither of these strengthenings.

C Proofs of the Main Results

This section gathers the proofs of Theorems 1 and 2. Most of the technical tools required for the demonstrations are introduced in Appendix A. Before delving into the proof of Theorem 1, let us first recall its formulation.

Theorem 1. *For $r > 0$, the representation function g_r is well-defined and belongs in $\mathcal{R}_{2r}(\mathbb{R}^n, d_2)$.*

Proof. Let $r > 0$ be a positive radius, we first verify that g_r is a well-defined representation function. Let S be a finite subset of $\mathbb{R}^n$, and x be an element of S, we show that $g_{r,S,x}$ is measurable with respect to the Lebesgue measure μ. Indeed, consider the following partition of its domain

$$B_r(S) = \bigcup_{U \subseteq S} A_r(U),$$

where $A_r(U) = \bigcap_{u \in U} B_r(u) \bigcap_{v \in S \setminus U} \big(B_r(S) \setminus B_r(v)\big)$ belongs in the Borel σ-algebra for all choice of $U \subseteq S$. Note that this partition is finite since there are at most $2^{|S|}$ choices for the subset U.

Moreover, $g_{r,S,x}$ is null on $B_r(S) \setminus B_r(x)$ and, for all $U \subseteq S$ with $x \in U$ and all y in $A_r(U)$, we have $g_{r,S,x}(y) = \frac{1}{|U|}$. Using the above partition, we rewrite

$$g_{r,S,x} : y \in B_r(S) \mapsto \sum_{\{x\} \subseteq U \subseteq S} \frac{\mathbb{1}_{A_r(U)}(y)}{|U|}, \qquad (1)$$

and recognize a simple non-negative function. As such, $g_{r,S,x}$ is both measurable and integrable, and $g_r(S)$ is well-defined. Moreover, we verify that $g_r(S)$ is normalized.

$$\begin{aligned}
\sum_{x \in S} g_r(S)(x) &= \sum_{x \in S} \int_{B_r(S)} \frac{g_{r,S,x}}{\mu(B_r(S))}\, d\mu, \\
&= \int_{B_r(S)} \frac{1}{\mu(B_r(S))} \sum_{x \in S} \frac{\mathbb{1}_{B_r(x)}(y)}{\sum_{z \in S} \mathbb{1}_{B_r(z)}(y)}\, d\mu(y), \\
&= \int_{B_r(S)} \frac{1}{\mu(B_r(S))}\, d\mu = 1.
\end{aligned}$$

Since $g_r(S)$ is also non-negative, it is a probability distribution over S, hence g_r is indeed a well-defined representation function over $\mathcal{P}(\mathbb{R}^n)$.

Axiom 1. Let S be a finite subset of $\mathbb{R}^n$, and x be an element of S. Using the expression of $g_{r,S,x}$ in Equation (1), we get for each y in $B_r(x)$ that $g_{r,S,x}(y) \geq \frac{1}{|S|}$. Since $\mu\big(B_r(S)\big) \leq |S| \cdot \mu\big(B_r(x)\big)$ by countable additivity and uniformity of μ, we finally obtain

$$\begin{aligned}
g_r(S)(x) &\overset{(a)}{=} \int_{B_r(x)} \frac{g_{r,S,x}}{\mu\big(B_r(S)\big)}\, d\mu, \\
&\geq \frac{1}{|S|} \int_{B_r(x)} \frac{1}{\mu\big(B_r(S)\big)}\, d\mu, \qquad (2) \\
&= \frac{1}{|S|} \frac{\mu\big(B_r(x)\big)}{\mu\big(B_r(S)\big)} \geq \frac{1}{|S|^2}.
\end{aligned}$$

Equality (a) holds since $g_{r,S,x}$ is null on $B_r(S) \setminus B_r(x)$. Hence we have $g_r(S)(x) > 0$ for all $x \in S$, and Axiom 1 holds.

Axiom 2. Let $S \subset \mathbb{R}^n$ be a finite subset, $\sigma_S : S \mapsto S$ be a self-isometry on S, and x be an element of S. Since $g_{r,S,x}$ is a simple function, we deduce from Equation (1) the following, i.e.,

$$\begin{aligned}
g_r(S)(\sigma(x)) &= \frac{1}{\mu\big(B_r(S)\big)} \sum_{\{\sigma(x)\} \subseteq U \subseteq S} \frac{\mu\big(A_r(U)\big)}{|U|}, \\
&\overset{(a)}{=} \frac{1}{\mu\big(B_r(S)\big)} \sum_{\{x\} \subseteq V \subseteq S} \frac{\mu\big(A_r(\sigma_S(V))\big)}{|\sigma_S(V)|}, \\
&\overset{(b)}{=} \frac{1}{\mu\big(B_r(S)\big)} \sum_{\{x\} \subseteq V \subseteq S} \frac{\mu\big(T \circ L(A_r(V))\big)}{|V|}, \\
&\overset{(c)}{=} \frac{1}{\mu\big(B_r(S)\big)} \sum_{\{x\} \subseteq V \subseteq S} \frac{\mu\big(A_r(V)\big)}{|V|}, \\
&= g_r(S)(x).
\end{aligned}$$

Equality (a) holds by writing $V = \sigma_S^{-1}(U) = \{\sigma_S^{-1}(u) \mid u \in U\}$ and noting that V contains $\{x\}$. Equality (b) uses that $|\sigma_S(V)| = |V|$ and the decomposition of the isometry σ_S as a linear transformation $L : x \mapsto Qx$ and a translation $T : x \mapsto x + t$, where Q and t are the orthogonal matrix and the vector of Lemma 5. Using that $T \circ L$ is an isometry on $\mathbb{R}^n$, we then rewrite $A_r(\sigma_S(V)) = T \circ L(A_r(V))$.

Finally, Equality (c) follows from the invariance of the Lebesgue measure by translation $\mu\big(T \circ L(A_r(V))\big) = \mu\big(L(A_r(V))\big)$, as well as its behavior under linear transformation $\mu\big(L(A_r(V))\big) = |\det(Q)| \cdot \mu\big(A_r(V)\big) = \mu\big(A_r(V)\big)$, where we used that Q is orthogonal. Hence Axiom 2 holds.

Since the proofs of Axioms 4, 5 and 6 use similar techniques, we first introduce the necessary tools in the more complex case of Axiom 5, and later use similar arguments to show Axiom 6 and then 4.

Axiom 5. Let X be a finite subset of $\mathbb{R}^n$, and let δ satisfy $\min\{r, \underline{d}(X)/2\} > \delta > 0$. Consider a subset Y of $\mathbb{R}^n$ satisfying $|X| = |Y|$ and $d_H(X, Y) \leq \delta$. Note that Lemma 1

ensures that $\pi = \pi_X|_Y$ the cardinal projection on X is bijective and verifies $d(y, \pi(y)) \leq \delta$ for all y in Y.

For U a subset of X, we denote by $\partial A_r(U)$ the boundary of $A_r(U)$. We moreover associate with $A_r(U)$ its "thick interior" $A_r^{-\delta}(U) = \big\{x \in S \mid d(x, \partial A_r(U)) > \delta\big\}$, as well as its "thick closure" $A_r^{+\delta}(U) = \big\{x \in \mathbb{R}^n \mid d(x, A_r(U)) \leq \delta\big\}$. Note that these definitions allow to obtain the "thick boundary" of $A_r(U)$ by set difference, i.e., $\partial A_r(U) + \overline{B}_\delta(o) = A_r^{+\delta}(U) \setminus A_r^{-\delta}(U)$.

We next show that $A_r^{-\delta}(U) \subseteq A_r(\pi(U))$; let z be an element of $A_r^{-\delta}(U)$.

- Let u be an element of U and let z' be the projection of z on the boundary of $B_r(u)$. Note in particular that z belongs to the segment $[z', u]$. This implies that $d(z, u) = d(z', u) - d(z, z') < r - \delta$ by definition of $A_r^{-\delta}(U)$. By triangle inequality, we then get $d(z, \pi(u)) \leq d(z, u) + d(u, \pi(u)) \leq \delta$. Note that this holds for all choices of u in U, hence also for all $\pi(u)$ in $\pi(U)$.
- Now let v be an element of $X \setminus U$ and let z' be the projection of z on the boundary of $B_r(v)$. Note this time that z' belongs in the segment $[z, v]$ and we similarly obtain $d(z, v) = d(z', z) + d(z', v) > r + \delta > r - \delta$ using the definition of $A_r^{-\delta}(U)$. The triangle inequality again gives $d(z, \pi(v)) \geq d(z, v) - d(v, \pi(v)) \geq \delta$ for all $\pi(v) \in Y \setminus \pi(U)$.

Together, this implies that z lies in $A_r(\pi(U))$ and we conclude that $A_r^{-\delta}(U) \subseteq A_r(\pi(U))$.

Since $A_r^{-\delta}(U)$ is also a subset of $A_r(U)$, Equation (1) implies, for all x in X and z in $A_r^{-\delta}(U)$, that $g_{r,X,x}(z) = \mathbb{1}_U(z)/|U| = g_{r,Y,\pi(x)}(z)$. Since this holds for all choices of $U \subseteq X$, we finally get that $g_{r,X,x} = g_{r,Y,\pi(x)}$ on $A_r^{-\delta,X} := \bigcup_{U \subseteq X} A_r^{-\delta}(U)$.

We now bound the difference of representation between x in X and $\pi(x)$ in Y using the partition of $B_r(X \cup Y)$ induced by $A_r^{-\delta,X}$.

$$
\begin{aligned}
&\big|g_r(X)(x) - g_r(Y)(\pi(x))\big| \\
&= \left| \int_{B_r(X)} \frac{g_{r,X,x}}{\mu\big(B_r(X)\big)} d\mu - \int_{B_r(Y)} \frac{g_{r,Y,\pi(x)}}{\mu\big(B_r(Y)\big)} d\mu \right|, \\
&\overset{(a)}{\leq} \left| \frac{\mu\big(B_r(X)\big)}{\mu\big(B_r(Y)\big)} - 1 \right| \cdot \left| \int_{A_r^{-\delta,X}} \frac{g_{r,X,x}}{\mu\big(B_r(X)\big)} d\mu \right| \\
&+ \int_{B_r(X \cup Y) \setminus A_r^{-\delta,X}} \frac{1}{\mu\big(B_r(x)\big)} d\mu, \\
&\overset{(b)}{\leq} \frac{\mu\big(B_{r+\delta}(X)\big) - \mu\big(B_r(X)\big)}{\mu\big(B_r(\pi(x))\big)} + \frac{\mu\big(B_{r+\delta}(X)\big) - \mu\big(A_r^{-\delta,X}\big)}{\mu\big(B_r(x)\big)}, \\
&\overset{(c)}{\leq} \frac{2}{V_r^n} \sum_{U \subseteq X} \mu\big(A_r^{+\delta}(U)\big) - \mu\big(A_r^{-\delta}(U)\big).
\end{aligned}
\tag{3}
$$

Inequality (a) follows from the functional equality $g_{r,X,x} = g_{r,Y,\pi(x)}$ on $A_r^{-\delta,X}$ and the fact that the two functions have their image in $[0, 1]$ otherwise; it also uses that both $\mu\big(B_r(X)\big)$ and $\mu\big(B_r(Y)\big)$ are greater than $\mu\big(B_r(x)\big)$. We establish inequality (b) by noting that $B_r(\pi(x)) \subseteq B_r(Y) \subseteq B_{r+\delta}(X)$ and that the integral in the first term is bounded by one. Finally, inequality (c) holds by writing V_r^n for the volume of the n-dimensional Euclidean ball of radius r, and using the inclusion $B_{r+\delta}(X) \subseteq \bigcup_{U \subseteq X} A_r^{+\delta}(U)$ in combination with Boole's inequality and the definition of $A_r^{-\delta,X}$.

For $U \subseteq X$, remark that the boundary $\partial A_r(U)$ is covered by the countable union of smooth $n-1$ manifolds $\bigcup_{x \in X} \partial B_r(x)$, hence it is $n-1$ rectifiable. By [Federer, 1996, Theorem 3.2.39], the Minkovski content of the boundary $\mathcal{M}^{n-1}\big(\partial A_r(U)\big)$ exists and is equal to its $n-1$ dimensional Hausdorff measure $\mathcal{H}^{n-1}\big(\partial A_r(U)\big)$, i.e.,

$$
\begin{aligned}
\mathcal{M}^{n-1}\big(\partial A_r(U)\big) &= \lim_{\delta \to 0} \frac{\mu\big(\partial A_r(U) + B_\delta(o)\big)}{2\delta}, \\
&\overset{(a)}{=} \lim_{\delta \to 0} \frac{\mu\big(A_r^{+\delta}(U)\big) - \mu\big(A_r^{-\delta}(U)\big)}{2\delta}, \\
&= \mathcal{H}^{n-1}\big(\partial A_r(U)\big).
\end{aligned}
$$

Equality (a) follows from the definitions of the sets $A_r^{+\delta}(U)$ and $A_r^{-\delta}(U)$ leading to the set equality $\partial A_r(U) + \overline{B}_\delta(o) = A_r^{+\delta}(U) \setminus A_r^{-\delta}(U)$.

We can hence take $\delta > 0$ small enough so as to satisfy, for all subsets $U \subseteq X$,

$$
\mu\big(A_r^{+\delta}(U)\big) - \mu\big(A_r^{-\delta}(U)\big) \leq 4\delta \mathcal{H}^{n-1}(\partial A_r(U)). \tag{4}
$$

With such a choice of δ, Equation (3) then becomes, i.e.,

$$
\begin{aligned}
\big|g_r(X)(x) - g_r(Y)(\pi(x))\big| &\leq \frac{8\delta}{V_r^n} \sum_{U \subseteq X} \mathcal{H}^{n-1}(\partial A_r(U)), \\
&\overset{(a)}{\leq} \frac{16\delta}{V_r^n} \sum_{x \in X} \mathcal{H}^{n-1}(\partial B_r(x)), \\
&\overset{(b)}{=} \frac{16\delta |X| n}{r}.
\end{aligned}
\tag{5}
$$

Inequality (a) follows from the set equality $\bigcup_{U \subseteq X} \partial A_r(U) = \bigcup_{x \subseteq X} \partial B_r(x)$, as well as by noting that each component of the boundary $\partial A_r(U)$ is counted at most twice over all subsets $U \subseteq X$. Equality (b) uses the relation between the surface and the volume of an n-dimensional sphere of radius r, i.e., $\mathcal{H}^{n-1}\big(\partial B_r(x)\big) = S_r^{n-1} = \frac{n}{r} V_r^n$.

For a fixed radius $r > 0$ and an arbitrary $\epsilon > 0$, we can finally choose $\delta \leq \epsilon r / \big(16|X|n\big)$ small enough, and Axiom 5 holds.

Axiom 6. Let $S \subseteq \mathbb{R}^n$ be a finite subset, x be an element of S and $x' \in \mathbb{R}^n \setminus S$ be a δ-clone of x, where δ verifies $\min\{r, \underline{d}(S)/2\} > \delta > 0$. We moreover denote by S' the union $S \cup \{x'\}$.

The proof follows the same structure as that of Axiom 5: for y an element of S such that $d(x, y) \geq 2r$, we first show that $g_{r,S,y}$ and $g_{r,S',y}$ are equal on a carefully chosen set. Indeed for z an element of $B_{r-\delta}(y)$, the triangle inequality gives $d(z, x') \geq d(x, y) - d(x, x') - d(x', y) > 2r - \delta - (r - \delta) = r$, and we obtain that $g_{r,S,y}(z) = g_{r,S',y}(z)$.

We then bound the difference of representation between S and S' in a similar fashion as in Equation (3).

$$\begin{aligned}
&\big|g_r(S)(y) - g_r(S')(y)\big| \\
&\overset{(a)}{\leq} \left(\frac{\mu\big(B_r(S')\big)}{\mu\big(B_r(S)\big)} - 1\right) \cdot \left| \int_{B_{r-\delta}(y)} \frac{g_{r,S',y}}{\mu\big(B_r(S')\big)} d\mu \right| \\
&+ \int_{B_r(y)\setminus B_{r-\delta}(y)} \frac{1}{\mu\big(B_r(S)\big)} d\mu, \\
&\overset{(b)}{\leq} \frac{\mu\big(B_{r+\delta}(S)\big) - \mu\big(B_r(S)\big)}{\mu\big(B_r(y)\big)} + \frac{\mu\big(B_r(y)\big) - \mu\big(B_{r-\delta}(y)\big)}{\mu\big(B_r(y)\big)}, \\
&\overset{(c)}{\leq} \frac{4\delta}{V_r^n}\Big(S_r^{n-1} + \sum_{U\subseteq S} \mathcal{H}^{n-1}(\partial A_r(U))\Big).
\end{aligned}$$

Inequalities (a) and (b) follow from similar arguments as Equation (3)(a) and (b) respectively; inequality (c) holds for δ small enough by combining the arguments of Equation (3)(c) and Equation (4).

A similar double-counting argument as in Equation (5)(a) finally gives

$$\big|g_r(S)(y) - g_r(S')(y)\big| \leq \frac{4\delta(2|S|+1)n}{r} \leq \epsilon, \qquad (6)$$

where the last inequality holds for an arbitrary $\epsilon > 0$ and a choice of $\delta \leq \epsilon r/\big(4n(2|S|+1)\big)$ small enough. Hence Axiom 6 holds with $\alpha = 2r$.

Axiom 4. Let δ be a positive number satisfying $r > \delta > 0$, and S be a finite subset of $\mathbb{R}^n$. We moreover let x, y be two elements of S such that $d(x,y) \leq \delta$.

Similarly as for the proof of Axiom 5, we first show that $B_{r-\delta}(x) \subset B_r(x) \cap B_r(y)$. Indeed, for z an element of $B_{r-\delta}(x)$, the triangle inequality gives $d(z,y) \leq d(z,x) + d(x,y) < r - \delta + \delta = r$. Equation (1) then implies the functional equality $g_{r,S,x} = g_{r,S,y}$ on $B_{r-\delta}(x)$.

We then bound the difference of representation between x and y in a similar fashion as in Equation (3).

$$\begin{aligned}
\big|g_r(S)(x) - g_r(S)(y)\big| &\overset{(a)}{\leq} \int_{B_r(x)\setminus B_{r-\delta}(x)} \frac{1}{\mu\big(B_r(S)\big)} d\mu, \\
&\overset{(b)}{\leq} \frac{\mu\big(B_r(x)\big) - \mu\big(B_{r-\delta}(x)\big)}{\mu\big(B_r(x)\big)} \\
&\overset{(c)}{\leq} \frac{4\delta S_r^{n-1}}{V_r^n} = \frac{4\delta n}{r},
\end{aligned} \qquad (7)$$

where inequality (a) and (b) use similar arguments as for Equation (3)(a) and (b) respectively; inequality (c) follows from combining the arguments of Equations (4) and (5). In particular, note that $\mu\big(B_r(x)\big) - \mu\big(B_{r-\delta}(x)\big) \leq \mu\big(B_{r+\delta}(x)\big) - \mu\big(B_{r-\delta}(x)\big) \leq 4\delta S_r^{n-1}$ holds for a value of δ small enough but independent of S since the Lebesgue measure is invariant by translation.

Then for an arbitrary $\epsilon > 0$, a choice of $\delta \leq \epsilon r/(4n)$ small enough ensures that Axiom 4 holds.

Conclusion. Since g_r is a representation function on $(\mathbb{R}^n, d_2)$ satisfying Axioms 1, 2, 4, 5 and 6 with $\alpha = 2r$, we conclude that g_r belongs in $\mathcal{R}_{2r}(\mathbb{R}^n, d_2)$.

□

We now recall the formulation of Theorem 2 before attacking its demonstration.

Theorem 2. *Let ν be a probability density function over $[0, \alpha/2]$. Then the representation function $f_\nu : S \in \mathcal{P}\big(\mathbb{R}^n\big) \mapsto \int_0^{\alpha/2} \nu(r) g_r(S)\, dr$ belongs in $\mathcal{R}_\alpha(\mathbb{R}^n, d_2)$.*

Proof. Let ν be a probability density function over $[0, \alpha/2]$, that is a non-negative Lebesgue-integrable function satisfying $\int_0^{\alpha/2} \nu(r) dr = 1$. Let S be a finite subset of $\mathbb{R}^n$ and x be an element of S.

First, note that $r \mapsto g_r(S)(x)$ is a non-negative step function over $\mathbb{R}_{>0}$, hence $r \mapsto \nu(r) \cdot g_r(S)(x)$ is Lebesgue-integrable and $f_\nu(S)(x)$ is non-negative. Moreover, we have

$$\begin{aligned}
\sum_{x\in X} f_\nu(S)(x) &= \int_{(0,\alpha/2]} \nu(r) \sum_{x\in X} g_r(S)(x)\, dr, \\
&\overset{(a)}{=} \int_{(0,\alpha/2]} \nu(r)\, dr = 1,
\end{aligned}$$

where equality (a) uses that g_r is a representation function. This ensures that f_ν is also a representation function of $(\mathbb{R}^n, d_2)$.

We now verify that Axioms 1, 2, 4, 5 and 6 are implied for f_ν from the fact that g_r belongs in $\mathcal{R}_\alpha(\mathbb{R}^n, d_2)$ for all $\alpha/2 \geq r > 0$. On the one hand, Axiom 1 is directly implied from Equation (2), i.e,

$$\begin{aligned}
f_\nu(S)(x) &= \int_{(0,\alpha/2]} \nu(r) g_r(S)(x)\, dr, \\
&\geq \int_{(0,\alpha/2]} \frac{\nu(r)}{|S|^2} dr, \\
&= \frac{1}{|S|^2} > 0.
\end{aligned}$$

Furthermore, for an isometry $\sigma_S : S \mapsto S$, we have

$$\begin{aligned}
f_\nu(S)(\sigma_S(x)) &= \int_{(0,\alpha/2]} \nu(r) g_r(S)(\sigma_S(x))\, dr, \\
&\overset{(a)}{=} \int_{(0,\alpha/2]} \nu(r) g_r(S)(x)\, dr, \\
&= f_\nu(S)(x),
\end{aligned}$$

where equality (a) uses that Axiom 2 holds for g_r, hence it also holds for f_ν.

On the other hand, Axioms 4, 5 and 6 require a little more work; we hereafter focus on Axiom 5. Let C be a positive constant in $(0, \alpha/2)$, and X be a finite subset of $\mathbb{R}^n$. For a constant δ satisfying $\min\{C, \underline{d}(X)/2\} > \delta > 0$ and a finite subset $Y \in \mathcal{P}(\mathbb{R}^n)$ satisfying $d_H(X,Y) \leq \delta$ and $|Y| = |X|$, Lemma 1 ensures that the canonical projection $\pi = \pi_X|_Y$ is

bijective and verifies $\max_{y\in Y} d(y,\pi(y)) \leq \delta$. The following then holds, i.e.,

$$\begin{aligned}
&\big|f_\nu(X)(x) - f_\nu(Y)(\sigma(x))\big| \\
&\overset{(a)}{\leq} \int_{(0,\alpha/2]} \nu(r)\big|g_r(X)(x) - g_r(Y)(\pi(x))\big|dr, \\
&\overset{(b)}{\leq} \int_0^C \nu(r)dr + \int_C^{\alpha/2} \nu(r)\frac{16\delta|X|n}{r}dr, \\
&\overset{(c)}{\leq} \mathcal{V}(C) + \frac{16\delta|X|n}{C}.
\end{aligned}$$

Inequality (a) uses the non-negativity of ν and the triangle inequality; inequality (b) holds for δ small enough by bounding the difference $\big|g_r(X)(x) - g_r(Y)(\sigma(x))\big|$ by one in the first term, and using Equation (5) for the second term. Inequality (c) is finally obtained by rewriting $\mathcal{V}(\cdot)$ for the cumulative distribution associated with the density $\nu(\cdot)$, and noting that $r \mapsto \frac{16\delta|X|n}{r}$ is decreasing on $[C, \alpha/2]$.

Now fix $\epsilon > 0$; by continuity of $\mathcal{V}$, there exists $C > 0$ such that $\mathcal{V}(C) = \epsilon/2$. Then for $0 < \delta \leq \epsilon C/(32|X|n)$ small enough, we get $\big|f_\nu(X)(x) - f_\nu(Y)(\sigma(x))\big| \leq \epsilon$ and Axiom 5 holds for f_ν.

Similar arguments combined with Equations (6) and (7) respectively show that Axioms 6 and 4 hold for f_ν. We then conclude that f_ν belongs in $\mathcal{R}_\alpha(\mathbb{R}^n, d_2)$.

□

D Self-Isometries in Euclidean Space

In this section, we show that a self-isometry σ_S on a finite subset $S \subseteq \mathbb{R}^n$ can be uplifted to a full-fledged isometry on $\mathbb{R}^n$, that is a rigid transformation.

Lemma 5. *Let $S = \{x_1, \dots, x_m\}$ be a finite subset of the Euclidean space $\mathbb{R}^n$ and $\sigma_S : S \mapsto S$ be a self-isometry on S. There then exists an $n \times n$-orthogonal matrix Q and an n-dimensional vector t such that, for all $i \in [m]$, we have $\sigma(x_i) = Qx_i + t$.*

Proof. For an index $2 \leq i \leq m$, we define $y_i = x_i - x_1$ and $z_i = \sigma(x_i) - \sigma(x_1)$. We then concatenate the y_i (resp. z_i) and define the $n \times (m-1)$ matrix $Y = [y_2, \dots, y_n]$ (resp. $Z = [z_2, \dots, z_m]$). For $i, j \in [m-1]$, note that the following holds:

$$\begin{aligned}
(Y^\top Y)_{i,j} &= y_{i+1} \cdot y_{j+1}, \\
&= \frac{1}{2}\Big(\|y_{i+1} + y_{j+1}\|^2 - \|y_{i+1}\|^2 - \|y_{j+1}\|^2\Big), \\
&= \frac{1}{2}\Big(d(x_{i+1}, x_{j+1})^2 - d(x_{i+1}, x_1)^2 \\
&\quad - d(x_{j+1}, x_1)^2\Big), \\
&\overset{(a)}{=} \frac{1}{2}\Big(d\big(\sigma(x_{i+1}), \sigma(x_{j+1})\big)^2 - d\big(\sigma(x_{i+1}), \sigma(x_1)\big)^2 \\
&\quad - d\big(\sigma(x_{j+1}), \sigma(x_1)\big)^2\Big), \\
&= z_{i+1} \cdot z_{j+1} = (Z^\top Z)_{i,j}.
\end{aligned}$$

Equality (a) holds since σ is an isometry on S.

By [Horn and Johnson, 2012, Theorem 3.7.11], there exists an orthogonal $n \times n$ matrix Q such that $Z = QY$, and we obtain for all $i \in [m]$ that $\sigma(x_i) = Q(x_i - x_1) + \sigma(x_1)$ (the case $i = 1$ holds trivially). Rewriting $t = \sigma(x_1) - Qx_1$ gives the desired result.

□

Note moreover that the Euclidean group $E(n)$, i.e., the group of isometries in Euclidean space, is exactly the semi-direct product of the orthogonal group $O(n)$ extended by the translational group $T(n)$. In other words, Lemma 5 ensures that all self-isometries $\sigma_S : S \mapsto S$ on a finite subset $S \in \mathbb{R}^n$ can be extended to a full-fledged isometry $\sigma : \mathbb{R}^n \mapsto \mathbb{R}^n$.

E Pseudo-metric and Metric Identification

In this section, we expand on the discussion in Section 5 regarding the extension of our framework to perfect clones.

Let (E, d) be a pseudo-metric space, that is an ordered pair where E is a set and $d : E \times E \mapsto \mathbb{R}_{\geq 0}$ is a pseudo-metric on E satisfying, for all $x, y, z \in E$, i.e.,

1. (*Non-negativity*) $d(x, y) \geq 0$,
2. (*Symmetry*) $d(x, y) = d(y, x)$,
3. (*Triangle inequality*) $d(x, z) \leq d(x, y) + d(y, z)$;
4. (*Identity*) $d(x, x) = 0$.

We next show that the pseudo-metric d implicitly defines an equivalence relation $\sim$ on E, which we refer to as the *metric identification*. We denote by $[x] = \{y \in E \mid x \sim y\}$ the equivalence class of x in E.

Lemma 6 (Metric Identification in (E, d))**.** *The binary relation defined for all $x, y \in E$ by $x \sim y$ if and only if $d(x, y) = 0$ is an equivalence relation. Moreover, for all $x \sim y$ and z in E, we have $d(x, z) = d(y, z)$.*

Proof. The symmetry and reflexivity of $\sim$ are directly implied by the symmetry and the identity of the pseudo-metric d; there only remains to verify that $\sim$ is transitive.

Let x, y be elements of E such that $x \sim y$, i.e., $d(x, y) = d(y, x) = 0$, and let z be in E. By triangle inequality, we have on one hand $d(x, z) \leq d(x, y) + d(y, z) = d(y, z)$. On the other hand, we also have $d(y, z) \leq d(y, x) + d(x, z) = d(x, z)$, hence we indeed get $d(x, z) = d(y, z)$.

Applying this to $z \in E$ such that $y \sim z$ finally implies $d(x, z) = 0$, i.e., $x \sim z$, and we verify that $\sim$ is transitive. □

We next consider the *quotient space* of E by the equivalence relation $\sim$, that is the set $E^* = E/\!\sim$ of all equivalence classes induced by $\sim$ on E. We may now define the metric $d^* : ([x], [y]) \in E^* \times E^* \mapsto d(x, y)$, and refer to the metric space (E^*, d^*) as the metric space induced by the pseudo-metric space (E, d).

For a finite subset $X \subseteq E$, we similarly define the quotient $X^* = X/\!\sim\, = \{[x] \mid x \in X\}$, and define $\mathcal{P}(E)^*$ to be the set containing all sets X^* , i.e., $\mathcal{P}(E)^* = \{X/\!\sim \mid X \in \mathcal{P}(E)\}$. We may then equip $\mathcal{P}(E)^*$ with the Hausdorff metric d_H^* associated to the metric d^*.

For representation functions f operating on a pseudo-metric space (E, d), we adapt our requirement and demand that Axiom 5 is respected after we collapse all perfect clones. This intuition is formalized as follows.

Axiom 11 (Individual Continuity for Pseudo-metric Spaces)**.** *There exists a representation function f^* of (E^*, d^*) verifying Axiom 5 that is equal to f after collapsing perfect clones, i.e., it holds for all $X \in \mathcal{P}(E)$ and $x \in X$ that $f^*(X^*)([x]) = \sum_{x' \in [x] \cap X} f(X)(x')$.*

In other words, Axiom 11 demands the following: for a finite subset $X \in \mathcal{P}(E)$ and a positive ϵ, there exists $\delta > 0$ such that, for all finite subset $Y \in \mathcal{P}(E)$ verifying $|Y^*| = |X^*|$ and $d_H(X, Y) \leq \delta$, it holds that

$$\max_{x \in X} \left| \sum_{x' \in [x] \cap X} f(X)(x') - \sum_{y \in \pi^{-1}([x]) \cap Y} f(Y)(y) \right| \leq \epsilon,$$

where $\pi = \pi_{X^*}|_{Y^*}$ denotes the canonical projection on X^* restricted to Y^*.