

Small Transformers Compute Universal Metric Embeddings

Anastasis Kratsios

McMaster University

Department of Mathematics

1280 Main Street West, Hamilton, Ontario, L8S 4K1, Canada

KRATSIOA@MCMASTER.CA

Valentin Debarnot

Universität Basel

Department of Computer Science

Basel, 4051, Switzerland

VALENTIN.DEBARNOT@UNIBAS.CH

Ivan Dokmanić

Universität Basel

Department of Computer Science

Basel, 4051, Switzerland

IVAN.DOKMANIC@UNIBAS.CH

Editor: Sayan Mukherjee

Abstract

We study representations of data from an arbitrary metric space \mathcal{X} in the space of univariate Gaussian mixtures equipped with a transport metric (Delon and Desolneux 2020). We prove embedding guarantees for feature maps implemented by small neural networks called *probabilistic transformers*. Our guarantees are of memorization type: we prove that a probabilistic transformer of depth about $n \log(n)$ and width about n^2 can bi-Hölder embed any n -point dataset from \mathcal{X} with low metric distortion, thus avoiding the curse of dimensionality. We further derive probabilistic bi-Lipschitz guarantees, which trade off the amount of distortion and the probability that a randomly chosen pair of points embeds with that distortion. If the geometry of \mathcal{X} is sufficiently regular, we obtain stronger bi-Lipschitz guarantees for all points. As applications, we derive neural embedding guarantees for datasets from Riemannian manifolds, metric trees, and certain types of combinatorial graphs. When instead embedding into multivariate Gaussian mixtures, we show that probabilistic transformers compute bi-Hölder embeddings with arbitrarily small distortion. Our results show that any finite metric dataset, from vertices on a graph to functions a function space, can be faithfully represented in a single representation space, and that the representation can be implemented by a simple transformer architecture. Thus one may only need a modular set of machine learning tools compatible with this one representation space, many of which already exist, for downstream supervised and unsupervised learning from a great variety of data types.

Keywords: Metric Embeddings, Geometric Deep Learning, Optimal Transport, Representation Learning, Transformers, Gaussian Mixtures.

1. Introduction

In machine learning practice we face a daunting variety of data sources, but all of them come equipped with some domain-specific notion of (dis)similarity. The successes of modern machine learning models may be in part attributed to the fact that they explicitly or implicitly learn representations of data that are compatible with these domain-specific notions. A key task in a successful machine learning pipeline is thus to identify a *representation space* \mathcal{R} and a *feature map* $\phi : \mathcal{X} \rightarrow \mathcal{R}$ which faithfully encodes \mathcal{X} . The data space \mathcal{X} and the dissimilarity $d_{\mathcal{X}}$ often form a metric space $(\mathcal{X}, d_{\mathcal{X}})$; “faithfully” then means with minimal distortion to $d_{\mathcal{X}}$. Once data is represented in \mathcal{R} , we can learn and predict from $\phi(\mathcal{X})$ using downstream machine learning tools for classification, regression models, clustering, dimension reduction models, etc.

For a long time, the default choice for \mathcal{R} has been a Euclidean space or a (reproducing kernel) Hilbert space. Yet there is now a body of work both in machine learning (Samal et al., 2018; Eidi and Jost, 2020; Muscoloni et al., 2017; Ganea et al., 2018; Giovanni et al., 2022) and in metric embedding theory (Bourgain, 1985; Matoušek, 1996; Naor et al., 2006) showing that Euclidean and Hilbert spaces are suboptimal for many practically-relevant notions of dissimilarity. This is because Euclidean spaces are “too small” and Hilbert spaces “too flat” to represent the great variety of geometries in modern datasets.

An archetypal situation where low-dimensional non-Euclidean representations outperform high-dimensional Euclidean is when working with hierarchical data, organized on trees or graphs with power-law degree distributions (Ravasz and Barabási, 2003), and equipped with the usual combinatorial (geodesic) metric. Any finite tree \mathcal{X} can be faithfully embedded into the two-dimensional *hyperbolic space* (Sarkar, 2011) while embeddings in d -dimensional Euclidean space require much greater distortion (Gupta, 2000). Since hierarchical data is ubiquitous, this favorable property of hyperbolic spaces has inspired a booming development of “hyperbolic” machine learning models. Examples include clustering (Chami et al., 2020; Tabaghi and Dokmanić, 2021, 2020), PCA-type methods (Chami et al., 2021), classification (López and Strube, 2020), hyperbolic analogues of feedforward networks (Ganea et al., 2018; Shimizu et al., 2021), and several Python packages supporting this Riemannian geometry (Miolane et al., 2020; Kochurov et al., 2020). These advances contributed to state-of-the-art performance when learning from natural language (Zipf, 1949; Dhingra et al., 2018; Le et al., 2019), knowledge graphs (Kolyvakis et al., 2020), social networks (Krioukov et al., 2010; Muscoloni et al., 2017), directed graphs (Munzner, 1997), scenario generation for stochastic phenomena (Pflug and Pichler, 2015), and combinatorial trees (Nickel and Kiela, 2017).

Hyperbolic representations perform well when \mathcal{X} possesses a “tree-like” geometric prior, but many datasets lack such structure or exhibit a more complex one. This has motivated the search for representations of combinatorial graphs in Riemannian manifolds beyond constant-curvature space forms, including products of space forms (mixed-curvature spaces) (Tabaghi et al., 2021; Samal et al., 2018; Eidi and Jost, 2020) and representations that match discrete curvature in heterogeneous rotationally-symmetric manifolds (Giovanni et al., 2022). While these proposals show great promise in addressing the emergent shortcomings of hyperbolic spaces, each of them introduces additional parameters (such as the choice of space forms or their dimensions and curvatures) without obvious nominal values on real-world datasets, and each requires a different set of downstream tools.

We are thus motivated to look for a “universal” representation space. Beyond the geometric considerations in the previous paragraphs, there may exist no machine learning models on \mathcal{X} with favourable approximation, generalization, and optimization properties even for data from a “nice” \mathcal{X} . And when good models do exist, we often work with datasets from some *unknown* low-dimensional subset of \mathcal{X} and this latent structure is not leveraged by the off-the-shelf models. We look for a representation that allows us to relay \mathcal{X} ’s structure to downstream tasks optimally. Furthermore, we seek a representation agnostic to the details of \mathcal{X} ’s geometry in the sense that it should work for any dataset equipped with a metric. The implications of such a “universal” representation space \mathcal{R} for machine learning is that machine learning tools could be “standardized” to process inputs from \mathcal{R} . Upon learning a faithful representation of any dataset \mathcal{X} in \mathcal{R} , these standardized tools could be used for downstream tasks like classification, regression, and dimensionality reduction.

1.1 “Good” Representation Spaces and Learnable Feature Maps, with Guarantees


We consider the problem of *representing* a general metric space $(\mathcal{X}, d_{\mathcal{X}})$ in a representation space \mathcal{R} , using a *feature map* $\phi : \mathcal{X} \rightarrow \mathcal{R}$. We are interested in “good” representation spaces and computationally tractable feature maps.

Let us take a moment to reflect on what a “good” \mathcal{R} and ϕ should entail. Empirically successful non-Euclidean embeddings (Lopez et al., 2021; Giovanni et al., 2022) and Bourgain’s impossibility result for metric embeddings in Hilbert spaces (Bourgain, 1985) suggest that \mathcal{R} must not be *flat*. On the other hand, we show in Section 4.3 that for any complete Riemannian manifold \mathcal{R} (including all manifolds of positive curvature) and for any positive integer n , there exists an n -point metric space which cannot be embedded in

\mathcal{R} with metric distortion smaller than $O(\log n)$. Finite-dimensional manifolds are thus *too small* to embed datasets from arbitrary metric spaces, so we require that \mathcal{R} be infinite-dimensional.

Given a finite dataset from \mathcal{X} , its representation in \mathcal{R} is often computed by solving an optimization problem. One example is Laplacian eigenmaps when \mathcal{X} is a Riemannian manifold and $d_{\mathcal{X}}$ the geodesic distance (Belkin and Niyogi, 2003, 2006). The feature map ϕ is thus defined implicitly via an optimization problem. A drawback of this strategy is that if a new sample is added to the dataset, computing its representation may require recomputing representations of all points. To avoid this, we are interested in efficient direct parameterizations of the feature map by deep neural networks with controlled complexity. In machine learning, this “amortization” of static optimization-based procedures is referred to as (out-of-sample) generalization (Bartlett et al., 2017; Agarwal and Niyogi, 2009; Bartlett and Long, 2021; Mei et al., 2022; Mei and Montanari, 2022). To make it possible, \mathcal{R} should be a familiar representation space with well-understood geometry, finitely parameterized points, and admitting efficient numerical algorithms. Since only finite-dimensional representations can be practically implemented, we ask that the feature map $\phi : \mathcal{X} \rightarrow \mathcal{R}$ implemented by our deep learning model take values in a finite-dimensional subspace of \mathcal{R} whenever \mathcal{X} has a suitable geometric prior (which we make precise shortly). We informally refer to the dimension of the image of ϕ as the *effective dimension* of the representation ϕ of \mathcal{X} .

We show that the requirements for \mathcal{R} are met by a certain space of probability measures equipped with an optimal transport metric, and that the corresponding feature map ϕ can be *exactly implemented* (in the sense of memorization capacity) by a deep neural network with controlled complexity. Concretely, we prove that a deep neural network with about $n \log(n)$ layers of width n^2 can embed any dataset of n points in \mathcal{X} with low metric distortion. Thus, unlike neural networks studied by most universal approximation theorems (Burger and Neubauer, 2001; Yarotsky, 2017; Kratsios and Papon, 2022; Acciaio et al., 2023; Kidger and Lyons, 2020; Puthawala et al., 2022; Shen et al., 2021), the PT does not face the curse of dimensionality. Its number of parameters is a (small-degree) polynomial in the embedded number of points which is independent of \mathcal{X} ’s dimension. In contrast, the networks constructed in universal approximation theorems usually require a number of parameters exponential in \mathcal{X} ’s dimension.



The particular n -point dataset (a subset \mathcal{X}_n of \mathcal{X}) encodes information about general points in \mathcal{X} . Its points act as reference points or *landmarks* for the remainder of \mathcal{X} . If \mathcal{X} is a polytope (for example, $[0, 1]^d$) then \mathcal{X}_n could be the set of extremal points in \mathcal{X} (for example, $\mathcal{X}_n = \{0, 1\}^d$). If \mathcal{X} is a connected compact Riemannian manifold then \mathcal{X}_n could be any maximal δ -dense subset; for a sufficiently small¹ $\delta > 0$ this set of n points encodes most of the metric information in $(\mathcal{X}, d_{\mathcal{X}})$. In machine learning, \mathcal{X}_n is the dataset to be analyzed or the training set which is randomly drawn from \mathcal{X} . In metric embedding theory suitable finite subsets play an analogous role to Schauder bases in the theory of Banach spaces (see (Naor, 2018)).

This places our work at the junction between *metric embedding theory* and *representation learning*. The former aims at proving the existence of low-dimensional and low-distortion embeddings of finite metric spaces (e.g. \mathcal{X}_n), while the latter seeks computationally tractable representations of non-linear data (e.g. \mathcal{X}) which are performant in downstream learning tasks. Our theoretical guarantees concern n -point datasets from \mathcal{X} because one cannot expect to embed general infinite metric spaces into any single “reasonable”² space. Finally, we complement the theory by stylized computer experiments which illustrate low metric distortion of our proposed representations (compared to Euclidean and hyperbolic) as well as out-of-sample generalization.

1. Let d_g be the geodesic distance on \mathcal{X} . The main result of (Katz and Katz, 2011) shows that if δ is less than 10 times the length of the smallest non-contractible loop in \mathcal{X} (i.e. \mathcal{X} ’s systole) then, map $x \mapsto (d_g(x, x_i))_{i=1}^n$ is a high-dimensional Euclidean bi-Lipschitz embedding of the Riemannian manifold \mathcal{X} ; where $\mathcal{X}_n := \{x_i\}_{i=1}^n$. A key point here is the trade-off between the embedding dimension n , which grows exponentially in δ^{-1} (via an elementary covering argument) and the Lipschitz constant of the map $x \mapsto (d_g(x, x_i))_{i=1}^n$ which is large whenever δ is small.
 2. By “reasonable” we mean a separable complete geodesic metric space.

1.2 The Representation Space

Our search for a “good” representation space begins by examining the Wasserstein space, denoted by $(\mathcal{P}_2(\mathbb{R}^d), \mathcal{W}_2)$, ($d \geq 1$), which is both infinite-dimensional and exhibits “*positive curvatures on all scales*” (Kloeckner and Bertrand, 2016, Section 4.1.1). Furthermore, even when $d = 1$, $(\mathcal{P}_2(\mathbb{R}^d), \mathcal{W}_2)$ contains an isometric copy of any finite-dimensional bounded Euclidean ball. Thus, it has plenty of room and a complex-enough geometry to make it a viable candidate for a universal representation space. A fortiori, this intuition has been confirmed by Andoni, Naor, and Neiman (Andoni et al., 2018) who show that $(\mathcal{P}_2(\mathbb{R}^d), \mathcal{W}_2)$ embeds any finite dataset with little or no distortion to \mathcal{X} ’s metric geometry when $d > 2$, with $d = 2$ still an open problem in geometric analysis.

A computational drawback of $(\mathcal{P}_2(\mathbb{R}^d), \mathcal{W}_2)$ is that its elements cannot be *exactly implemented* by a computer (capable of processing real values) but rather only approximated by an expressive class of probability measures such as Gaussian mixtures or empirical distributions (atomic probability measures) (Chevallier, 2018). This makes them ill-suited for sharp embedding guarantees since the optimal approximation rate using, for example, atomic measures with M atoms in $\mathcal{P}_2(\mathbb{R}^d)$ with respect to the Wasserstein distance is $\mathcal{O}(M^{-1/d})$ (Graf and Luschgy, 2000; Kloeckner, 2012; Liu and Pagès, 2020). Moreover, restricting the Wasserstein-2 distance to either the set of Gaussian mixtures or empirical (atomic) probability measures no longer yields a complete geodesic space, resulting in an ill-behaved geometry.

We sidestep these issues by instead working on the smaller optimal transport-theoretic space $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$, introduced by (Delon and Desolneux, 2020), whose elements are (finitely parameterized) *univariate Gaussian mixtures* and whose distance is a strengthening of the Wasserstein distance obtained by considering only those couplings which are themselves Gaussian mixtures. This space, metrized by restricting the optimal couplings, is part of a broader line of recent work which encodes additional structure into the transport distance by constraining the couplings; we highlight adapted optimal transport (Backhoff-Veraguas et al., 2020b; Acciaio et al., 2020; Backhoff et al., 2022), martingale optimal transport (Dolinsky and Soner, 2014; Beiglböck et al., 2017b; Guo and Obłój, 2019; Backhoff-Veraguas and Pammer, 2022b), and semi-martingale optimal transport (Liu and Neufeld, 2019). The marked advantage of $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ over $(\mathcal{P}_2(\mathbb{R}^d), \mathcal{W}_2)$ is that its elements can be exactly implemented (up to numerical precision) which eliminates the quantization error of $\mathcal{O}(M^{-1/d})$. Geometrically, $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ shares many of the appealing properties of the classical Wasserstein space $(\mathcal{P}_2(\mathbb{R}^d), \mathcal{W}_2)$ (e.g., it is a complete separable *geodesic space*); thus, there are no geometric drawbacks of working with $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ instead of with $(\mathcal{P}_2(\mathbb{R}^d), \mathcal{W}_2)$. As byproduct, since the distance \mathcal{MW}_2 is strictly stronger than the classical Wasserstein distance \mathcal{W}_2 , our deep neural embedding results into $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ automatically imply embeddings in $(\mathcal{P}_2(\mathbb{R}), \mathcal{W}_2)$.

To summarize, we move to the space of univariate Gaussian mixtures with transport distances of (Delon and Desolneux, 2020) for which there is a well-developed machine learning machinery (including non-linear dimension reduction (Bigot et al., 2017), clustering (Mi et al., 2018), and regression (Chen et al., 2021)) and which is neither “too small” nor “too flat”. Furthermore, as we show, the feature maps are computationally tractable and can be parameterized by deep neural networks. Finally, several maintained Python libraries streamline the implementation of deep learning algorithms in these spaces (Flamary et al., 2021; Delon and Desolneux, 2021).

1.3 Universal Feature Maps

We implement the feature maps using an extension of the *probabilistic transformer* (PT) model (Kratsios et al., 2022; Kratsios, 2023). PT extends the classical transformer network used in natural language processing³ (Vaswani et al., 2017); in particular, the PT implements the classical transformer “on average”. It can be shown to universally approximate any continuous function while simultaneously *exactly* implementing arbitrary compact (possibly non-convex) constraints (Kratsios et al., 2022).

We adapt the probabilistic transformer by composing it with a non-parameterized input layer which translates metric information into vectorial data by mapping a point in \mathcal{X} to a vector of its distances to a set of

3. And now in many other places in deep learning.




Figure 1: Schematic view of a probabilistic transformer (gray) representing a metric space \mathcal{X} (orange) while bi-Lipschitz embedding the distinguished points (landmarks, data) \mathcal{X}_n (black).




Figure 2: Embedding landscape of probabilistic transformers. Data from structured spaces such as trees or manifolds is “easiest” to embed in terms of feature map complexity. Data from unstructured spaces (general metric graphs, expanders) requires more complex (but still uncursed) neural feature maps.

landmarks $\{x_l\}_{l=1}^n \subseteq \mathcal{X}$ (cf. Figure 6a). The resulting process of representing a metric space \mathcal{X} with the PT with embedding guarantees for any distinguished n -point dataset $\mathcal{X}_n \subseteq \mathcal{X}$ is illustrated in Figure 1.

1.4 Summary of Contributions

Our results can be summarized with reference to the PT’s embedding landscape in Figure 2. Our first two *universal representation theorems* relate to the first row, where no “geometric prior” is assumed on the metric space \mathcal{X} . The second and third rows relate to the sharper embedding guarantees we obtain when \mathcal{X} does have a regular geometry. In particular, this regularity allows us to explicitly bound the embedding’s effective dimension which becomes independent of n , and depends only on \mathcal{X} ’s latent geometry.

We begin by showing that, for any finite dataset \mathcal{X}_n in \mathcal{X} and any geometric perturbation parameter $\alpha \in (1/2, 1)$ (Hölder coefficient), there is a PT which bi- α -Hölder embeds \mathcal{X}_n into \mathcal{R} with metric distortion equal to that of the quantitative version of Assouad’s embedding theorem (Assouad, 1979) as derived by Naor and Neiman (Naor and Neiman, 2012). We note that the result of Naor and Neiman is an existence theorem, while we prove that the embedding can be implemented by a concrete deep neural network (the PT) with explicitly controlled width and depth. Of course, it is desirable to transition from $\alpha < 1$ to a bi-Lipschitz embedding with $\alpha = 1$. One reason is that Lipschitz feature maps are easiest to learn from (Kratsios and Papon, 2022). Our second main result shows that for any $D > 1$, still without geometric assumptions on \mathcal{X} , there exists a PT which represents \mathcal{X}_n in \mathcal{R} and for which a pair of uniformly randomly sampled points are bi-Lipschitz embedded with metric distortion at most D with probability at least $\mathcal{O}(n^{-4e/(1+D)})$. Thus the PT exhibits a natural trade-off between embedding quality and satisfiability of the embedding.

It is natural to expect that if \mathcal{X} has a regular geometric prior, it should be possible to obtain a deterministic small-distortion bi-Lipschitz guarantee for all points in \mathcal{X}_n . This is indeed the case: we show that if \mathcal{X} is a compact d -dimensional Riemannian manifold with controlled Ricci curvature, then there is a PT which bi-Lipschitz embeds \mathcal{X}_n into \mathcal{R} with effective dimension $3d + 6$. A similar fact holds for combinatorial trees, with PT’s metric distortion coinciding with that of Gupta’s embeddings of trees in Euclidean space (Gupta, 2000). In these cases, we also obtain explicit depth, width, and effective dimension estimates. It is worth emphasizing that the depth and the width only depend on the number of points n for which we seek a memorization guarantee (the number of landmarks) and the effective dimension only depends on the regularity and geometry of \mathcal{X} (cf. Figure 2). Lastly, we show that if one embeds into Gaussian mixtures on Euclidean space of dimension at least three, then there are probabilistic transformers which implement bi-Hölder embeddings of arbitrarily low distortion.

The significance of these results for machine learning practice is that $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ indeed acts as a “universal” embedding space with the feature map implemented by a probabilistic transformer. This motivates the development of downstream tools to work with data points that are probability measures, and in particular Gaussian mixtures. In fact, some of those tools are already available⁴, e.g. (Bigot et al., 2017; Mi et al., 2018; Ye et al., 2017; Ho et al., 2017; Chen et al., 2021). Furthermore, deep neural networks implicitly or explicitly compute representations of data, but, to the best of our knowledge, there are no prior theoretic results for such deep neural representations of embedding maps which guarantee low metric distortion. We thus provide the first theoretical results on embeddings of datasets from general metric spaces using deep neural networks.

Independently of how the embeddings are implemented (by probabilistic transformers or otherwise), these are also the first guarantees that the discrete geometries covered by our results can be embedded into the space of Gaussian mixtures on \mathbb{R} with an optimal transport-type distance. Finally, we develop new proof techniques that opportunely combine metric embedding theory (Heinonen, 2003; Krauthgamer et al., 2005; Naor and Neiman, 2012; Ostrovskii, 2013; Eriksson-Bique, 2018; Andoni et al., 2018), memorization theory of deep feedforward networks (Sontag, 1997b; Park et al., 2020; Feldman and Zhang, 2020; Daniely, 2020; Vershynin, 2020; Bubeck et al., 2020; Vardi et al., 2022), and computational optimal transport (Delon and Desolneux, 2020). Finally, we note that our results immediately imply embeddings into the usual Wasserstein-2 space over \mathbb{R} .

2. Geometric Background

We briefly overview some of the relevant terminologies from metric embedding theory. For further details we recommend the book (Ostrovskii, 2013) or the lecture notes (Matoušek, 2013; Naor, 2015).

A metric space is a pair $(\mathcal{X}, d_{\mathcal{X}})$ of a set \mathcal{X} and a *distance function* $d_{\mathcal{X}} : \mathcal{X} \times \mathcal{X} \rightarrow [0, \infty)$ which is symmetric in its arguments, is 0 if and only if $x = \tilde{x}$ (thus can uniquely identify points), and satisfies the triangle inequality: for every three points x_1, x_2, x_3 in \mathcal{X} $d_{\mathcal{X}}(x_1, x_3) \leq d_{\mathcal{X}}(x_1, x_2) + d_{\mathcal{X}}(x_2, x_3)$. The prototypical metric space is the Euclidean space; i.e. \mathbb{R}^n with Euclidean distance $d^2(x, \tilde{x}) \stackrel{\text{def.}}{=} \sum_{i=1}^n (x_i - \tilde{x}_i)^2$. We now review three other examples which are central to our analysis.

2.1 Combinatorial Graphs

A (*combinatorial*) *graph* is a pair $G \stackrel{\text{def.}}{=} (V, E)$ of a set of *vertices* (or vertices) V and a set of unordered pairs⁵ $E \subseteq \binom{V}{2}$, called *edges*, such that $\{u, v\} \in E$ if there is an edge between $u, v \in V$. We will always assume G to be *connected*, which means that for any $u, v \in V$ there is a sequence of edges $\{u, v_1\}, \dots, \{v_N, v\}$ in E “linking u to v ”. Every graph $G = (V, E)$ induces a metric space $\mathcal{X}_G \stackrel{\text{def.}}{=} (V, d_G)$ whose *graph geodesic* distance

4. All our results regarding embeddings into $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ yield embeddings into $(\mathcal{GM}_2(\mathbb{R}), \mathcal{W}_2)$ with the same distortion using the same probabilistic transformer network. Therefore, one can deploy any such available machine learning tool on $(\mathcal{W}_2(\mathbb{R}), \mathcal{W}_2)$.

5. For a set V , $\binom{V}{2}$ denotes the quotient of the Cartesian product $V \times V$ under the equivalence relation \sim defined on any $(v_1, v_2), (v_3, v_4) \in V \times V$ by $(v_1, v_2) \sim (v_3, v_4)$ if and only if $(v_1, v_2) = (v_3, v_4)$ or $(v_1, v_2) = (v_4, v_3)$.

d_G is defined for any $u, v \in V$ by

$$d_G(v, u) \stackrel{\text{def}}{=} d_G(u, v) \stackrel{\text{def}}{=} \inf \{N : \exists \{v, v_1\}, \dots, \{v_{N-1}, u\} \in E\}.$$

We only consider *simple graphs*, meaning that any two vertices can be connected by at most one edge.

2.2 Riemannian Manifolds with Lower-Bounded Ricci Curvature

We overview the “smooth geometric prior” considered in this paper; namely, we review the notion of a Riemannian manifold. For a more detailed exposition on Riemannian geometry, we refer the reader to (Jost, 2017, Chapter 1) on the topic.

A d -dimensional Riemannian manifold (M, g) is a pair of a d -dimensional *smooth manifold* M and smoothly varying family of positive-definite inner-products $g \stackrel{\text{def}}{=} (g_x)_{x \in M}$, with each g_x mapping pairs of vectors in the tangent space $T_x(M)$ above the point $x \in M$ to \mathbb{R} . Equivalently, by (Nash, 1954, Theorem 2) we may describe any d -dimensional Riemannian manifold “extrinsically” by viewing it as a Riemannian submanifold of the Euclidean space \mathbb{R}^{2d+1} .

We say that (M, g) is complete if every two points in M can be reached by a (continuous) curve and if the *geodesic distance*, defined for $x, \tilde{x} \in M$ by

$$d_g(x, \tilde{x}) \stackrel{\text{def}}{=} \inf_{\gamma} \int_0^1 \langle \gamma'(t), \gamma'(t) \rangle_{\gamma(t)} dt, \quad (1)$$

makes (M, d_g) into a complete metric space; where the infimum in (1) is taken over all piece-wise continuously differentiable curves $\gamma : [0, 1] \rightarrow M$ starting at x and ending at \tilde{x} ; i.e. $\gamma(0) = x$ and $\gamma(1) = \tilde{x}$. By the main result of (Hopf and Rinow, 1931), the completeness of (M, d_g) is equivalent to the existence of unique smooth curves minimizing the objective function (1); these curves are called *geodesics*.

In this paper, we focus on Riemannian manifolds satisfying a lower-bounded *Ricci curvature condition*, which we express using the *curvature-dimension inequality* of (Baudoin et al., 2014): for every infinitely differentiable map $f : M \rightarrow \mathbb{R}$ the following holds

$$\frac{1}{2} \Delta \|\nabla f\|_g^2 - \langle \nabla f, \nabla(\Delta f) \rangle_g \geq \frac{1}{d} \|\Delta f\|_g^2 + r \|\nabla f\|_g, \quad (2)$$

where Δ is the Laplacian on (M, g) (which describes heat flow on (M, g)) and $\|\cdot\|_g \stackrel{\text{def}}{=} \langle \cdot, \cdot \rangle_g$. To make matters concrete, one can show that a d -dimensional sphere of radius $\rho > 0$ satisfies (2) with $r = \rho^{-2}d(d-1) > 0$, the hyperbolic plane satisfies (2) with $r = -1 < 0$, and the Euclidean space satisfies (2) with $r = 0$. Since Δ dictates the heatflow on (M, g) , then r in (2) can be interpreted as describing the rate at which heat dissipates on (M, g) with small r being fast (e.g. the hyperbolic plane) and large r being slow (e.g. the sphere).

2.3 Mixed-Gaussian Optimal Transport

We rely on one more example of a non-Euclidean metric space, which is a variant of the space of probability measures introduced by (Wasserstein, 1969), in the context of optimal transport theory. This is the optimal transport-theoretic space of Gaussian mixtures introduced in (Delon and Desolneux, 2020), which specializes in the classical optimal transport constructions to the class of Gaussian mixtures. This metric space of probability measures is part of a broader area of contemporary research encoding structure into transport plans by restricting the classes of admissible transport plans; see (Beiglböck et al., 2017a; Bion-Nadal and Talay, 2019; Acciaio et al., 2020; Backhoff-Veraguas et al., 2020a; Backhoff-Veraguas and Pammer, 2022a)). In what follows, we use $\mathcal{P}_p(\mathbb{R}^d)$ to denote the set of probability measures on the Euclidean space \mathbb{R}^d admitting a p^{th} -moment; i.e. $\mathbb{P} \in \mathcal{P}_p(\mathbb{R}^d)$ if $\mathbb{E}_{X \sim \mathbb{P}}[\|X\|^p] < \infty$.

We refer to the metric space $\mathcal{W}_2(\mathbb{R}^d) \stackrel{\text{def}}{=} (\mathcal{P}_2(\mathbb{R}^d), \mathcal{W}_2)$ as the Wasserstein space, where \mathcal{W}_2 is defined for any two \mathbb{P} and \mathbb{Q} in $\mathcal{P}_2(\mathbb{R}^d)$ by

$$\mathcal{W}_2^2(\mathbb{P}, \mathbb{Q}) \stackrel{\text{def}}{=} \inf_{\pi \in \text{Cpl}(\mathbb{P}, \mathbb{Q})} \mathbb{E}_{(X, Y) \sim \pi} [\|X - Y\|^2], \quad (3)$$

where the set $\text{Cpl}(\mathbb{P}, \mathbb{Q})$ consists of all probability distributions on \mathbb{R}^{2d} with marginals \mathbb{P} and \mathbb{Q} . A key point, highlighting our interest in the case in embedding of metric spaces into univariate distributions, is that when $d = 1$ then

$$\mathcal{W}_2(\mathbb{P}, \mathbb{Q}) = \int_0^1 \left(Q_{\mathbb{P}}(t) - Q_{\mathbb{Q}}(t) \right)^2 dt,$$

where $Q_{\mathbb{P}}$ and $Q_{\mathbb{Q}}$ are the respective quantile functions of \mathbb{P} and of \mathbb{Q} . Thus, for empirical distributions, unlike in the multivariate case ($d > 1$) where computing $\mathcal{W}_2(\mathbb{P}, \mathbb{Q})$ has super-cubic complexity (Orlin, 1988), in the univariate case ($d = 1$) \mathcal{W}_2 can be easily computed via a simple sorting procedure; i.e. with near linear complexity.

In the case where \mathbb{P} and \mathbb{Q} are *Gaussian mixtures* (Delon and Desolneux, 2020, Section 4.1) propose a stronger distance function than the classical Wasserstein-2 distance $\mathcal{W}_2(\mathbb{P}, \mathbb{Q})$ on the set of subset of $\mathcal{P}_2(\mathbb{R})$ consisting of univariate Gaussian mixtures; we denote this set by $\mathcal{GM}_2(\mathbb{R})$. This distance is defined by restricting the couplings in optimization problem (3) to the smaller class couplings π which are themselves are Gaussian mixtures (on \mathbb{R}^{2d}). Denoted by $\mathcal{MW}_2(\mathbb{P}, \mathbb{Q})$, the mixed-Gaussian Wasserstein distance between two Gaussian mixtures $\mathbb{P} = \sum_{i=1}^I w_{1,i} N(\mu_{1,i}, \Sigma_{1,i})$ and $\mathbb{Q} = \sum_{j=1}^J w_{2,j} N(\mu_{2,j}, \Sigma_{2,j})$ is equivalently expressed in the following computationally tractable formulation. Denote the Wasserstein-2 distance between $N(\mu_{1,i}, \Sigma_{1,i})$ and $N(\mu_{2,j}, \Sigma_{2,j})$ by ω_{ij} and let $\Omega = (\omega_{ij}^2)_{i,j=1}^{I,J}$. Then $\mathcal{MW}_2^2(\mathbb{P}, \mathbb{Q})$ is given by

$$\begin{aligned} \mathcal{MW}_2^2(\mathbb{P}, \mathbb{Q}) &\stackrel{\text{def.}}{=} \min_{V \in [0, \infty)^{I \times J}} \text{tr}(V^{\top} \Omega) \\ &\text{subject to } V \mathbf{1}_J = w_1 \\ &\quad V^{\top} \mathbf{1}_I = w_2 \end{aligned}$$

where $\mathbf{1}_L = (1, \dots, 1)^{\top}$ denotes the L -dimensional vector of ones. Central to the tractability of \mathcal{MW}_2^2 is that, by (Cuesta and Matran, 1989), the Wasserstein-2 distance between two Gaussians $N(\mu_1, \Sigma_1)$ and $N(\mu_2, \Sigma_2)$ is given in closed form by

$$\mathcal{W}_2^2(N(\mu_1, \Sigma_1), N(\mu_2, \Sigma_2)) = \|\mu_1 - \mu_2\|^2 + \text{tr} \left(\Sigma_1 + \Sigma_2 - 2\sqrt{\sqrt{\Sigma_1}\Sigma_2\sqrt{\Sigma_1}} \right),$$

where \sqrt{A} denotes the square-root of a symmetric positive-definite matrix A .

There are several relationships between the two optimal transport distances. First, clearly by definition one has $\mathcal{W}_2 \leq \mathcal{MW}_2$; moreover, the inequality is often strict. Conversely, it holds that

$$\mathcal{MW}_2(\mathbb{P}, \mathbb{Q}) \leq \mathcal{W}_2(\mathbb{P}, \mathbb{Q}) + \sqrt{2} \left(\left(\sum_{i=1}^I w_{1,i} \text{tr}(\Sigma_{1,i}) \right)^{1/2} + \left(\sum_{j=1}^J w_{2,j} \text{tr}(\Sigma_{2,j}) \right)^{1/2} \right);$$

from which we see that \mathcal{MW}_2 equals to \mathcal{W}_2 for finitely supported measures. We use $(\mathcal{GM}(\mathbb{R}^d), \mathcal{MW}_2)$ to denote the space of Gaussian mixtures with mixed-Gaussian Wasserstein metric.

2.4 Bi-Hölder Embeddings

A map $\phi : \mathcal{X} \hookrightarrow \mathcal{R}$ from a metric space (\mathcal{X}, d) and to a (metric) representation space $(\mathcal{R}, d_{\mathcal{R}})$ is called a bi-Hölder embedding if: there is an $0 < \alpha \leq 1$, a *scale* $s > 0$, and a *distortion* $D \geq 1$ such that for every $x, \tilde{x} \in \mathcal{X}$

$$sd^{\alpha}(x, \tilde{x}) \leq d_{\mathcal{R}}(\phi(x_1), \phi(x_2)) \leq sD d^{\alpha}(x, \tilde{x}). \quad (4)$$

The smallest value of D for which (4) holds quantifies the worst-case dilation between any two points induced by the feature map ϕ , upon re-scaling⁶ by s . As an example, if $\mathcal{R} = \mathbb{R}^n$, $\mathcal{X} \subseteq \mathbb{R}^m$, and if ϕ is smooth then

6. If \mathcal{R} is a normed space and ϕ is invariant under linear maps then we can always absorb s into the model. However, for general representation spaces re-scaling is not meaningful without the scale s .

D is roughly $\max_x \|\nabla\phi(x)\|$; however, for general metric spaces these quantities are meaningless. Thus, a large distortion D only *globally* distorts \mathcal{X} 's geometry by stretching all of it. As illustrated in Figure 3, a small value of α relatively distorts the space's geometry by placing more importance (illustrated by bright colours) on nearby pairs of points and less importance on distant pairs of points (relative to the color redred vertex in the center of Figure 3). This has the effect that large scales become more minor and small scales become larger when $\alpha \ll 1$, with this effect disappearing as α approaches 0. Effectively α plays a comparable role to




Figure 3: Effect of varying α on the geometry of the embedding metric space. a) $\alpha = 1$ (Lipschitz): all distances are equally distorted. b) $\alpha \ll 1$ (fractional): small distances are prolonged, large distances are shortened.

tuning the *neighbourhood size* in classical iterations of graph attention (Veličković et al., 2018) or the role of neighborhood size when defining the *mean average precision* performance metric of (Cruceru et al., 2020; Giovanni et al., 2022).

The parameter $\alpha \in (0, 1]$ can also be interpreted through its role in the loss function, which we use to train our model. If \mathcal{X}_n is a set of points which we want to embed into $\mathcal{MW}(\mathbb{R})$ then in our numerical experiments, we will train a probabilistic transformer to learn a low distortion bi-Hölder embedding by numerically optimizing the following loss function which is a proxy for condition (4):

$$\sum_{x, \tilde{x} \in \mathcal{X}_n} |d_{\mathcal{X}}^{\alpha}(x, \tilde{x}) - \mathcal{MW}_2(\phi(x), \phi(\tilde{x}))|. \quad (5)$$

Here, the minimization is taken over the set of a set of models. When $\alpha \ll 1$ then, (5) over emphasizes embedding nearby pairs of points over embedding distant pairs of points. In contrast, when $\alpha = 1$, all pairs of points are given equal importance in the embedding. Thus, through the loss function (5), α implicitly plays the role of the neighborhood size defining the graph attention mechanism of (Veličković et al., 2018) or the average (distance) distortion loss function (used in e.g. (Samal et al., 2018; Giovanni et al., 2022)).

(Krauthgamer et al., 2005) introduce the notion of an aspect ratio of a measure space as the ratio of total mass over the minimum mass at any point. Similarly, we define the *aspect ratio* of a finite metric space (\mathcal{X}_n, d_n) as the ratio of the maximum distance between any two points therein (its diameter) over the minimum separation between any two distinct points. We define

$$\text{aspect}(\mathcal{X}_n, d_n) \stackrel{\text{def.}}{=} \frac{\max_{x, \tilde{x} \in \mathcal{X}_n} d_n(x, \tilde{x})}{\min_{x, \tilde{x} \in \mathcal{X}_n; x \neq \tilde{x}} d_n(x, \tilde{x})}.$$

We note that closely related concepts to the aspect ratio can be found in the (approximate) memorization results for deep feedforward networks (Park et al., 2021). Furthermore, requiring that the aspect ratio is positive and finite circumvents otherwise memorization by a feedforward neural network may be impossible (Sontag, 1997a).

A related quantity is the *diameter* of the n -point metric space (\mathcal{X}_n, d_n) , defined by $\text{diam}(\mathcal{X}_n, d_n) \stackrel{\text{def.}}{=} \max_{x, \tilde{x} \in \mathcal{X}_n} d_n(x, \tilde{x})$. Thus, the aspect ratio can be interpreted as (\mathcal{X}_n, d_n) 's diameter re-scaled by the reciprocal smallest distance in \mathcal{X}_n .




Figure 4: \mathcal{X} 's capacity is the number of small balls of radius $r/5$ that can fit in any ball of radius r .

Lastly, we will make use of the notion of “metric capacity” $\text{cap}(\mathcal{X}, d_{\mathcal{X}})$ of a metric space $(\mathcal{X}, d_{\mathcal{X}})$. This “dimensional” metric invariant is defined similarly to box or Hausdorff dimension (in fractal geometry (Falconer, 1986)), is a formalization of the notion of dimension of a general metric space \mathcal{X} . When \mathcal{X} is the m -dimensional Euclidean space or when m -dimensional compact Riemannian manifold then, the logarithm of its capacity can be shown to be proportional to m . Illustrated in Figure 4, the *capacity* $\text{cap}(\mathcal{X}, d_{\mathcal{X}})$ of a metric space $(\mathcal{X}, d_{\mathcal{X}})$ is equal to

$$\sup \{d \in \mathbb{N}_+ : \exists (x_i)_{i=0}^d \in \mathcal{X}^{d+1}, \exists r > 0 \text{ s.t. } \sqcup_{i=1}^d B_{\mathcal{X}}(x_i, r/5) \subseteq B_{\mathcal{X}}(x_0, r)\},$$

where $B_{\mathcal{X}}(x, r) \stackrel{\text{def.}}{=} \{u \in \mathcal{X} : d_{\mathcal{X}}(u, z) < r\}$ and \sqcup denotes the *union* of disjoint sets. We refer the reader to (Heinonen, 2001, Chapter 10) and (Bruè et al., 2021, Proposition 1.7) for further details on metric capacity and to (Le Donne and Rajala, 2015) for an exposition of metric-theoretic notions of dimension.

3. The Probabilistic Transformer Model

Since any good representation space should admit “nice” feature maps which can be approximated by standard neural networks. Our main result shows that this holds for $\mathcal{M}\mathcal{W}_2(\mathbb{R})$, with the neural network being the probabilistic transformer. The PT model, introduced below and illustrated in Figure 5, processes a dataset (\mathcal{X}, d_n) from a metric space $(\mathcal{X}, d_{\mathcal{X}})$ in three phases before producing a probabilistic output in $\mathcal{G}\mathcal{M}_2(\mathbb{R}^d)$.




Figure 5: The architecture of a probabilistic transformer.

First, data in (\mathcal{X}_n, d_n) is encoded into an L -dimensional vector of distances to a set of landmarks $x_1, \dots, x_L \in \mathcal{X}_n$. The encoding, illustrated by the $\Phi_{\{x_i\}_{i=1}^L}$ layer in Figure 5, is reminiscent of the graph attention mechanism of (Veličković et al., 2018). The *un-normalized graph attention* mechanism is defined as

$$\Phi_{\{x_i\}_{i=1}^L} : \mathcal{X}_n \ni x \mapsto (d_n(x, x_i))_{i=1}^L \in \mathbb{R}^L. \quad (6)$$

The landmarks act as an intrinsic reference frame, they identify points in \mathcal{X}_n by their relative positions to $\{x_i\}_{i=1}^L$. This can be seen as a non-linear counterpart to a truncated orthonormal basis of a Hilbert space. When all points in \mathcal{X} are chosen as landmarks, this un-normalized graph attention mechanism is a finite-dimensional version of the embedding of (\mathcal{X}_n, d_n) into $(\mathbb{R}^n, \ell^\infty)$ (Fréchet, 1906).

Remark 1 *The points in x_1, \dots, x_L in the definition of $\Phi_{\{x_i\}_{i=1}^L}$ are in principle trainable and could be discovered algorithmically instead of being specified by the user.*

Next, this vectorial data is processed by a set of *independent, parallel* ReLU-feedforward networks $\mathcal{N}\mathcal{N}_w : \mathbb{R}^L \rightarrow \mathbb{R}^K$, and K ReLU-feedforward networks $\mathcal{N}\mathcal{N}_1 : \mathbb{R}^L \rightarrow \mathbb{R}^d \times \mathbb{R}^{d \times d}, \dots, \mathcal{N}\mathcal{N}_K : \mathbb{R}^L \rightarrow \mathbb{R}^d \times \mathbb{R}^{d \times d}$. The role of the networks $\{\mathcal{N}\mathcal{N}_k\}_{k=1}^K$ is to process inputs to a d -dimensional Gaussian measure-valued outputs

$N_d(\mu, \Sigma)$ by outputting a mean vector μ and a square-root of a covariance matrix $\Sigma^\top \Sigma$. The role of the network \mathcal{NN}_w is to generate input-dependent weights in the K -simplex $\{w \in [0, 1]^K : \sum_{k=1}^K w_k = 1\}$ used to mix the Gaussians into a Gaussian mixture.

Remark 2 *It is theoretically sufficient for the ReLU network to be independent. In practice, this can be relaxed by merging them into a single network as we do in Section 5.*

The predictions of all networks are then combined using the probabilistic attention mechanism (Kratsios et al., 2022) which generalizes the construction of (Bishop, 1994). The (Gaussian) probabilistic attention mechanism is defined for any $K \in \mathbb{N}_+$, any weight $w \in \mathbb{R}^K$, and any set of keys $Y \stackrel{\text{def.}}{=} ((\mu_1, \Sigma_1), \dots, (\mu_K, \Sigma_K)) \in \mathbb{R}^{K \times (d+d^2)}$ by

$$\text{P-Attn}(w, Y) \stackrel{\text{def.}}{=} \sum_{k=1}^K \sigma_K(w)_k N(\mu_k, \Sigma_k^\top \Sigma_k) \in \mathcal{P}_2(\mathbb{R}^d);$$

where we have identified \mathbb{R}^{d+d^2} with $\mathbb{R}^d \times \mathbb{R}^{d \times d}$ (where $\mathbb{R}^{N \times M}$ is the set of $N \times M$ matrices) and $\sigma_K(u) \stackrel{\text{def.}}{=} (e^{u_k} / \sum_{i=1}^K e^{u_i})_{k=1}^K$ is the *softmax function*. The probabilistic attention mechanism’s place in our model is illustrated by the last layer in Figure 6a, which synthesizes the mixtures from the $(\mu_k)_{k=1}^K$, $(\Sigma_k)_{k=1}^K$, and the $(w_k)_{k=1}^K$ parameters produced by the feedforward networks.

The parameters defining the Gaussian mixtures are implemented by deep feedforward networks with ReLU activation function $\text{ReLU}(x) \stackrel{\text{def.}}{=} \max\{0, x\}$. Let $D \in \mathbb{N}_+$. Fix a *depth*⁷ $J \in \mathbb{N}_+$ and a multi-index $[d] \stackrel{\text{def.}}{=} (d_0, \dots, d_J)$ of integers with $d_0 = d$ and $d_J = D$. A function $\mathcal{NN} : \mathbb{R}^d \rightarrow \mathbb{R}^D$ is said to be a deep feedforward network if for every $j = 1, \dots, J$ there are $d_j \times d_{j-1}$ -dimensional matrices $A^{(j)}$ called weights and $b^{(j)} \in \mathbb{R}^{d_j}$ called biases, such that \mathcal{NN} admits the iterative representation

$$\begin{aligned} \mathcal{NN}(x) &\stackrel{\text{def.}}{=} A^{(J)} x^{(J)} + b^{(J)} \\ x^{(j)} &\stackrel{\text{def.}}{=} \text{ReLU} \bullet (A^{(j)} x^{(j-1)} + b^{(j)}) \quad \text{for } j = 1, \dots, J-1, \\ x^{(0)} &\stackrel{\text{def.}}{=} x, \end{aligned} \tag{7}$$

where \bullet denotes component-wise composition. We now formalize the probabilistic transformer in Figure 6a.

Definition 3 (Probabilistic Transformer (PT)) *Fix a metric space (\mathcal{X}, d) and a $D \in \mathbb{N}_+$. A probabilistic transformer is a map $T : \mathcal{X} \rightarrow \mathcal{P}_2(\mathbb{R}^D)$ with representation*

$$T(x) = \text{P-Attn} \left(\mathcal{NN}_w(u), (\mathcal{NN}_k(u))_{k=1}^K \right), \quad u \stackrel{\text{def.}}{=} \Phi_{\{x_l\}_{l=1}^L}(x), \tag{8}$$

where $L \in \mathbb{N}_+$, $x_1, \dots, x_L \in \mathcal{X}$, and $\mathcal{NN}_w : \mathbb{R}^L \rightarrow \mathbb{R}^K$ and $\mathcal{NN}_1 : \mathbb{R}^L \mapsto \mathbb{R}^D \times \mathbb{R}^{D \times D}, \dots, \mathcal{NN}_K : \mathbb{R}^L \mapsto \mathbb{R}^D \times \mathbb{R}^{D \times D}$ are ReLU feedforward networks.

The *depth* of a probabilistic transformer T (with representation (8)) is defined as the maximum of the depth of each of the feedforward networks $\mathcal{NN}_w, \mathcal{NN}_1, \dots, \mathcal{NN}_K$. Similarly, the *width* of a transformer T is defined as the maximum of the widths of each feedforward network $\mathcal{NN}_w, \mathcal{NN}_1, \dots, \mathcal{NN}_K$. These definitions are natural since each of the networks $\mathcal{NN}_w, \mathcal{NN}_1, \dots, \mathcal{NN}_K$ defining T are independent, in the sense that their layers do not pass inputs to or from one another and their parameters have no explicit relations.

The *effective dimension* of a probabilistic transformer T , is the maximum number of mixtures output by T for any input scaled by the number of parameters defining each Gaussian probability measure in

$$\text{effdim}(T) := K(D + D^2),$$

where we use the notation of Definition 3. Thus, $d_T \stackrel{\text{def.}}{=} K$. We define the number of trainable parameters in T , denoted by $\text{par}(T)$, as the aggregate number of *trainable parameters* amongst the networks

7. The depth is the total number of affine layers defining the network and not the number of “hidden layers”; that is, the depth is one more than the number of ReLU activation functions applied component-wise in (7).

$\mathcal{NN}_w, \mathcal{NN}_1, \dots, \mathcal{NN}_K$; i.e. $\text{par}(T) \stackrel{\text{def.}}{=} \text{par}(\mathcal{NN}_w) + \sum_{k=1}^K \text{par}(\mathcal{NN}_k)$. As in (Bartlett et al., 2019), the number of trainable parameters encodes “network topology” by only counting the number of non-zero entries, thus only counting the number of weighted and biases between any two neurons with a direct connection. In the notation of (7), the number of trainable parameters of a neural network \mathcal{NN} is defined as

$$\text{par}(\mathcal{NN}) \stackrel{\text{def.}}{=} \|c_0\|_0 + \sum_{j=1}^J (\|A^{(j)}\|_0 + \|b^{(j)}\|_0),$$

where $\|\cdot\|_0$ counts the number of non-zero entries in a matrix or vector of a given size. For fully connected feedforward networks of constant width $\text{par}(\mathcal{NN})$ is roughly equal to the network’s depth multiplied by its width squared; but in general $\text{par}(\mathcal{NN})$ will be much smaller since our networks will tend to have a *sparsely connected* architecture analogously to (Suzuki, 2019) or in the convolutional neural network theory (Zhou, 2020). If one replaces $\|\cdot\|_0$ with the Fröbenius and Euclidean norms (for the weight and biases respectively) then, $\text{par}(\mathcal{NN})$ becomes a path norm (Neysshabur et al., 2015; Zheng et al., 2019; E. and Wojtowytsch, 2022) typically used in neural network regularization.

4. Main Results

Our *quantitative* results concerning the embedding capabilities of probabilistic transformers in mixed-Gaussian Wasserstein spaces fall into two classes. The first set of results concerns the possibility of embedding and the complexity of the model performing an embedding when no “latent geometric structures” are present in the dataset, and the latter class of results concerns the embedding of datasets with a latent geometry.

Table 1 summarizes the rates in our main results. We emphasize that explicit constants are also available and we refer the reader interested in detailed constants, and precise \mathcal{O} rates instead of $\tilde{\mathcal{O}}$ to Tables 3 and 4 in Appendix A. Note that we use the following asymptotic notation. Given $g, f : \mathbb{N} \rightarrow [0, \infty)$ will say that g is $\tilde{\mathcal{O}}(f)$ if g is $\mathcal{O}(f)$ where $f = \tilde{f} \times$ logarithmic terms.

Table 1: Complexity of the probabilistic transformers performing n -point embeddings.

Latent Geometry	Effective Dimension d_T	Depth $\in \tilde{\mathcal{O}}(\cdot)$	Width $\in \Omega(\cdot)$	Result
General	$\Omega\left(\frac{\log(\text{cap}(\mathcal{X}_n, d_n))}{\alpha}\right)$	$n(1 + \log(n^{5/2} \text{aspect}(\mathcal{X}_n, d_n)))$	$\max\{d_T, n^2\}$	Theorem 4
General	$\mathcal{O}((D-2)^{-2} \theta_D \log_2(n))$	$n^{\Theta_D} (1 + \theta_D + \log(\text{aspect}(\mathcal{X}_n, d_n)))$	$\max\{d_T, n^{2\theta_D}\}$	Theorem 5
General - Multivariate Mixtures	$\mathcal{O}\left(\frac{n^6}{(D-1)} \text{aspect}(\mathcal{X}_n, d)^2\right)$	$n^2 + \log(n^{5/2} \text{aspect}(\mathcal{X}_n, d_n))$	$\mathcal{O}\left(\frac{n^5(n-1)}{(D-1)} \text{aspect}(\mathcal{X}_n, d)^2\right)$	Theorem 7
Discrete: Trees	M	$n(1 + \log(n^{5/2} \text{diam}(\mathcal{X}_n, d_n)))$	$\max\{M, n^2\}$	Proposition 11
Discrete: 2-Hop Graphs	$\Omega\left(\frac{\log(1+\rho(A_G))}{\alpha}\right)$	$n(1 + \log(n^{5/2} \text{diam}(\mathcal{X}_n, d_n)))$	$\max\{d_T, n^2\}$	Corollary 12
Manifold: Riemannian, Bounded Curvature	$\Omega\left(\frac{m^{1+\alpha}}{\alpha(1-\alpha)^{1+\alpha}}\right)$	$n(1 + \log(n^{5/2} \text{aspect}(\mathcal{X}_n, d_n)))$	$\max\{d_T, n^2\}$	Corollary 9
Manifold: Riemannian and Compact	$\Omega(d)$	$n(1 + \log(n^{5/2} \text{aspect}(\mathcal{X}_n, d_n)))$	$\max\{d, n^2\}$	Proposition 8

4.1 Embedding Guarantees for General Datasets

Our first main result shows that any metric space \mathcal{X} can be represented in $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ with a small PT guaranteeing that any fixed finite set of points can be embedded with little distortion, if we allow for minor perturbations to \mathcal{X} ’s geometry. By a “small PT”, we mean that the representation implemented by the PT of Theorem 5 does not face the *the curse of dimensionality*.

Theorem 4 (Deterministic Fractional Embeddings of Large Data by Small Transformers)

There is an absolute constants $C > 0$ such that, for every metric space $(\mathcal{X}, d_{\mathcal{X}})$, every finite subset $\mathcal{X}_n \subseteq \mathcal{X}$ with $n \geq 2$ points, and every “geometric perturbation parameter” $\frac{1}{2} < \alpha < 1$, there exists a probabilistic transformer T for which: for every $x, \tilde{x} \in \mathcal{X}$

$$d_{\mathcal{X}}^{\alpha}(x, \tilde{x}) \leq \mathcal{W}_2(T(x), T(\tilde{x})) \leq \mathcal{MW}_2(T(x), T(\tilde{x})) \leq \lceil CC_{\mathcal{X}_n} \rceil d_{\mathcal{X}}^{\alpha}(x, \tilde{x}),$$

where $C_{\mathcal{X}_n} \stackrel{\text{def.}}{=} \left(\frac{12 \log(\text{cap}(\mathcal{X}_n))}{(1-\alpha)} \right)^{1+\alpha}$. Furthermore, T 's effective dimension, depth, and width are recorded in Table 1 (with explicit constant given in Tables 3 and 4).

Independently of any machine learning model, results such as (Bartal et al., 2005) confirm there are n -point metric spaces \mathcal{X}_n with the property that any finite subset $\tilde{\mathcal{X}}_n$ that bi-Lipschitz embeds into ℓ^2 with distortion D cannot contain more than $n^{1-\vartheta_D}$ points, for some $\vartheta_D > 0$ depending only on D . Equivalently, if we uniformly and independently pick a pair of points in \mathcal{X}_n at random then they belong to $\tilde{\mathcal{X}}_n$ with probability at-most $\frac{1}{n^{2\vartheta_D}}$. This places a fundamental limitation on the number of pairs of points which can be bi-Lipschitz embedded with distortion D using any model. In particular, this type of limitation must translate to the PT model. In contrast, this is not the case, when working with almost bi-Lipschitz (α -Hölder with $\alpha \approx 1$) embeddings.

Nevertheless, the next result provides a probabilistic guarantee that a PT embeds pairs of points in the set of landmarks \mathcal{X}_n with a given distortion. As we may expect, the likelihood that any two given data points are embedded with a prescribed (maximal) distortion is proportional to the size of the distortion. In other words if we do not allow any perturbation to the landmark set's geometry then, with high probability, most pairs of points can be bi-Lipschitz embedded into $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ with large distortion and with small probability most pairs of points in the landmark set can be embedded with small distortion.

Theorem 5 (PAC-Type Embeddings of Large Data by Small Transformers) *Let $(\mathcal{X}, d_{\mathcal{X}})$ be a metric space, fix a finite subset \mathcal{X}_n of \mathcal{X} with $n \geq 2$ points, and let \mathbb{P} be the uniform probability measure on \mathcal{X}_n . There is a $\frac{1}{e^2} < \delta_n < 1$ depending only on n such that for every $\delta_n < \delta < 1$, there exists a “scale” $s > 0$ and a probabilistic transformer $T : (\mathcal{X}_n, d_n) \hookrightarrow (\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ satisfying*

$$\mathbb{P} \left(sd_{\mathcal{X}}(x, \tilde{x}) \leq \mathcal{W}_2(T(x), T(\tilde{x})) \text{ and } \mathcal{MW}_2(T(x), T(\tilde{x})) \leq sD(D-1)d_{\mathcal{X}}(x, \tilde{x}) \right) \geq \delta,$$

Abbreviate $\lambda \stackrel{\text{def.}}{=} \log_n(\sqrt{\delta})$. The distortion parameter $D > 2$ is given explicitly by

$$D = -2 \frac{(1 + \lambda)^{\frac{1+\lambda}{\lambda}}}{\lambda}$$

Furthermore, T 's effective dimension, depth, and width are recorded in Table 1 (explicit constants are in Tables 3 and 4).

Remark 6 (The Minimum Valid Probability δ_n in Theorem 5) *The quantity δ_n in Theorem 5 is defined by $\delta_n \stackrel{\text{def.}}{=} \max \left\{ \frac{1}{e^2}, \tilde{\delta}_n \right\}$ where $\tilde{\delta}_n$ is the unique $\frac{1}{e^2} \leq \delta < 1$ solving $\log_n(\sqrt{\delta}) = (1 + \log_n(\sqrt{\delta})) \frac{1 + \log_n(\sqrt{\delta})}{\log_n(\sqrt{\delta})}$.*

Even though the multivariate transport distances are more expensive to compute, it is natural to ask whether there are any advantages in representing metric spaces as mixtures of multivariate instead of univariate Gaussians. The answer is yes, as shown by the following deep neural analogue of the main finding of (Andoni et al., 2018). The result shows that if one considers multivariate Gaussian mixtures, there are probabilistic transformers that implement metric embeddings of arbitrarily low distortion on any finite subset of a given metric space; however, the effective dimension of these embeddings does become large. We note that it still depends only polynomially on the number of points being embedded.

Theorem 7 (Distortionless Fractional Embeddings into Multivariate Gaussian Mixtures) *Let (\mathcal{X}, d) be a metric space and $\mathcal{X}_n \subseteq \mathcal{X}$ a finite subset with at least two points. For every distortion $D > 1$ there exists a probabilistic transformer $T : \mathcal{X} \rightarrow \mathcal{GM}_2(\mathbb{R}^3)$ for which there is a “scale” $s > 0$ such that for every $x, \tilde{x} \in \mathcal{X}_n$*

$$s d^{1/2}(x, \tilde{x}) \leq \mathcal{MW}_2(T(x), T(\tilde{x})) \leq D s d^{1/2}(x, \tilde{x}). \quad (9)$$

Moreover, the complexity of T is given in Table 1.

Theorem 7 shows that one can achieve arbitrarily low distortion with embeddings into the space of multivariate Gaussian mixtures. In the more computationally efficient setting, where we consider embeddings into univariate Gaussian mixtures, Theorems 4 and 5 provide efficient embedding guarantees for arbitrary finite subsets of general metric spaces with small but non-vanishing distortion. It is natural to ask whether we can obtain stronger guarantees if \mathcal{X}_n is a subset of a well-behaved \mathcal{X} with a “geometric prior”. Indeed, we show below that in this case the guarantees in Theorem 1 can be improved to deterministic bi-Lipschitz guarantees.

4.2 Consequences: Improved Guarantees When (\mathcal{X}_n, d_n) has a Geometric Priors

Theorems 5 and 4 can be straightforwardly applied to any finite metric space, in the sense that there is no (\mathcal{X}_n, d_n) quantities to “plug-in”, and this is true regardless of if d_n encoded some additional “latent geometric priors”. Nevertheless, a fully explicit use of Theorem 4 requires an explicit estimate on (\mathcal{X}_n, d_n) ’s metric capacity. Conveniently, such estimates can be derived when (\mathcal{X}_n, d_n) has some latent geometry.

We consider two different types of geometries from which the n -point set of landmarks is drawn. For instance, one may alternatively interpret \mathcal{X}_n as a training dataset. The first is the smooth geometries arising from suitable Riemannian manifolds and the second is discrete geometries arising from well-behaved types of combinatorial graphs.

4.2.1 SMOOTH GEOMETRIC PRIORS

In the case of a “smooth geometric prior”, meaning that points are drawn from a Riemannian manifold with suitable curvature, the required number of mixtures and the embedding distortion can be made independent of the number of datapoints n . A fortiori, the necessary number of mixtures is proportional to the latent Riemannian manifold’s dimension.

Proposition 8 (Representation of Riemannian Manifolds with bi-Lipschitz Guarantees) *Fix $d \in \mathbb{N}_+$, let (M, g) be a d -dimensional Riemannian manifold with geodesic distance d_g , and let $K \subseteq M$ be compact. For any finite subset $\mathcal{X}_n \subseteq K$ with $n \geq 2$ points, there is a constant $C_K > 0$ depending only on K and a PT as in Theorem 4 satisfying: for every $x, \tilde{x} \in \mathcal{X}$ the following holds*

$$d_g(x, \tilde{x}) \leq \mathcal{W}_2(\varphi(x), \varphi(\tilde{x})) \leq \mathcal{M}\mathcal{W}_2(\varphi(x), \varphi(\tilde{x})) \leq C_K d_g(x, \tilde{x}).$$

Furthermore, T ’s effective dimension, depth, and width are recorded⁸ in Table 1.

Proposition 8 provides bi-Lipschitz embedding guarantees for finite subsets of a general Riemannian manifold via a small PT. If the latent Riemannian manifold has bounded Ricci curvature, then we find that the constant C_K can be made explicit and independent of any fixed compact subset of the Riemannian manifold. For this, we apply Theorem 4, and we turn to representations by our PT model with n -point (non-Lipschitz) bi-Hölder embedding guarantees.

Corollary 9 (Representation of Riemannian Manifolds with Bounded Ricci Curvature) *Fix $d \in \mathbb{N}_+$, let (M, g) be a complete m -dimensional Riemannian manifold lower-bounded Ricci curvature and with geodesic distance d_g . There is an absolute constant $\tilde{C} > 0$ such that, for every finite subset $\mathcal{X}_n \subseteq M$ with $n \geq 2$ points, there is a PT as in Theorem 4 satisfying*

$$d_g^\alpha(x, \tilde{x}) \leq \mathcal{W}_2(T(x), T(\tilde{x})) \leq \mathcal{M}\mathcal{W}_2(T(x), T(\tilde{x})) \leq \left\lceil \tilde{C} \left(\frac{m}{1-\alpha} \right)^{1+\alpha} \right\rceil d_g^\alpha(x, \tilde{x}).$$

Furthermore, T ’s effective dimension, depth, and width are recorded in Table 1 (with explicit constant given in Tables 3 and 4).

8. Explicit constant are recorded in Tables 3 and 4.

As a concrete example of Corollary 9 we obtain the following embedding guarantee for spherical data. It is worth noting that the sphere is both a simple yet interesting example since even in a Euclidean representation space, the result of (Robinson, 2006a) implies that the sphere cannot be embedded into any Euclidean space with no distortion⁹.

Example 1 (Embedding Spherical Data) Consider the m -dimensional sphere $S^m \stackrel{\text{def.}}{=} \{x \in \mathbb{R}^{m+1} : \|x\| = 1\}$ with Riemannian geometry inherited from its inclusion into the Euclidean space \mathbb{R}^{m+1} . In this geometry, the distance between any two $x_1, x_2 \in S^m$ is $d_g(x_1, x_2) = \arccos(x_1^\top x_2)$ and S^m has lower-bounded Ricci curvature (see (Jost, 2017, page 276)). Therefore, for every n -point subset $\mathcal{X}_n \subseteq S^m$ ($n \in \mathbb{N}_+$) there exists a PTT with width, depth, and number of mixture components as in Corollary 9 satisfying: for every $x, \tilde{x} \in S^m$ it holds that

$$d_g^{3/4}(x, \tilde{x}) \leq \mathcal{M}W_2(T(x), T(\tilde{x})) \leq \left[C \binom{4m}{1}^{1+3/4} \right] d_g^{3/4}(x, \tilde{x}),$$

and $C > 0$ is independent of n and of m .

4.2.2 COMBINATORIAL PRIORS

A particularly useful class of combinatorial graphs are *trees*; i.e. graphs in which no path originating at one vertex can ever come back to itself without passing the same edge at least twice. The simplicity of the combinatorial geometry of trees begs the question: can the transition between the deterministic embeddings of Theorem 4 (with geometric perturbation parameter $\alpha < 1$) and the probabilistic embeddings of Theorem 5 (with geometric perturbation parameter $\alpha = 1$) be avoided, and can the simple geometry of combinatorial tree allows us to obtain purely deterministic embedding results? The following result gives an affirmative answer. Furthermore, an embedding with $\mathcal{O}(n^{1/(M-1)})$ distortion is possible by using a probabilistic transformer with at most $\mathcal{O}(M)$ mixture components.

Remark 10 (Simplified Result Formulation for Graphs) For simplicity, we consider the case where the graph in question is finite and where each point in \mathcal{X} is a landmark. However, one can easily see how the statement extends to the general case considered in Theorem 4.

Proposition 11 (Bi-Lipschitz Representations of Combinatorial Trees) Let $G = (V, E)$ be an n -vertex tree and let (V, d_G) be its associated metric space (as in Example 2.1). There is an absolute constant $C > 0$ such that, for any positive integer M where $M, n \geq 2$, there exists a PTT with width and depth as in Theorem 4 and such that

$$d_G(x, \tilde{x}) \leq W_2(T(x), T(\tilde{x})) \leq \mathcal{M}W_2(T(x), T(\tilde{x})) \leq Cn^{1/(M-1)}d_G(x, \tilde{x}).$$

Furthermore, T 's effective dimension, depth, and width are recorded in Table 1 (with explicit constant given in Tables 3 and 4).

Therefore, for combinatorial graphs with simple geometric priors, the deterministic to probabilistic transition between Theorem 4 and 5 can be avoided. Next, we consider where we wish to represent data with manifold priors using the probabilistic transformer architecture.

Proposition 11 was derived using the same method of proof as for Theorem 4, but specialized to the geometry of combinatorial trees. Instead, the next result is a consequence of Theorem 4.

Many well-studied graphs are such that each vertex is reachable by hopping over at most one other. We refer to these types of graphs as “2-hop graphs”. Examples of 2-hop combinatorial graphs include, friendship graphs (Erdős et al., 1966), complete bipartite graphs as in Mantel’s Theorem (Erdős, 1964/65), cocktail-party graphs (Jungerman and Ringel, 1978), wheel graphs (Buckley and Harary, 1988), and various others.

To state our guarantees for 2-hop graphs we require some additional terminology. The adjacency matrix A_G of a graph $G = (V, E)$, whose vertices we enumerate by $V = \{v_i\}_{i=1}^n$, is the $n \times n$ -matrix with binary

⁹. We note this, to clarify the misconception that the “embeddings” of (Nash, 1954) are not in the metric sense (as considered in this article) but are in a different (Riemannian sense).

with $(A_G)_{i,j} = 1$ if and only if there is an edge between the vertices v_i and v_j . The spectral radius of A_G , denoted by $\rho(A_G)$, is the largest absolute eigenvalue of A_G .

Corollary 12 (Representation for 2-Hop Combinatorial Graphs) *Let $G = (V, E)$ be an n -vertex 2-hop graph and let (V, d_G) be its associated metric space (as in Example 2.1). Fix $\frac{1}{2} < \alpha < 1$, and let T be the PT of Theorem 4. Each $T(x)$ is comprised of no more than $\lceil 12C\alpha^{-1} \log(1 + \rho(A_G)) \rceil$ mixture components and satisfies: for each $x, \tilde{x} \in \mathcal{X}_n$ the following holds*

$$d_G^\alpha(x, \tilde{x}) \leq \mathcal{W}_2(T(x), T(\tilde{x})) \leq \mathcal{MW}_2(T(x), T(\tilde{x})) \leq \left[C \left(\frac{12 \log(1 + \rho(A_G))}{1 - \alpha} \right)^{1+\alpha} \right] d_G^\alpha(x, \tilde{x}).$$

Furthermore, T 's effective dimension, depth, and width are recorded in Table 1 (with explicit constant given in Tables 3 and 4).

The embedding guarantee in Corollary 12 can be further approximated using additional information about G 's connectivity. Specifically, define G 's *maximum degree* to be the maximum number of edges connected to any of its vertices, denoted by $\deg^+(G)$, and its *minimum degree* to be the smallest number of edges connected to any of its vertices, denoted by $\deg^-(G)$.

We use the following example to showcase the explicit constants and PT complexity estimates which are derived through our analysis (also recorded Appendix A).

Example 2 (Embedding of 2-Hop Graphs in Terms of Connectivity Information) *Consider the setting of Corollary 12, with $\alpha = \frac{3}{4}$, $n \geq 12$, and let T be the probabilistic transformer described by that result. The upper-bound on the spectral radius $\rho(A_G)$ given in (Das and Kumar, 2004, Theorem 2.7) implies the following embedding guarantee*

$$d_G^{3/4}(x, \tilde{x}) \leq \mathcal{MW}_2(T(x), T(\tilde{x})) \leq \lceil CC_G \rceil d_G^{3/4}(x, \tilde{x}),$$

where $C_G = \left(48 \log \left(1 + \sqrt{2(\#E)(n-1) \deg^+(G) + (\deg^+(G) - 1) \deg^-(G)} \right) \right)^{7/4}$. Furthermore, the ‘‘complexity’’ of T can be estimated in terms of G 's degree, its number of edges, and its diameter; namely,

$$\begin{aligned} \text{effdim}(T) &= \mathcal{O} \left(\log \left(1 + \sqrt{2(\#E)(n-1) \deg^+(G) + (\deg^+(G) - 1) \deg^-(G)} \right) \right) \\ \text{depth}(T) &= \mathcal{O} \left(n \left\{ 1 + \sqrt{n \log(n)} \left[1 + \frac{\log(2)}{\log(n)} \left(c + \frac{\log(n^{5/2} \text{diam}(\mathcal{X}_n, d_n))}{\log(2)} \right) \right] \right\} \right). \end{aligned}$$

We conclude our theoretical analysis by illustrating the fundamental limits of what can be achieved when embedding *arbitrary* finite metric spaces. Our results highlight the limitations of our model when globally embedding an (\mathcal{X}_n, d_n) with no geometric prior into $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$, as well as the mathematical limitations of any model embedding such a space into any smooth and finite-dimensional representation space.

4.3 Limitations to Global Embedding Capacity for Datasets with no Geometric Priors

It is natural to ask if the transition between the probabilistic bi-Lipschitz embedding of Theorem 5 and the deterministic bi-Hölder embedding of Theorem 4 is just an artifact of our analysis. The answer is no. The next result confirms as the phenomenon is generic and persists for any ‘‘regular’’ finite-dimensional embedding.

We begin by considering the case of complete Riemannian manifold \mathcal{R} of bounded negative curvature for which the frequently used hyperbolic spaces in non-Euclidean representation learning are prototypical. Our next result combines the results of (Naor et al., 2006) and (Biggs and Hoare, 1983) to construct an explicit sequence of bipartite combinatorial graphs (namely the sextet graphs introduced by (Biggs and Hoare, 1983)) which cannot be embedded into \mathcal{R} with non-diverging distortion.

Proposition 13 (Impossibility of Manifold Embedding of Fixed Finite Dimension) *Let \mathcal{R} be a complete Riemannian manifold with negative sectional curvature bounded in $[-C, -c]$ where $0 < c \leq C$. Then there exists a constant $c_{\mathcal{R}} > 0$ depending only on \mathcal{R} such that, for every prime integer $n \geq 2$ satisfying $n \bmod 16 \in \{\pm 3, \pm 5, \pm 7\}$ there is a combinatorial graph $(\mathcal{X}, d_{\mathcal{X}})$ with n -vertices and degree exactly 3 such that every bi-Lipschitz embedding $f : \mathcal{X} \rightarrow \mathcal{R}$ must have distortion D bounded below by*

$$D \geq c_{\mathcal{R}} \sqrt{\frac{4}{3} \log_2(n) - 2}. \quad (10)$$

Remark 14 (Improved Lower Bound Via Expander Graphs) *One can explicitly obtain larger lower bounds than in 10 by using the result of (Margulis, 1982, 1988) (both discovered independently and improving a result of (Erdős and Sachs, 1963)) stating that: for every $\delta \geq 3$ there is a diverging sequence $\{n_k\}_{k \in \mathbb{N}}$ and graphs $\{G_{n_k}\}_{n \in \mathbb{N}}$ each with n_k -vertices such that each G_{n_k} has girth at-least $\frac{3}{4} \log_{\delta-1}(n_k)$. Using this sequence of expander graphs, instead of the sequence of Sxetel graphs used in the proof of Proposition 13, implies that there is no bi-Lipschitz embedding of G_{n_k} into \mathcal{R} with distortion less than*

$$\mathcal{O}\left((\delta - 2) \sqrt{\frac{4}{3} \log_{\delta-1}(n) + 2}\right).$$

The advantage of the construction in Proposition 13 is a simple explicit class of graphs.

We complement Proposition 13 by considering the case of compact Riemannian manifolds or Euclidean spaces. We again find that there is a sequence of pathological finite metric spaces which cannot be bi-Lipschitz embedded into such representation spaces with non-diverging distortion.

Proposition 15 (Impossibility of Manifold Embedding of Fixed Finite Dimension) *Let \mathcal{R} be a Riemannian manifold which is bi-Lipschitz embedded in the Hilbert space ℓ^2 . There is a constant $c_{\mathcal{R}} > 0$ (depending only on \mathcal{R}) such that for every positive integer n , there is an n -point metric space $(\mathcal{X}, d_{\mathcal{X}})$ for which every bi-Lipschitz embedded $f : \mathcal{X} \rightarrow \mathcal{R}$ has distortion D satisfying*

$$D \geq c_{\mathcal{R}} \frac{\log(n)}{\log(\log(n))}.$$

Example 3 (Compact Manifolds Satisfy Proposition 15) *If \mathcal{R} is a compact Riemannian manifold then Whitney's embedding theorem implies that there is a smooth diffeomorphism (onto its image) $\varphi : \mathcal{R} \rightarrow \mathbb{R}^{2d}$. Therefore, Rademacher's theorem and the compactness of \mathcal{R} implies that φ is a bi-Lipschitz embedding. Thus, \mathcal{R} satisfies the conditions of Proposition 15 by the Whitney embedding theorem.*

The PAC-type bound of Theorem 5 can be reformulated as a function of the distortion D (instead of the probability δ). Doing so allows us to compare our bi-Lipschitz embedding results with the above two impossibility results. We now record this inverted formulation.

Remark 16 (Inverted Form of Theorem 5 where Distortion is Given Instead of δ) *The relationship between the distortion D and the "probabilistic level" $0 < \delta < 1$ in Theorem 5 can be inverted. In that case, our proof implies that for any prespecified distortion level $D > 2$ there is a PTT as in Theorem 5 satisfying:*

$$\mathbb{P}\left(sd_{\mathcal{X}}(x, \tilde{x}) \leq W_2(T(x), T(\tilde{x})) \text{ and } \mathcal{M}W_2(T(x), T(\tilde{x})) \leq sD(D-1)d_{\mathcal{X}}(x, \tilde{x})\right) \geq n^{-2+2\theta_D},$$

where θ_D is the unique solution to $\frac{2}{D} = (1 - \theta)^{\theta/(1-\theta)}$ over $[1 - 2e/D, 1)$.

Juxtaposing either of the impossibility results in Proposition 15 or Proposition 13 with Theorem 4 (as formulated in Remark 16), we find that, even if there are n -point metric spaces which simply cannot be well-represented in most Riemannian representation spaces, we can always identify a subset of *any such*

space which *can* be arbitrarily well-represented in $\mathcal{MW}_2(\mathbb{R})$. Furthermore, the embeddings are explicitly implemented by PT models whose required number of parameters we now know. The distortion $D > 2$ controls the size of this subspace we are willing to accept. For instance when comparing against (13) we can select any distortion $D \in (2, c_{\mathcal{R}} \sqrt{\frac{4}{3} \log_2(n) - 2}]$ and always deduce the existence of an embedding on a subset of those sextet graphs (this is not necessarily the case for a complete Riemannian representation space with pinched negative curvature).

5. Experiments

In this section we complement the theoretical embedding guarantees from Section 4 by preliminary computer experiments on synthetic data. We show that the proposed feature maps can indeed be trained in a standard deep learning framework, that the theoretical advantages of PT mixture-Wasserstein embeddings over Euclidean and hyperbolic carry over to practice, and that the PT-based feature maps generalize beyond \mathcal{X}_n . We compare our proposed representation maps to learned representation maps in Euclidean and hyperbolic space when representing trees, random graphs and graphs “sampled from” Riemannian manifolds.^{10,11}

Our first experiment compares the capacity of deep neural embedding in the space $(\mathcal{GM}_2(\mathbb{R}), \mathcal{WG}_2(\mathbb{R}))$ to embed combinatorial trees against the state of the art; namely embeddings into the hyperbolic plane. Our second experiment examines the efficiency with which PTs utilize dimension. We compare how well the PT embeds points on a simple high-dimensional compact Riemannian manifold, namely the sphere, to how well a standard feed-forward network can embed those same points into a Euclidean space of equal dimension.

The probabilistic transformer we implement is illustrated in Figure 6a. The weights of the network are computed following (Delon and Desolneux, 2020), by minimizing

$$\mathcal{L}_{\text{PT}}(\theta) = \sum_{x, y \in \mathcal{X}_{\text{train}}} (\mathcal{MW}_2^2(T_\theta(x), T_\theta(y)) - d_{\mathcal{X}}^{2\alpha}(x, y))^2. \quad (\text{Loss})$$

In practice we set $\alpha = 1$ as it does not have a strong influence on empirical performance. We conjecture that this is due to the influence of sampling and optimization which result in an “error floor”.

Numerical parameters All networks are trained by the Adam optimizer in `pytorch`, with weight decay parameter 10^{-6} , initial learning rate 10^{-4} and final learning rate 10^{-6} .

5.1 Hyperbolic vs. Gaussian Mixtures Geometry: Embedding Trees

Hyperbolic embeddings are the state-of-the-art when embedding metric trees. Indeed, results such as (Sarkar, 2011) show that two dimensions is enough to embed any metric tree with low distortion whereas finite trees can only be embedded into d -dimensional Euclidean space with low distortion if d is large (Gupta, 2000).

We compare the ability of $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ ’s geometry to accommodate deep neural embeddings of metric trees against that of the hyperbolic plane. We implement the embedding maps using (universal (Kratsios and Papon, 2022; Acciaio et al., 2023)) overparameterized neural network models, in an effort to minimize the chance that differences in expressivity obscure the geometric effects. As we show below, the additional flexibility of $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ over the hyperbolic plane is manifested even when embedding regular binary trees.

EMBEDDING INTO THE HYPERBOLIC PLANE

We first consider the hyperbolic plane (Ganea et al., 2018; Cruceru et al., 2021) which theoretically embeds trees with arbitrarily small distortion (Sarkar, 2011, Theorem 6). In the next section we look at a higher-dimensional

10. More precisely, on neighborhood graphs constructed over points sampled from manifolds.

11. The Python codes used to produce the results of this section are available at <https://github.com/swing-research/Universal-Embeddings>.

hyperbolic space of the same dimension as the number of the degrees of freedom of our mixture-Wasserstein embedding.

The hyperbolic space has several isometric (up to a scalar) representations. In this section, we first rely on its canonical (in the sense of (Dowty, 2018)) *information geometric* representation¹². The set of *non-degenerate Gaussian* probability measures on \mathbb{R} is defined by

$$\mathcal{G}_1^+ \stackrel{\text{def}}{=} \left\{ \mathcal{N}(\mu, \sigma) \in \mathcal{P}_1(\mathbb{R}) : \frac{N(\mu, \sigma)}{dx} \propto \exp\left(-\frac{(x - \mu)^2}{2\sigma^2}\right), \mu \in \mathbb{R}, \sigma \in (0, \infty) \right\}.$$

This set can be made into a Riemannian manifold with the Riemannian metric defined by the *Fisher information matrix* (Amari, 2016). As shown in (Costa et al., 2015, Equation (7)), the geodesic distance (called the Fisher-Rao distance) between two Gaussian probability measures $N(\mu_1, \sigma_1), N(\mu_2, \sigma_2) \in \mathcal{G}_1^+$ in this Riemannian geometry is

$$d_F(N(\mu_1, \sigma_1), N(\mu_2, \sigma_2)) = \sqrt{2} \ln \left(\frac{\|(\frac{\mu_1}{\sqrt{2}}, \sigma_1) - (\frac{\mu_2}{\sqrt{2}}, -\sigma_2)\| + \|(\frac{\mu_1}{\sqrt{2}}, \sigma_1) - (\frac{\mu_2}{\sqrt{2}}, \sigma_2)\|}{\|(\frac{\mu_1}{\sqrt{2}}, \sigma_1) - (\frac{\mu_2}{\sqrt{2}}, -\sigma_2)\| - \|(\frac{\mu_1}{\sqrt{2}}, \sigma_1) - (\frac{\mu_2}{\sqrt{2}}, \sigma_2)\|} \right).$$

Moreover, importantly for our comparison, d_F is (up to a constant factor of $\sqrt{2}$) equal to the hyperbolic distance on $\mathbb{H}^2 \stackrel{\text{def}}{=} \mathbb{R} \times (0, \infty)$ which is precisely the parameter space of \mathcal{G}_1^+ as well as the upper half-plane model of the hyperbolic space. There are two reasons to first look at the 2D hyperbolic embeddings. First, they are easy to visualize; second, it allows us to some extent to dissociate the effects of geometry and expressivity and ease of optimization of neural networks. We benchmark our representations against the metric transformer model of (Acciaio et al., 2023, Example 11) (which for clarity we call the \mathcal{G}_1^+ -transformer) since that model is universal (Acciaio et al., 2023, Theorem 3.6). It is a variant of hyperbolic neural networks (Cruceru et al., 2021), which are known to be universal (Bilokopytov and Kratsios, 2020, Corollary 3.16).

The weights of the metric transformer T_θ are obtained by minimizing

$$L_{\mathcal{G}_1^+}(\theta) = \sum_{x, y \in \mathcal{X}_{\text{train}}} \left(d_F^2(T_\theta(x), T_\theta(y)) - d_{\mathcal{X}}^2(x, y) \right)^2. \quad (11)$$

The main difference with our model is that the metric transformer makes convex combinations in the parameter space of Gaussian measures, not in the space of probability measures itself.

EMBEDDING INTO THE d -DIMENSIONAL HYPERBOLIC SPACE

Even though the hyperbolic plane is theoretically sufficient to embed trees with small distortion (Sarkar, 2011), experimental evidence indicates that neural networks representations into higher dimensional hyperbolic space perform better (Giovanni et al., 2022). We therefore also compare our representations with representations in d -dimensional hyperbolic space represented by the Poincaré upper-half space model, whose underlying set it

$$\mathbb{H}^d = \{x \in \mathbb{R}^{d+1} \text{ such that } x_{d+1} > 0\},$$

and whose geodesic distance is given by

$$d_{\mathbb{H}^d}(x, y) = \text{arccosh} \left(1 + \frac{\|x - y\|_2^2}{2x_{d+1}y_{d+1}} \right), \forall x, y \in \mathbb{H}^d.$$

We adapt the output of the probabilistic transformer to return vectors in \mathbb{H}^d , by swapping its last layer for a trainable linear readout map (see Figure 6c) and modifying the embedding objective to minimize the distortion with respect to the hyperbolic distance. The d -dimensional hyperbolic model is then trained using the following loss function:

$$L_{\mathbb{H}^d}(\theta) = \sum_{x, y \in \mathcal{X}_{\text{train}}} \left(d_{\mathbb{H}^d}^2(T_\theta(x), T_\theta(y)) - d_{\mathcal{X}}^2(x, y) \right)^2. \quad (12)$$

12. Because it shares many similarities with our embedding space $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$

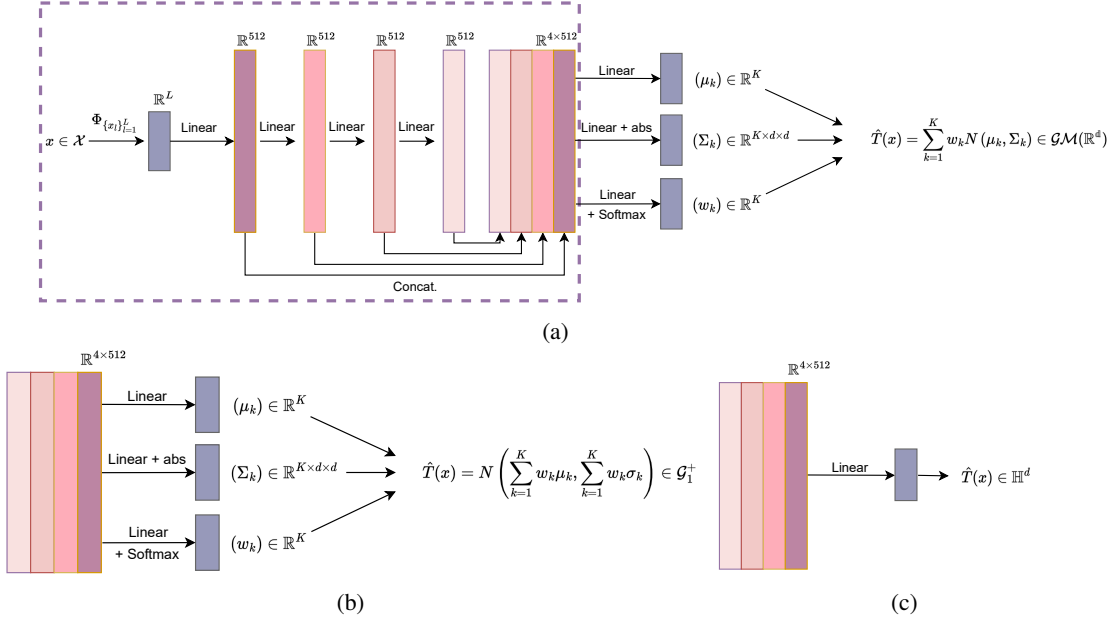


Figure 6: Two different networks implementing feature (embedding) maps. The first block (violet dashed line) is common to all networks; the output blocks change to accommodate different target spaces. a) Probabilistic transformer of (8). b) \mathcal{G}_1^+ -transformer. c) \mathbb{H}^d -transformer.

EXPERIMENTAL RESULTS

We consider a regular binary tree $\mathcal{X} = (V, E)$ (Figure 7a) of depth six with a total of $|V| = 127$ vertices. The distance $d_{\mathcal{X}}$ is the shortest path distance. We partition the vertices V into training and testing sets, $V_{\text{train}} \cup V_{\text{test}} = V$, with $|V_{\text{train}}| = 111$ and $|V_{\text{test}}| = 16$. The test vertices (colored white in Figure 7a) are used to evaluate the quality of out-of-sample representations (that is to say, the generalization) computed by the different representation maps. We designate $L = 20$ training vertices as landmarks (colored black in Figure 7a). The two probabilistic transformers take as input the distances between a vertex and the L landmarks. The neural networks contains more than 7,000 parameters with only the last layer being adapted to produce an output in the desired space. We use $K = 5$ mixture components and $d = 15$ for the dimension of the hyperbolic space to ensure a fair comparison with the probabilistic transformer’s effective dimension.

		$\mathcal{GM}_2(\mathbb{R})$	\mathbb{H}^2	\mathbb{H}^{15}
Training points	Mean error	0.03	0.15	0.04
	Max. error	0.25	0.44	0.26
Testing points	Mean error	0.02	0.19	0.05
	Max. error	0.25	0.44	0.26

Table 2: Average and maximum relative embedding errors for binary tree representations.

Table 2 shows that our neural representations in the space of Gaussian mixtures perform better than \mathbb{H}^2 even when embedding simple trees, and that they perform on par or slightly better than \mathbb{H}^{15} . The difference in performance between \mathbb{H}^2 and \mathbb{H}^{15} epitomizes the difficulty in verifying theoretical facts with computer experiments: even though theoretically both spaces can achieve low distortion, it is easier to approximate a good feature map for \mathbb{H}^{15} . We report the average relative absolute difference between the pairwise distances in

the binary tree as compared to embeddings in the respective representation spaces. The training set evaluates the embedding quality of the given set of landmarks. The role of the test set is to check whether the learned representation still provides a good embedding using the set of landmarks in the binary graph¹³.

We visually assess the quality of the computed representations in Figure 7: Figure 7b displays the probability density functions of the $\mathcal{GM}(\mathbb{R})$ representations and Figure 7c displays the information-theoretic hyperbolic representations using Gaussian measures. In both cases the networks learn representations that change continuously with the shortest path distance in \mathcal{X} .

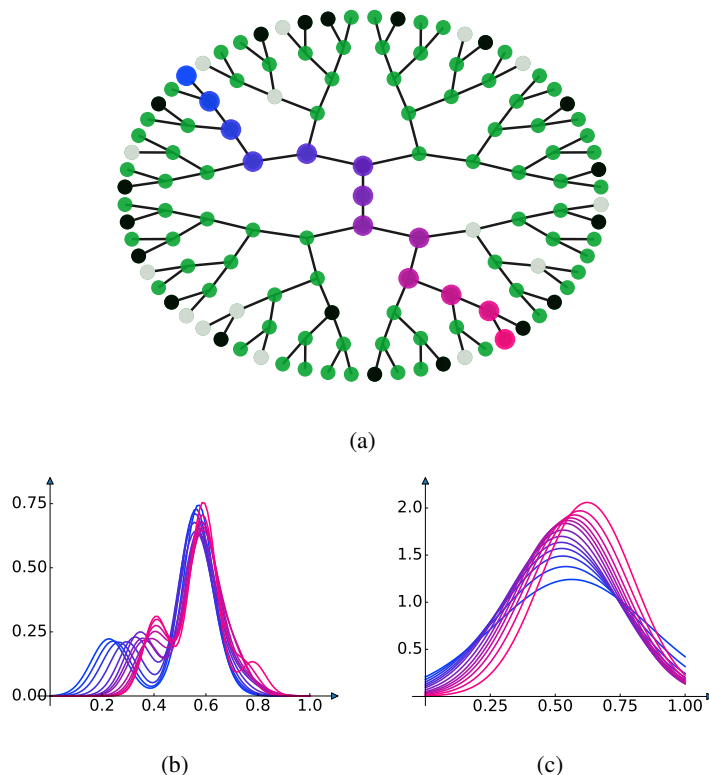


Figure 7: Two embeddings of colored vertices in the binary tree \mathcal{X} shown in (a); white vertices are not seen during training; landmarks are colored black. We apply the transformation $\sigma \mapsto \log_{10}(\sigma + 2.1)$ to aid visualization. (b) Mixture Gaussian Wasserstein embeddings represented by their density. (c) hyperbolic (univariate Gaussian) embeddings represented by their density.

Figure 7 visually examines the stability of the embeddings produced by the PT against those learned by its hyperbolic counterpart. Since both output univariate probability distributions, a visual comparison can easily be drawn. When comparing Figure 7c to Figure 7b we see that the embedding of the path in Figure 7 is only a bit more contorted than the embedding produced by the PT’s hyperbolic counterpart. This can explain why both embeddings generalize similarly well, but that the gap in embedding quality is explained by the richness of the geometry of $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ as compared to the rigidity of the hyperbolic plane’s geometry.

13. Our theorems give “memorization” guarantees, or, phrased differently, approximation guarantees for embeddings of finite metric spaces. In computer experiments we thus deliberately explore the generalization aspect. The related theoretical questions remain open.

5.2 PTs Leverage Dimension Efficiently: Manifold Embeddings

We now compare the capacity of the Euclidean space with $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ to represent smooth geometries using few effective dimensions. We represent datasets from N -dimensional spheres equipped with the geodesic distance. We remark that, as shown in (Robinson, 2006b), there exist no bi-Lipschitz embeddings of the sphere into any Euclidean space with arbitrarily low distortion.¹⁴ In this case our findings corroborate known facts and also show that the PT leverages the rich geometry of $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ to produce better low-dimensional embeddings standard feedforward neural networks taking values in high-dimensional Euclidean spaces.

EMBEDDING INTO EUCLIDEAN SPACE

Despite negative theoretical results, Euclidean representations are attractive in practice thanks to the wealth of available machine learning and numerical tools. It is thus desirable to compare them with our proposed representations in the space of measures. To this end we modify the output of our probabilistic transformer to output vectors in \mathbb{R}^d , by swapping its last layer for a trainable linear readout map (see Figure 8) and modifying the embedding objective to minimize the distortion with respect to the Euclidean distance $\|\cdot - \cdot\|_{\ell_d^2}$. Naively, one may expect a well-performing Euclidean representation for sufficiently high embedding dimension d . However, as shown below, even for very large d there is a hard limit to representation quality, reflecting the introductory theoretical notes.

We let $\mathcal{S}^N \stackrel{\text{def.}}{=} \{z \in \mathbb{R}^{N+1} : \|z\| \leq 1\}$ denote the N -dimensional sphere and equip it with the geodesic distance $d_{\mathcal{S}^N}(x, \tilde{x}) = \cos^{-1}(x^\top \tilde{x})$, $x, \tilde{x} \in \mathcal{S}^N$ (Dai and Müller, 2018; Fletcher et al., 2009). We randomly sample data points $\{x_i\}_{i=1}^n$ from the uniform probability measure on \mathcal{S}^N . The resulting Euclidean model is a feedforward network with a suitable Euclidean loss function,

$$\mathcal{L}_{\text{Euclid}}(\theta) = \sum_{x, y \in \mathcal{X}_{\text{train}}} \left(\|T_\theta(x) - T_\theta(y)\|_{\ell_d^2}^2 - d_{\mathcal{X}}^2(x, y) \right)^2. \quad (13)$$

5.2.1 2-DIMENSIONAL SPHERE \mathcal{S}^2

We begin by visualizing the $\mathcal{GM}_2(\mathbb{R})$ embeddings of points on the 2-sphere to develop intuition. A good representation should achieve comparably low distortion across the entire manifold and vary continuously with respect to the geodesic distance on \mathcal{S}^2 . The learned feature map should generalize well from a set of landmarks to arbitrary points on the sphere.

We train a PT for 160 iterations with the Adam optimizer (Kingma and Ba, 2015). In each iteration, we use random batch of 32 points among the 10,000 fixed training points chosen from a uniform measure on the sphere. The 13 landmarks are not updated during training. The network contains about 200,000 parameters and returns a Gaussian mixture with $K = 5$ components.

14. We note that this is different from positive results on Riemannian embeddings of the sphere in Euclidean spaces as exemplified by Nash’s theorems (Nash, 1954). These “embeddings” view the sphere as a submanifold of the Euclidean space but make no guarantees that the embedding preserves the geodesic distance.




Figure 8: Replacing the PT’s final layer by a linear readout layer yields our feedforward network benchmark.

The results are shown in Figure 9. For each point x_k , colors in Figures 9a and 9b represent the quantities

$$\frac{1}{K} \sum_{j=1}^K \left| d_{\mathcal{M}}^2(x_k, x_j) - \mathcal{MW}_2^2(\hat{T}_\theta(x_k), \hat{T}_\theta(x_j)) \right| \quad \text{and} \quad \sup_{1 \leq i, j \leq K} \left| d_{\mathcal{M}}^2(x_i, x_j) - \mathcal{MW}_2^2(\hat{T}_\theta(x_i), \hat{T}_\theta(x_j)) \right|.$$

We can see from the figure that the distances in the representation space faithfully represent those on the sphere. This is further detailed in Figure 9c where color of a point encodes the embedding error between the point and the red cross. Notice that even if the maximum error between points may seem large (about 0.35), it is reached only for a very small number of point pairs; most pairs have a relative error below 10^{-2} .




Figure 9: Embedding quality of 2-dimensional sphere S^2 in $\mathcal{GM}_2(\mathbb{R})$. a) Average error of a point and all the other elements. b) Maximum error of a point and all the other elements. c) Error between the point marked by a red cross and all other points.

We then assess the continuity of the embedding in Figure 10. We plot the probability density function of a sequence of points spiralling up from the bottom of the sphere to the top, which we colour-code as progressing from red to blue. As expected, the Gaussian mixtures representing points on the sphere progressively change. We also see that only a few significant mixture components are required to perform the embedding. To aid visualization, we normalize the support of the densities and transform the standard deviation of the Gaussian components using the mapping $\sigma \mapsto \log_{10}(\sigma + 2.1)$.




Figure 10: Probability density functions of the embeddings of points on S^2 into $\mathcal{GM}_2(\mathbb{R})$.

5.2.2 EMBEDDING THE N -DIMENSIONAL SPHERE

The motivation to use Gaussian mixtures instead of Euclidean embeddings is that they easily capture complex geometries; achieving comparable distortion would either require a very high-dimensional Euclidean embedding or is simply impossible due to curvature. To illustrate this, we set up an experiment on N -spheres where we vary N . We fix the number of mixture components at the output of the PT to $K = 5$. The Euclidean network outputs vectors in \mathbb{R}^{15} , thus having the same number of degrees of freedom. We also compare these two networks with embedding in the hyperbolic space \mathbb{H}^{15} . All networks contain about 200,000 parameters. We set the number of landmarks to $L = N + 10$ (note that we need at least N landmarks to uniquely localize a point on \mathcal{S}^N). We distribute the landmarks quasi-uniformly on the N -sphere adapting a procedure from (Debarnot and Weiss, 2022).

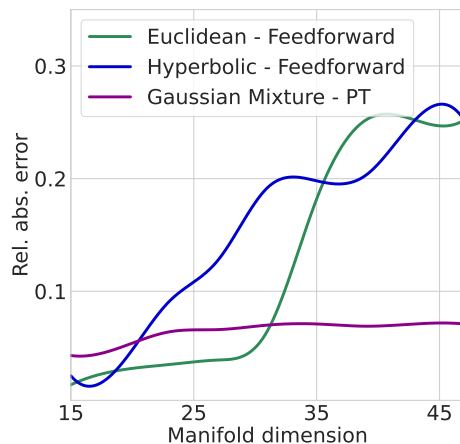


Figure 11: Impact of dimension.

Figure 11 shows the results. As the manifold dimension N increases, the quality of the Euclidean and Hyperbolic embeddings implemented by the deep feedforward networks rapidly deteriorate. In contrast, the distortion of the embedding implemented by the probabilistic transformer into the space $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ of (Delon and Desolneux, 2020) remains stable.

6. Outline of the Proofs of Main Results

This section overviews the main steps toward proving Theorem 4. Similar techniques are used in deriving the probabilistic Theorem 5. The latter proof is more technical as it relies on the non-linear Dvoretzky theorem (Naor and Tao, 2012) which is a metric version of the classical result of Dvoretzky (Dvoretzky, 1961); we defer it to Appendix B.

Our proofs use a novel two-step approach. In *step one*, we use a key lemma (Lemma 17 below) to show that certain probability-measure-valued functions can be implemented by metric transformers with few parameters. In *step two*, given any of the three aforementioned “geometric priors”, we demonstrate the existence of “optimal” embeddings of the required form. Each of the proofs then concludes by applying the key lemma to show that the embedding can be *exactly* implemented (memorized) by the probabilistic transformer. The proof is completed by remarking that the PT we have built is defined on the entire metric space \mathcal{X} .

We note that the d -dimensional Gaussian measure with mean $\mu \in \mathbb{R}^d$ and $d \times d$ covariance matrix Σ is denoted by $N_d(\mu, \Sigma)$.

Lemma 17 (Key Lemma: Probabilistic Transformers Memorize Empirical Markov Kernels) *Let d, K be a positive integer, let $1 \leq p < \infty$, $\#\mathcal{X}_n = n > 2$, and let $\sigma \stackrel{\text{def.}}{=} \text{ReLU}$. Let (\mathcal{X}_n, d_n) be an n -point metric*

space and for $k = 1, \dots, K$, let $\mu_k : \mathcal{X}_n \rightarrow \mathbb{R}^d$. Define the map $\varphi : \mathcal{X}_n \rightarrow \mathcal{GM}_2(\mathbb{R}^d)$ by

$$\varphi(x) \stackrel{\text{def.}}{=} \frac{1}{K} \sum_{k=1}^K N_d(\mu_k(x), 0).$$

Then, there exists a probabilistic transformer with unnormalized graph attention $T : \mathcal{X}_n \rightarrow \mathcal{GM}_2(\mathbb{R}^d)$, which exactly implements φ ; i.e.:

$$\varphi(x) = T(x) \quad \text{for every } x \in \mathcal{X}_n$$

such that

$$\text{width}(T) = \max\{n, K, d^2, n(n-1) + \max\{d, 12\}\}, \quad (14)$$

$$\text{depth}(T) = \mathcal{O} \left(n \left\{ 1 + \sqrt{n \log(n)} \left[1 + \frac{\log(2)}{\log(n)} \left(C_n + \frac{\log(n^{5/2} \text{aspect}(\mathcal{X}_n, d_n))}{\log(2)} \right)_+ \right] \right\} \right) \quad (15)$$

$$\text{effdim}(T) \leq K(d + d^2) \quad (16)$$

$$\begin{aligned} \text{par}(T) = \mathcal{O} \left(K n \left(\frac{11}{4} \max\{n, d\} n^2 - 1 \right) \left\{ d + \max\{d, 12\}^2 \sqrt{n \log(n)} \right. \right. \\ \left. \left. \times \left[1 + \frac{\log(2)}{\log(n)} \left(C_n + \frac{\log(n^{5/2} \text{aspect}(\mathcal{X}_n, d_n))}{\log(2)} \right)_+ \right] \right\} \right). \end{aligned} \quad (17)$$

Moreover, the ‘‘dimensional constant’’ $C_n > 0$ is defined by $C_n \stackrel{\text{def.}}{=} \frac{2 \log(5\sqrt{2\pi}) + \frac{3}{2} \log(n) - \frac{1}{2} \log(n+1)}{2 \log(2)}$.

Before stating the next lemma, let us recall the notion of a *doubling metric space*. Briefly, these are metric spaces for which a maximal ratio determines the number of open balls of finite radius required to cover an open ball of twice their radius. More precisely, a metric space (\mathcal{X}, d) is said to be doubling if and only if, there is a constant $C_d > 0$ for every point $x \in \mathcal{X}$ and every radius $r > 0$ the open metric ball about x of radius r , i.e. the set $B(x, r) \stackrel{\text{def.}}{=} \{u \in \mathcal{X} : d(u, x) < r\}$, there are at most C_d points $x_1, \dots, x_n \in \mathcal{X}$ (i.e. $n \leq C_d$) such that the collection of open metric balls $\{B(x_i, r/2)\}_{i=1}^n$ cover $B(x, r)$. The smallest constant $C_d > 0$ for which the above relation holds is called the *doubling constant* of (\mathcal{X}, d) . Examples of doubling metric spaces are any finite metric space, Carnot Groups such as the Heisenberg group (Pansu, 1989), and more generally, any metric space that can be bi-Hölder embedded into some Euclidean space (see (Heinonen, 2001, Theorem 12.2) and (Naor and Neiman, 2012)).

The doubling constant C_d of a doubling metric space (\mathcal{X}, d) plays the role of dimension; for example, for a compact n -dimensional Riemannian manifold $\log_2(C_d)$ is of the order $\mathcal{O}(n)$ and the same is true of the n -dimensional Euclidean space. However, in general it can be difficult to compute the doubling constant of (\mathcal{X}, d) . Thus, we turn to the notion of *metric capacity* as defined in (Bruè et al., 2021). The reason is that, as shown in (Bruè et al., 2021, Proposition 1.7) the doubling constant and the metric capacity are proportional (we refer the reader to that paper or the proof of Lemma 21 for details).

Our second main embedding lemma concerns the case where (\mathcal{X}_n, d_n) has its points drawn from a doubling metric space. Since the logarithm of the doubling constant of such spaces plays the role of topological dimension for manifolds (see (Heinonen, 2001, Chapter 12)), it is intuitive that the number of (parameterized) point masses required to embed (\mathcal{X}_n, d_n) into the Wasserstein space depends on the logarithm of the doubling constant.

Lemma 18 (Datasets in Doubling Metric Spaces bi-Hölder Embed into $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$) *Let (\mathcal{X}, d) be a doubling metric space with doubling constant $C_d > 0$, $\mathcal{X}_n \subseteq \mathcal{X}$, and $d_n \stackrel{\text{def.}}{=} d_{\mathcal{X}_n}$. There are constants $c, C > 0$, such that for any $\frac{1}{2} < \alpha < 1$, there exists an injective map $\varphi : (\mathcal{X}_n, d_n) \hookrightarrow (\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ of the form*

$$\varphi(x) = \frac{1}{\lceil c\alpha^{-1} \log(C_d) \rceil} \sum_{k=1}^{\lceil c\alpha^{-1} \log(C_d) \rceil} N_1(\phi(x)_k, 0),$$

satisfying: for every $x, \tilde{x} \in \mathcal{X}$ the following holds

$$d_{\mathcal{X}}^{\alpha}(x, \tilde{x}) \leq \mathcal{W}_2(\varphi(x), \varphi(\tilde{x})) \leq \mathcal{M}\mathcal{W}_2(\varphi(x), \varphi(\tilde{x})) \leq \left[C \left(\frac{\log(C_d)}{1-\alpha} \right)^{1+\alpha} \right] d_{\mathcal{X}}^{\alpha}(x, \tilde{x}),$$

where $\phi : (\mathcal{X}_n, d_n) \rightarrow (\mathbb{R}^{\lceil c \log(C_d) \rceil}, \ell_2)$ is a α -Hölder with α -Hölder constant $\left[C \left(\frac{\log(C_d)^2}{1-\alpha} \right)^{1+\alpha} \right]$.

Since every *finite* metric space is a doubling space, then Lemma 18 applies to any finite subset \mathcal{X}_n of any general metric space \mathcal{X} . Consequentially, Theorem 4 is obtained by applying Lemma 17 to the embedding in Lemma 18.

7. Limitations and Possible Improvements

As mentioned in the introduction, our results are the first theoretical guarantees for deep neural representations of finite datasets from arbitrary metric spaces. We prove several embedding guarantees which pleasingly reflect the impact of the ambient geometry of space on the complexity of the representations and the number of parameters of the feature map. These results can be seen as deep neural approximation theorems for embeddings of finite metric spaces which are usually studied in metric geometry non-constructively. To obtain these results we developed a new proof technique which combines elements of metric embedding theory and memorization theory for deep neural networks; this technique may be a tool of independent interest.

At the same time, the preludial nature of our work makes it likely that a number of improvements are possible. One natural question is whether we can extend the n -point guarantees to the whole \mathcal{X} when \mathcal{X} has a regular geometry (for example, when it is a compact Riemannian manifold). While this may require rephrasing some of our questions and using different mathematical tools, it does seem plausible.

Further, there are several aspects of the proofs, for embeddings into the space of Univariate Gaussian mixtures, which suggest improvements are possible: 1) a key step passes through a high-dimensional Euclidean space, but we know experimentally that Euclidean embeddings under-perform and theoretically (Naor, Andoni, Neiman (Andoni et al., 2018)) that this is not necessary; 2) the Gaussian mixtures are often degenerated into empirical measures with equal atomic weights, but numerical experiments suggest a clear advantage of using proper Gaussian mixtures. While at the moment we do not know how to leverage these observations, they remain a subject of ongoing work.

FUNDING

This research was supported by the European Research Council (ERC) Starting Grant 852821-SWING. Anastasis was also supported by the NSERC Discovery grant (RGPIN-2023-04482) and by their McMaster University Startup Grant.

ACKNOWLEDGMENTS

The authors would like to thank Giulia Livieri for her very helpful feedback.

A. Explicit Network Complexities: Theorems and Examples

The next table summarizes the precise complexities of the probabilistic transformer representing the metric spaces with generic n -point embedding guarantees, which are covered in any one of the theoretical results in this paper. This table is an explicit analogue of Table 1.

Remark 19 (Approximate Complexity of T For Low-Distortion Embeddings) *The proof of Theorem 5 and a remark on (Naor and Tao, 2012, page 492) shows if $D \approx 2$ (but $D > 2$) then, the probabilistic*

Table 3: Width of the probabilistic transformers performing n -point embeddings, with explicit constants.

Latent Geometry	Effective Dimension d_T	Width	Result
General	$2\lceil 12C\alpha^{-1}(\log(\text{cap}(\mathcal{X}_n, d_n))) \rceil$	$\max\{d_T, n(n-1) + 12\}$	Theorem 4
General	$\mathcal{O}((D-2)^{-2}\theta_D \log_2(n))$	$\mathcal{O}(\max\{d_T, n^{2\theta_D} + 12\})$	Theorem 5
General - Multivariate Mixtures	$12n(n-1)/2(5\frac{5n^4}{2(D-1)} \text{aspect}(\mathcal{X}_n, d)^2 + 2)$	$\max\left\{n, \frac{25n^5(n-1)}{4(D-1)}(\text{aspect}(\mathcal{X}_n, d)^2 + 2), 9, n(n-1) + 12\right\}$	Theorem 7
Discrete: Trees	$2M$	$\max\{d_T, n(n-1) + 12\}$	Proposition 11
Discrete: 2-Hop Graphs	$2\lceil 12C\alpha^{-1} \log(1 + \rho(A_G)) \rceil$	$\max\{d_T, n(n-1) + 12\}$	Corollary 12
Manifold: Riemannian, Bounded Curvature	$2\lceil \tilde{C} \frac{m^{1+\alpha}}{\alpha(1-\alpha)^{1+\alpha}} \rceil$	$\max\{d_T, n(n-1) + 12\}$	Corollary 9
Manifold: Riemannian & Compact	$4d$	$\max\{d_T, n(n-1) + 12\}$	Proposition 8

 Table 4: Depth of the probabilistic transformers performing n -point embeddings, with explicit constants.

Here $c \stackrel{\text{def}}{=} \frac{2 \log(5\sqrt{2\pi}) + \frac{3}{2} \log(n) - \frac{1}{2} \log(n+1)}{2 \log(2)}$.

Latent Geometry	Depth	Result
General	$\mathcal{O}\left(n \left(1 + \sqrt{n \log(n)} \left(1 + \frac{\log(2)}{\log(n)} \left[c + \frac{\log(n^{5/2} \text{aspect}(\mathcal{X}_n, d_n))}{\log(2)}\right]_+\right)\right)\right)$	Theorem 4
General	$\mathcal{O}\left(n^{\theta_D} \left(1 + \sqrt{\theta_D \log(n)} \left(1 + \frac{\log(2)}{\theta_D \log(n)} \left[c + \frac{3\theta_D \log(n) + \log(\text{aspect}(\mathcal{X}_n, d_n))}{\log(2)}\right]_+\right)\right)\right)$	Theorem 5
General - Multivariate Mixtures	$\mathcal{O}\left(n \left\{1 + \sqrt{n \log(n)} \left[1 + \frac{\log(2)}{\log(n)} \left(C_n + \frac{\log(n^{5/2} \text{aspect}(\mathcal{X}_n, d_n))}{\log(2)}\right)\right]_+\right\}\right)$	Theorem 7
Discrete: Trees	$\mathcal{O}\left(n \left(1 + \sqrt{n \log(n)} \left(1 + \frac{\log(2)}{\log(n)} \left[c + \frac{\log(n^{5/2} \text{diam}(\mathcal{X}_n, d_n))}{\log(2)}\right]_+\right)\right)\right)$	Proposition 11
Discrete: 2-Hop Graphs	$\mathcal{O}\left(n \left(1 + \sqrt{n \log(n)} \left(1 + \frac{\log(2)}{\log(n)} \left[c + \frac{\log(n^{5/2} \text{diam}(\mathcal{X}_n, d_n))}{\log(2)}\right]_+\right)\right)\right)$	Corollary 12
Manifold: Riemannian, Bounded Curvature	$\mathcal{O}\left(n \left(1 + \sqrt{n \log(n)} \left(1 + \frac{\log(2)}{\log(n)} \left[c + \frac{\log(n^{5/2} \text{aspect}(\mathcal{X}_n, d_n))}{\log(2)}\right]_+\right)\right)\right)$	Corollary 9
Manifold: Riemannian & Compact	$\mathcal{O}\left(n \left(1 + \sqrt{n \log(n)} \left(1 + \frac{\log(2)}{\log(n)} \left[c + \frac{\log(n^{5/2} \text{aspect}(\mathcal{X}_n, d_n))}{\log(2)}\right]_+\right)\right)\right)$	Proposition 8

transformer T in Theorem 5 has depth about the order¹⁵ $\tilde{\mathcal{O}}\left(n^{\frac{11(D-2)}{\log(1/(D-2)^4)}} + \mathcal{O}\left(\frac{11(D-2) \log \log(1/(D-2))}{2 \log(1/(D-2))^2}\right)\right)$. When n is also large then it has width of the order $\tilde{\mathcal{O}}\left(n^{\frac{D-2}{\log(D-2)}} + \mathcal{O}\left(\frac{(2D-4) \log \log(1/(D-2))}{\log(1/(D-2))^2}\right)\right)$.

B. Proof Details

This section contains the paper’s main proofs.

We begin with the following memorization lemma, on which the proof of Lemma 17 is developed. This result extends the quantitative “class memorization” result of (Vardi et al., 2022) to a quantitative vector-valued interpolation result. The next Lemma can be contrasted with the VC-bounds of (Bartlett et al., 2019) which imply that a ReLU feedforward network of depth D memorizing how to match N inputs in \mathbb{R}^n with C classes must have at-least $\Omega\left(\frac{D}{\ln(D)}\right)$ parameters. We recall that the following convenient notation is used in the next lemma. For each $u \in \mathbb{R}$ we denote $[u]_+ \stackrel{\text{def}}{=} \max\{u, 1\} = \text{ReLU}(u - 1)$.

Lemma 20 (Memory Capacity of Deep ReLU Regressors) *Let $n, d, N \in \mathbb{N}_+$, let $f : \mathbb{R}^n \rightarrow \mathbb{R}^d$ be some function, and consider distinct $x_1, \dots, x_N \in \mathbb{R}^n$. There exists a deep ReLU network $\mathcal{NN} : \mathbb{R}^n \rightarrow \mathbb{R}^d$ satisfying*

$$\mathcal{NN}(x_i) = f(x_i),$$

15. We use the notation $\tilde{\mathcal{O}}$ to hide logarithmic factors.

for every $i = 1, \dots, N$. Furthermore, we have the following quantitative model complexity estimates:

$$\begin{aligned} \text{width}(\mathcal{NN}) &= n(N-1) + \max\{d, 12\}, \\ \text{depth}(\mathcal{NN}) &= \mathcal{O} \left(N \left\{ 1 + \sqrt{N \log(N)} \left[1 + \frac{\log(2)}{\log(N)} \left(C_n + \frac{\log(N^2 \text{aspect}(\mathcal{X}_N, \|\cdot\|_2))}{\log(2)} \right) \right]_+ \right\} \right), \\ \text{par}(\mathcal{NN}) &= \mathcal{O} \left(N \left(\frac{11}{4} \max\{n, d\} N^2 - 1 \right) \left\{ d + \sqrt{N \log(N)} \left[1 + \frac{\log(2)}{\log(N)} \right. \right. \right. \\ &\quad \left. \left. \left. \times \left(C_n + \frac{\log(N^2 \text{aspect}(\mathcal{X}_N, \|\cdot\|_2))}{\log(2)} \right) \right] \max\{d, 12\} [1 + \max\{d, 12\}] \right\} \right). \end{aligned}$$

The “dimensional constant” is defined by $C_n \stackrel{\text{def}}{=} \frac{2 \log(5\sqrt{2\pi}) + \frac{3}{2} \log(n) - \frac{1}{2} \log(n+1)}{2 \log(2)} > 0$.

Proof [Proof of Lemma 20] Step 1: Centering Data

Let $\mathcal{X}_N \stackrel{\text{def}}{=} \{x_i\}_{i=1}^N$ and for any $\bar{x} \in \mathbb{R}^n$ and any $r > 0$ denote $\overline{B_2(\bar{x}, r)} \stackrel{\text{def}}{=} \{u \in \mathbb{R}^n : \|u - \bar{x}\|_2 \leq r\}$. Since \mathcal{X}_N is finite it is compact and therefore by Jung’s Theorem (see (Jung, 1901)) there exists some $\bar{x} \in \mathbb{R}^n$ such that

$$\mathcal{X}_N \stackrel{\text{def}}{=} \overline{B_2(\bar{x}, r)} \text{ and } r \stackrel{\text{def}}{=} \text{diam}(\mathcal{X}_N, \|\cdot\|_2) \frac{n^{1/2}}{(2(n+1))^{1/2}},$$

where $\text{diam}(\mathcal{X}_N, \|\cdot\|_2) \stackrel{\text{def}}{=} \max_{1 \leq i, j \leq N} \|x_i - x_j\|_2$. Thus, the (affine) isometry $W : \mathbb{R}^n \rightarrow \mathbb{R}^n$ sending any $x \in \mathbb{R}^n$ to $x - \bar{x}$, bijectively maps \mathcal{X}_N to the subset $\tilde{\mathcal{X}}_N \stackrel{\text{def}}{=} \{u \in \mathbb{R}^n : (\exists x_i \in \mathcal{X}_N) u = x_i - \bar{x}\}$ of $\overline{B_2(0, r)}$. In particular, since W is an isometry then

$$\text{aspect}(\mathcal{X}_N, \|\cdot\|_2) = \text{aspect}(\tilde{\mathcal{X}}_N, \|\cdot\|_2).$$

Since W is an affine map then, the pre-composition of any deep feedforward network by W yields a deep feedforward network of the same depth and width.

Step 2: Independently Memorizing Classes

By (Vardi et al., 2022, Theorem 3.1), for every $x \in \mathcal{X}_N$ there exists a feedforward network $\tilde{\mathcal{N}}_x : \mathbb{R}^n \rightarrow \mathbb{R}^d$ with ReLU activation function satisfying the “class memorization property”

$$\tilde{\mathcal{N}}_x(u) = \begin{cases} 1 & : \text{if } u = x - \bar{x} \\ 0 & : \text{if } u \neq x - \bar{x} \end{cases} \quad (18)$$

for every $u \in \tilde{\mathcal{X}}_N$. Furthermore, $\tilde{\mathcal{N}}_x$ has width 12 and depth

$$\mathcal{O} \left(\left(N \log(N) \right)^{1/2} + \left(\frac{N}{\log(N)} \right)^{1/2} \max\{\log(R), \log(C)\} \right), \quad (19)$$

where $C = 2$ and $R = 10\pi^{1/2} n^{1/2} N^2 r \delta^{-1}$ and where $\delta \stackrel{\text{def}}{=} \min_{1 \leq i, j \leq N; i \neq j} \|(x_i - \bar{x}) - (x_j - \bar{x})\|_2 = \min_{1 \leq i, j \leq N; i \neq j} \|x_i - x_j\|_2$. We may therefore rewrite (19) as

$$\mathcal{O} \left(\sqrt{N \log(N)} \left(1 + \frac{\log(2)}{\log(N)} \left[C_n + \frac{\log(N^2 \text{aspect}(\mathcal{X}_N, \|\cdot\|_2))}{\log(2)} \right]_+ \right) \right), \quad (20)$$

where the “dimensional” constant $C_n > 0$ is defined by $C_n \stackrel{\text{def}}{=} \frac{2 \log(5\sqrt{2\pi}) + \frac{3}{2} \log(n) - \frac{1}{2} \log(n+1)}{2 \log(2)}$ and where, for any $u \in \mathbb{R}$ we define $[u]_+ \stackrel{\text{def}}{=} \max\{u, 1\}$.

Step 3: Parallel Interpolating ReLU Network for each x in \mathcal{X}_N

For every $x \in \mathcal{X}_N$, we view the vector $f(x) \in \mathbb{R}^d$ as a $d \times 1$ matrix in $\mathbb{R}^{d \times 1}$. Now, since the composition

of affine maps is again affine and since the first and last layers of each $\tilde{\mathcal{N}}\mathcal{N}_x$ are affine maps then, for every $x \in \mathcal{X}_N$, the map $\mathcal{N}\mathcal{N}_x : \mathbb{R}^n \rightarrow \mathbb{R}^d$ defined by

$$\mathcal{N}\mathcal{N}_x(u) \stackrel{\text{def.}}{=} f(x) [\tilde{\mathcal{N}}\mathcal{N}_x \circ W^{-1}(u)], \quad (21)$$

is a ReLU feedforward network satisfying

1. $\text{depth}(\mathcal{N}\mathcal{N}_x) = \text{depth}(\tilde{\mathcal{N}}\mathcal{N}_x)$ where each $\text{depth}(\tilde{\mathcal{N}}\mathcal{N}_x)$ is as in (20),
2. $\text{width}(\mathcal{N}\mathcal{N}_x) = \max\{d, \text{width}(\tilde{\mathcal{N}}\mathcal{N}_x)\} = \max\{d, 12\}$,
3. $\text{par}(\mathcal{N}\mathcal{N}_x) \leq d + \text{par}(\tilde{\mathcal{N}}\mathcal{N}_x)$.

Moreover, by construction, for each $x \in \mathcal{X}_N$ it holds that

$$\mathcal{N}\mathcal{N}_x(u) = \begin{cases} f(x) & : \text{if } u = x \\ 0 & : \text{if } u \neq x, \end{cases} \quad (22)$$

for every $u \in \mathcal{X}_N$. Furthermore, we can bound the number of parameters $\text{par}(\mathcal{N}\mathcal{N}_x)$ defining each of the ReLU networks $\mathcal{N}\mathcal{N}_x$ above by

$$\begin{aligned} \text{par}(\mathcal{N}\mathcal{N}_x) &\leq d + \text{par}(\tilde{\mathcal{N}}\mathcal{N}_x) \\ &\leq \text{depth}(\tilde{\mathcal{N}}\mathcal{N}_{x_1}) \text{width}(\tilde{\mathcal{N}}\mathcal{N}_{x_1}) (\text{width}(\mathcal{N}\mathcal{N}_{x_1}) + 1) + d \\ &\leq \mathcal{O} \left(d + \sqrt{N \log(N)} \left(1 + \frac{\log(2)}{\log(N)} \left[C_n + \frac{\log(N^2 \text{aspect}(\mathcal{X}_N, \|\cdot\|_2))}{\log(2)} \right] \right)_+ \right) \times \\ &\quad \left(\max\{d, 12\} (\max\{d, 12\} + 1) \right). \end{aligned} \quad (23)$$

Step 4: Assembling the Parallel Interpolating ReLU Networks

Next, we assemble each of the ReLU networks $\{\mathcal{N}\mathcal{N}_x\}_{x \in \mathcal{X}_N}$ into a single ReLU network interpolating f on \mathcal{X}_N . Since $\text{ReLU}(x) \stackrel{\text{def.}}{=} \max\{0, x\}$ has the 2-identity property (as defined in (Cheridito et al., 2021, Definition 4) and discussed on page 3 of that manuscript) then, (Cheridito et al., 2021, Proposition 5) implies that there exists a “deep diagonalized parallelization”; that is, a map $\|\mathcal{N}\mathcal{N} : \mathbb{R}^n \rightarrow \mathbb{R}^d$ satisfying the following for every $u \in \mathbb{R}^n$

$$\|\mathcal{N}\mathcal{N}(u) = \begin{pmatrix} \mathcal{N}\mathcal{N}_{x_1}(u) \\ \vdots \\ \mathcal{N}\mathcal{N}_{x_N}(u) \end{pmatrix}. \quad (24)$$

Furthermore, (Cheridito et al., 2021, Proposition 5) shows that $\|\mathcal{N}\mathcal{N}$ has width, depth, and number of non-zero parameters are estimated by

1. **Width:** $\|\mathcal{N}\mathcal{N}(u)$ has width $n(N - 1) + \max\{d, 12\}$
2. **Depth:** $\|\mathcal{N}\mathcal{N}(u)$ has depth

$$\mathcal{O} \left(N \left(1 + \sqrt{N \log(N)} \left(1 + \frac{\log(2)}{\log(N)} \left[C_n + \frac{\log(N^2 \text{aspect}(\mathcal{X}_N, \|\cdot\|_2))}{\log(2)} \right] \right)_+ \right) \right), \quad (25)$$

3. **Number of non-zero parameters:** The number of non-zero parameters determining $\|\mathcal{N}\mathcal{N}$ is

$$\begin{aligned} &\left(\frac{11}{4} \max\{n, d\} N^2 - 1 \right) \sum_{i=1}^N \text{par}(\mathcal{N}\mathcal{N}_{x_i}). \text{ When expanded, this can be rewritten as} \\ &\mathcal{O} \left(N \left(\frac{11}{4} \max\{n, d\} N^2 - 1 \right) \left(d + \sqrt{N \log(N)} \left(1 + \frac{\log(2)}{\log(N)} \left[C_n + \frac{\log(N^2 \text{aspect}(\mathcal{X}_N, \|\cdot\|_2))}{\log(2)} \right] \right)_+ \right) \left(\max\{d, 12\} (\max\{d, 12\} + 1) \right) \right). \end{aligned}$$

Consider the $d \times (Nd)$ block-matrix $A \stackrel{\text{def.}}{=} (I_d, \dots, I_d)$. Multiplying the outputs of $\|\mathcal{NN}(u)$ on the left by the matrix A defines a map $\mathcal{NN}^* : \mathbb{R}^n \rightarrow \mathbb{R}^d$; i.e. for every $u \in \mathbb{R}^n$ the map \mathcal{NN}^* is defined by

$$\mathcal{NN}^*(u) \stackrel{\text{def.}}{=} A\|\mathcal{NN}(u). \quad (26)$$

Furthermore, since the post-composition of ReLU feedforward networks is again a ReLU feedforward network then \mathcal{NN}^* is a ReLU feedforward network. Moreover, by construction, the width, depth, and number of parameters defining \mathcal{NN}^* satisfy the same estimates as $\|\mathcal{NN}$, respectively.

Step 5: Verifying that \mathcal{NN}^* Implements f on \mathcal{X}_N

It remains only to verify that \mathcal{NN}^* equals to f on \mathcal{X}_N . Upon combining the identities in (22), (26), and (24) we conclude that: for each $x \in \mathcal{X}$ the following holds

$$\mathcal{NN}^*(x) = \sum_{i=1}^N \mathcal{NN}_{x_i} = \sum_{i=1}^N f(x_i)I_{x_i=x} = f(x). \quad (27)$$

Relabeling “ $\mathcal{NN} \stackrel{\text{def.}}{=} \mathcal{NN}^*$ ” yields the lemma’s statement. ■

We now move to the proof of the main memorization lemma, namely Lemma 17, which guarantees that probabilistic transformer networks can memorize “conditionally Gaussian” Markov kernels for any finite number of inputs. In the following, we make use of the finite-dimensional Fréchet embedding

$$\tilde{\iota} : (\mathcal{X}_n, d_n) \ni x \mapsto (d(x, \tilde{x}))_{\tilde{x} \in \mathcal{X}_n}. \quad (28)$$

Proof [Proof of Lemma 17] Set $L \stackrel{\text{def.}}{=} n$ (where L is as in (6)) and identify the space of $d \times d$ -matrices with Fröbenius norm $\mathbb{R}^{d \times d}$ with \mathbb{R}^{d^2} with Euclidean norm (*NB, these are isometric metric spaces*). Then, $\tilde{\iota} = \Phi_{\{x_i\}_{i=1}^n}$, where $\tilde{\iota}$ is the Fréchet isometric embedding defined in (28).

For each $k = 1, \dots, K$ define the maps $\tilde{\mu}_k \stackrel{\text{def.}}{=} \mu_k \circ \iota^{-1} : \iota(\mathcal{X}_n) \rightarrow \mathbb{R}^d$. By Lemma 20, for every $k = 1, \dots, K$ there exists a ReLU network $\hat{\mu}_k : \mathbb{R}^n \rightarrow \mathbb{R}^d$ satisfying the following

$$\hat{\mu}_k(x) = \tilde{\mu}_k(x) = \mu_k \circ \iota^{-1}(x), \quad (29)$$

for every $x \in \mathcal{X}_n$. The following equivalence of the norms $\|\cdot\|_\infty \leq \|\cdot\|_2 \leq \sqrt{n}\|\cdot\|_\infty$ and the fact that ι is an isometric embedding of (\mathcal{X}_n, d_n) into $(\mathbb{R}^n, \|\cdot\|_\infty)$ implies

$$\begin{aligned} \text{aspect}(\mathcal{X}_n, d_n) &\stackrel{\text{def.}}{=} \frac{\max_{1 \leq i, j \leq n} d_n(x_i, x_j)}{\min_{1 \leq k, s \leq n; k \neq s} d_n(x_k, x_s)} \\ &= \frac{\max_{1 \leq i, j \leq n} \|\iota(x_i) - \iota(x_j)\|_\infty}{\min_{1 \leq k, s \leq n; k \neq s} \|\iota(x_k) - \iota(x_s)\|_\infty} \\ &\leq \sqrt{n} \frac{\max_{1 \leq i, j \leq n} \|\iota(x_i) - \iota(x_j)\|_2}{\min_{1 \leq k, s \leq n; k \neq s} \|\iota(x_k) - \iota(x_s)\|_2} \\ &= \sqrt{n} \frac{\max_{u, \tilde{u} \in \iota(\mathcal{X}_n)} \|u - \tilde{u}\|_2}{\min_{v, \tilde{v} \in \iota(\mathcal{X}_n); v \neq \tilde{v}} \|v - \tilde{v}\|_2} \\ &\stackrel{\text{def.}}{=} \sqrt{n} \text{aspect}(\iota(\mathcal{X}_n), \|\cdot\|_2). \end{aligned} \quad (30)$$

Thus, Lemma 20 (i)-(iii) and (30) imply that for each $k = 1, \dots, K$ the ReLU network $\hat{\mu}_k$ satisfies

- (i) **Width:** \mathcal{NN} has width $n(n-1) + \max\{d, 12\}$,
- (ii) **Depth:** \mathcal{NN} has depth of the order of

$$\mathcal{O} \left(n \left(1 + \sqrt{n \log(n)} \left(1 + \frac{\log(2)}{\log(n)} \left[C_n + \frac{\log(n^{5/2} \text{aspect}(\mathcal{X}_n, d_n))}{\log(2)} \right]_+ \right) \right) \right),$$

(iii) **Number of non-zero parameters:** The number of non-zero parameters in \mathcal{NN} is at most

$$\mathcal{O}\left(n\left(\frac{11}{4}\max\{n,d\}n^2-1\right)\left(d+\sqrt{n\log(n)}\left(1+\frac{\log(2)}{\log(n)}\left[C_n+\frac{\log(n^{5/2}\text{aspect}(\mathcal{X}_n,d_n))}{\log(2)}\right]\right)\right)\left(\max\{d,12\}(\max\{d,12\}+1)\right)\right),$$

and where the ‘‘dimensional constant’’ $C_d > 0$ is defined by $C_d \stackrel{\text{def}}{=} \frac{2\log(5\sqrt{2\pi})+3\log(d)-\log(d+1)}{2\log(2)}$.

Next, for every $k = 1, \dots, K$, define the ReLU network $\hat{\Sigma}_k : \mathbb{R}^n \rightarrow \mathbb{R}^{d^2}$ by

$$\hat{\Sigma}_k(x) \stackrel{\text{def}}{=} \bar{0}_{d^2} \text{ReLU} \bullet \left(\bar{0}_{d^2 \times n} x + \bar{0}_{d^2} \right) + \bar{0}_{d^2}, \quad (31)$$

where $\bar{0}_{d^2 \times n}$ is the $d^2 \times n$ -matrix with all entries 0, $\bar{0}_{d^2} \in \mathbb{R}^{d^2}$ is the vector with all entries 0. NB, by construction

$$\max_{k=1,\dots,K} \text{par}(\hat{\Sigma}_k) = 0. \quad (32)$$

Similarly, define the ReLU network $\hat{w} : \mathbb{R}^n \rightarrow \mathbb{R}^K$ by

$$\hat{w}_k(x) \stackrel{\text{def}}{=} \bar{0}_K \text{ReLU} \bullet \left(\bar{0}_{K \times n} x + \bar{0}_K \right) + \bar{0}_K. \quad (33)$$

Combining (29), (31), and (33) with the fact that $\Phi_{\{x_l\}_{l=1}^N} = \tilde{t} = \tilde{t}|_{\tilde{t}(\mathcal{X}_n)}^{-1}$ we have that: for every $x \in \mathcal{X}_n$ the following holds

$$\begin{aligned} T(x) &\stackrel{\text{def}}{=} \sum_{k=1}^K [\sigma_K \circ \hat{w} \circ \Phi_{\{x_l\}_{l=1}^N}(x)]_k \times \\ &\quad N_d \left(\hat{\mu}_k \circ \Phi_{\{x_l\}_{l=1}^N}(x), \hat{\Sigma}_k \circ \Phi_{\{x_l\}_{l=1}^N}(x) \right) \\ &= \sum_{k=1}^K \sigma_K(0)_k N_d(\mu_k \circ \tilde{t}^{-1} \circ \tilde{t}(x), 0) \\ &= \sum_{k=1}^K \frac{e^0}{\sum_{j=1}^K e^0} N_d(\mu_k \circ \tilde{t}^{-1} \circ \tilde{t}(x), 0) \\ &= \sum_{k=1}^K \frac{1}{K} N_d(\mu_k \circ \tilde{t}^{-1} \circ \tilde{t}(x), 0) \\ &= \phi(x). \end{aligned}$$

It remains only to compute the width, depth, and number of non-zero parameters determining T . Since T has width

$$\max \left\{ \text{width}(\hat{w}), \max_{k=1,\dots,K} \text{width}(\hat{\Sigma}_k), \max_{k=1,\dots,K} \text{width}(\mu_k) \right\}.$$

Therefore, we deduce that

$$\begin{aligned} \text{width}(T) &= \max\{\max\{n, K\}, \max\{n, d^2\}, n(n-1) + \max\{d, 12\}\} \\ &= \max\{n, K, d^2, n(n-1) + \max\{d, 12\}\}, \end{aligned}$$

and T has depth $\max\{\text{depth}(\hat{w}), \max_{k=1,\dots,K} \text{depth}(\hat{\Sigma}_k), \max_{k=1,\dots,K} \text{depth}(\mu_k)\}$. Since $\text{depth}(\hat{w}) = \max_{k=1,\dots,K} \text{depth}(\hat{\Sigma}_k) = 1$ then, T has depth of the following order

$$\mathcal{O}\left(n\left(1+\sqrt{n\log(n)}\left(1+\frac{\log(2)}{\log(n)}\left[C_n+\frac{\log(n^{5/2}\text{aspect}(\mathcal{X}_n,d_n))}{\log(2)}\right]\right)\right)\right).$$

Furthermore, the number of parameters defining T equals

$$\begin{aligned} \text{par}(T) &\stackrel{\text{def.}}{=} \text{par}(\hat{w}) + \sum_{k=1}^K \left[\text{par}(\hat{\mu}_k) + \text{par}(\hat{\Sigma}_k) \right] \\ &= 0 + K \left[0 + \text{par}(\hat{\Sigma}_1) \right]. \end{aligned}$$

Thus, the number of non-zero parameters defining T is

$$\mathcal{O} \left(Kn \left(\frac{11}{4} \max\{n, d\} n^2 - 1 \right) \left(d + \sqrt{n \log(n)} \left(1 + \frac{\log(2)}{\log(n)} \left[C_n + \frac{\log(n^{5/2} \text{aspect}(\mathcal{X}_n, d_n))}{\log(2)} \right] \right)_+ \max\{d, 12\}^2 \right) \right).$$

■

We now turn our attention to the proof of Theorem 5.

B.1 Proof of PAC-Type Representation Theorem for “Unstructured Case”

The following argument uses the notation of an *ultrametric space*; briefly, an ultrametric space $(\mathcal{Z}, d_{\mathcal{Z}})$ is a metric space which satisfies the *strong triangle inequality*: for all $z_1, z_2, z_3 \in \mathcal{Z}$ one has $d_{\mathcal{Z}}(z_1, z_2) \leq \max\{d_{\mathcal{Z}}(z_1, z_2), d_{\mathcal{Z}}(z_2, z_3)\}$. Note that, $\max\{d_{\mathcal{Z}}(z_1, z_2), d_{\mathcal{Z}}(z_2, z_3)\} \leq d_{\mathcal{Z}}(z_1, z_2) + d_{\mathcal{Z}}(z_2, z_3)$ so the strong triangle inequality implies the familiar triangle inequality.

Proof [Proof of Theorem 5] Let $n \geq 2$ and fix an n -point metric subset \mathcal{X}_n of \mathcal{X} . Denote $(\mathcal{X}_n, d_n) \stackrel{\text{def.}}{=} (\mathcal{X}_n, d_{\mathcal{X}})$.

Step 1 - Embedding of a Dvoretzky-Type Subspace into Low-Dimensional Euclidean Space

Fix $\epsilon \in (0, \infty]$ and let (\mathcal{X}_n, d_n) be an n -point metric space. By (Naor and Tao, 2012, Theorem 1.2) there exists a $\theta_\epsilon \in (0, 1)$ (depending only on ϵ) and subset $\tilde{\mathcal{X}}_n \subseteq \mathcal{X}_n$ of cardinality

$$\#\tilde{\mathcal{X}}_n \geq n^{\theta_\epsilon} \geq n^{1 - \frac{2\epsilon}{2+\epsilon}}, \quad (34)$$

with the following property: there exists a separable ultrametric space $(\mathcal{Z}, d_{\mathcal{Z}})$ and a bi-Lipschitz map $\varphi_1 : (\tilde{\mathcal{X}}_n, d_n) \hookrightarrow (\mathcal{Z}, d_{\mathcal{Z}})$ such that for every $x, \tilde{x} \in \tilde{\mathcal{X}}_n$ it holds that

$$d_n(x, \tilde{x}) \leq d_{\mathcal{Z}}(\varphi_1(x), \varphi_1(\tilde{x})) \leq (2 + \epsilon) d_n(x, \tilde{x}). \quad (35)$$

Since $(\mathcal{Z}, d_{\mathcal{Z}})$ is a separable ultrametric space then (Vestfrid and Timan, 1979) (or (Shkarin, 2004)) implies that there exists an isometric embedding $\varphi_2 : (\mathcal{Z}, d_{\mathcal{Z}}) \hookrightarrow (\ell_2^\infty, \ell_2)$, where the norm $\|\cdot\|_{\ell_2}$ is defined for any $x \in \mathbb{R}^{\mathbb{N}}$ by $\|x\|_{\ell_2}^2 \stackrel{\text{def.}}{=} \sum_{k=1}^\infty x_k^2$, $\ell_2^\infty \stackrel{\text{def.}}{=} \{x \in \mathbb{R}^{\mathbb{N}} : \|x\|_{\ell_2} < \infty\}$, and thus (ℓ_2^∞, ℓ_2) the infinite dimensional separable Hilbert space. Define $\varphi_3 \stackrel{\text{def.}}{=} \varphi_2 \circ \varphi_1$. It follows from (35) that: for all $x, \tilde{x} \in \tilde{\mathcal{X}}_n$

$$d_n(x, \tilde{x}) \leq \|\varphi_3(x) - \varphi_3(\tilde{x})\|_{\ell_2} \leq (2 + \epsilon) d_n(x, \tilde{x}). \quad (36)$$

By the Johnson-Lindenstrauss Flattening Theorem (Matoušek, 2002, Theorem 15.2.1) there exist an $s > 0$ and a map $\varphi_4 : (\ell_2^\infty, \ell_2) \rightarrow (\ell_2^{n_\epsilon}, \ell_2)$, where $n_\epsilon \in \mathcal{O}(\epsilon^{-2} \log_2(\#\tilde{\mathcal{X}}_n)) = \mathcal{O}(\epsilon^{-2} \theta_\epsilon \log_2(n))$ (thus n_ϵ is at-most $\mathcal{O}(\epsilon^{-2} \log_2(n))$ since $\theta_\epsilon < 1$) satisfying

$$s \|z - \tilde{z}\|_{\ell_2} \leq \|\varphi_4(z) - \varphi_4(\tilde{z})\|_{\ell_2} \leq s(1 + \epsilon) \|z - \tilde{z}\|_{\ell_2}, \quad (37)$$

for all $z, \tilde{z} \in \varphi_3(\tilde{\mathcal{X}}_n)$. Set $\varphi_5 \stackrel{\text{def.}}{=} \varphi_4 \circ \varphi_3 : (\tilde{\mathcal{X}}_n, d_n) \hookrightarrow (\mathbb{R}^{n_\epsilon}, \ell_2)$. Together (36) and (37) imply that: for every $x, \tilde{x} \in \tilde{\mathcal{X}}_n$ it holds that

$$s d_n(x, \tilde{x}) \leq \|\varphi_5(x) - \varphi_5(\tilde{x})\|_{\ell_2} \leq s(1 + \epsilon)(2 + \epsilon) d_n(x, \tilde{x}). \quad (38)$$

Step 2 - Constructing The Embedding into $(\mathcal{GM}(\mathbb{R})_2, \mathcal{MW}_2)$

Since $\tilde{\mathcal{X}}_n$ is finite, then $(\tilde{\mathcal{X}}_n, d_n)$ is compact. Since φ_5 is a bi-Lipschitz map then it is continuous. Therefore, by (Munkres, 2000, Theorem 26.5) $\varphi_5(\tilde{\mathcal{X}}_n)$ is a compact subset of $(\mathbb{R}^{n_\epsilon}, \ell_2)$. By the Heine-Borel theorem (Munkres, 2000, Theorem 27.3), there exists an $r > 0$ such that $\varphi_5(\tilde{\mathcal{X}}_n) \subseteq \overline{\text{Ball}_{\mathbb{R}^{n_\epsilon}, \ell_2}(0, r)}$; where, $\overline{\text{Ball}_{\mathbb{R}^{n_\epsilon}, \ell_2}(0, r)} \stackrel{\text{def.}}{=} \{u \in \mathbb{R}^{n_\epsilon} : \|u\|_{\ell_2} \leq r\}$. Since $\varphi_5(\tilde{\mathcal{X}}_n) \subseteq \overline{\text{Ball}_{\mathbb{R}^{n_\epsilon}, \ell_2}(0, r)}$ then (Kloeckner, 2010, Proposition 7.4) applies. More specifically, as formulated in (Bertrand and Kloeckner, 2012, Example 5.5) there exists some $b \in \mathbb{R}^{n_\epsilon}$ (which can be computed explicitly using Algorithm 1) such that the map

$$\iota : \overline{\text{Ball}_{\mathbb{R}^{n_\epsilon}, \ell_2}(0, r)} \ni x \mapsto \frac{1}{n_\epsilon} \sum_{k=1}^{n_\epsilon} \delta_{\sqrt{n_\epsilon}(x+b)_k} \in (\mathcal{P}_2(\mathbb{R}), \mathcal{W}_2),$$

is an isometry. Since the composition of isometries is itself an isometry then, the map $\varphi \stackrel{\text{def.}}{=} \iota \circ \varphi_5 : (\tilde{\mathcal{X}}_n, d_n) \hookrightarrow (\mathcal{P}_2(\mathbb{R}), \mathcal{W}_2)$ is an isometry and is of the required form. Lastly, since φ has range in the set of “degenerate Gaussian mixtures” $\mathcal{Z} \stackrel{\text{def.}}{=} \{\mathbb{P} \in \mathcal{GM}_2(\mathbb{R}) : \exists N \in \mathbb{N}_+ \exists w \in \Delta_N \exists \mu_1, \dots, \mu_N \in \mathbb{R} \text{ s.t. } \mathbb{P} = \sum_{n=1}^N w_n \delta_{\mu_n}\}$ then, we obtain the conclusion by applying the comparison result in (Delon and Desolneux, 2020, Proposition 6) which states that the “identity map” on $(\mathcal{W}_2(\mathbb{R}), \mathcal{W}_2) \ni \mathbb{P} \mapsto \mathbb{P} \in (\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ is an isometry on the collection of finitely supported probability measures (i.e. “degenerate Gaussian mixtures”). Applying Lemma 17 to the map $\phi : (\mathcal{X}_n, d_n) \rightarrow (\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$, defined for each $x \in \mathcal{X}_n$ by $\phi(x) \stackrel{\text{def.}}{=} \varphi \circ [x \mapsto (d(x, \tilde{x})_{\tilde{x} \in \tilde{\mathcal{X}}_n})]$, yields the desired transformer T satisfying

$$sd_n(x, \tilde{x}) \leq \mathcal{MW}_2(T(x), T(\tilde{x})) \leq s(2 + \epsilon)^2 d_n(x, \tilde{x}),$$

for all $x, \tilde{x} \in \tilde{\mathcal{X}}_n$.

Comment: From here, it now only remains to estimate the probability that a uniformly and independently chosen pair of random points in \mathcal{X}_n lie in $\tilde{\mathcal{X}}_n$.

Step 3 - Estimating Probability of Independently and Uniformly Sampling two Points in the Dvoretzky-type Subspace

Let \mathbb{P} be the uniform probability distribution on \mathcal{X}_n and let $X, \tilde{X} \sim \mathbb{P}$ be independent \mathcal{X}_n -valued random elements. By independence and the fact that X and \tilde{X} have the same law we can compute the following

$$\mathbb{P}(X \in \tilde{\mathcal{X}}_n \text{ and } \tilde{X} \in \tilde{\mathcal{X}}_n) = \mathbb{P}(X \in \tilde{\mathcal{X}}_n) \mathbb{P}(\tilde{X} \in \tilde{\mathcal{X}}_n) = \mathbb{P}(X \in \tilde{\mathcal{X}}_n)^2. \quad (39)$$

Together, the lower-bound in (34), the definition of \mathbb{P} , and (39) imply that

$$\begin{aligned} \mathbb{P}(X \in \tilde{\mathcal{X}}_n \text{ and } \tilde{X} \in \tilde{\mathcal{X}}_n) &= \mathbb{P}(X \in \tilde{\mathcal{X}}_n)^2 \\ &= \left(\frac{\#\tilde{\mathcal{X}}_n}{\#\mathcal{X}_n} \right)^2 \\ &\geq \left(\frac{n^{\theta_\epsilon}}{n} \right)^2 \\ &\geq \left(\frac{n^{1-(2\epsilon/2+\epsilon)}}{n} \right)^2 \\ &= n^{-4\epsilon/(2+\epsilon)}. \end{aligned} \quad (40)$$

Moreover, by (Naor and Tao, 2012, Theorem 1.2) the precise value of θ_ϵ is given by the unique solution $\theta \in (0, 1)$ to $\frac{2}{2+\epsilon} = (1-\theta)\theta^{\theta/(1-\theta)}$; nb, $\theta_\epsilon \in [1 - 2\epsilon/(2+\epsilon), 1)$. Labeling $D \stackrel{\text{def.}}{=} 2 + \epsilon$, $\theta_D \stackrel{\text{def.}}{=} \theta_\epsilon$, observing that T is defined on all of \mathcal{X} , and noting that by construction $d_T = k$ yields the theorem’s statement and concludes our proof.

Step 4 - Inverting the relation between D and θ_D :

Fix $\delta \in (0, 1)$ such that $n^{-2+2\theta_D} = \delta$. Solving for θ_D yields

$$\theta_D = 1 + \log_n(\sqrt{\delta}). \quad (41)$$

Since θ_D must belong to $(0, 1)$ then (41) only gives a valid θ_D if $\delta \in (1/e^2, 1)$; thus, we now impose that constraint. Since θ_D solves $\frac{2}{D} = (1 - \theta)\theta^{\theta/(1-\theta)}$ where $\theta \in [1 - \frac{2e}{D}, 1)$ then (41) implies that

$$D = -2 \frac{(1 + \log_n(\sqrt{\delta}))^{\frac{1 + \log_n(\sqrt{\delta})}{\log_n(\sqrt{\delta})}}}{\log_n(\sqrt{\delta})}. \quad (42)$$

In order for D in (42) to be a valid distortion (compatible with our argument) we need that $D > 2$. Therefore, we require that $\delta > \delta_n$; where $\tilde{\delta}_n \in [\frac{1}{e^2}, 1)$ is the unique solution to

$$\log_n(\sqrt{\tilde{\delta}_n}) = (1 + \log_n(\sqrt{\tilde{\delta}_n}))^{\frac{1 + \log_n(\sqrt{\tilde{\delta}_n})}{\log_n(\sqrt{\tilde{\delta}_n})}}.$$

Set $\delta_n \stackrel{\text{def}}{=} \max\{\frac{1}{e^2}, \tilde{\delta}_n\}$. Restricting $\delta \in (\delta_n, 1)$ yields a δ for which θ_D and D both satisfy the required assumptions of our argument. This complete the proof. \blacksquare

B.2 Proof of Distortionless Multivariate Embeddings

Proof [Proof of Theorem 7] Fix a positive integer n and let \mathcal{X}_n be a n -point subset of a metric space (\mathcal{X}, d) which we view as an n -point metric space (\mathcal{X}_n, d) . Let $\iota : (\mathcal{X}_n, d) \rightarrow (\mathbb{R}^n, \|\cdot\|_\infty)$ be the Fréchet embedding defined by $\iota(x) \stackrel{\text{def}}{=} (d(x, x_i))_{i=1}^n$ and denote $\tilde{\mathcal{X}}_n \stackrel{\text{def}}{=} \iota(\mathcal{X}_n)$. Clearly, (\mathcal{X}_n, d) and $(\tilde{\mathcal{X}}_n, \|\cdot\|_\infty)$ are isometric with isometry given by ι . It is therefore sufficient to bi-Hölder embed $(\mathcal{X}_n, \|\cdot\|_\infty)$ into $(\mathcal{P}_2(\mathbb{R}^3), \mathcal{W}_2)$; we do this now. In the proof of (Andoni et al., 2018, Theorem 1) from (Andoni et al., 2018, Equation (12) to Equation(22)), the authors define a map $\tilde{\phi} : \tilde{\mathcal{X}}_n \rightarrow (\mathcal{P}_2(\mathbb{R}^3), \mathcal{W}_2)$ with the property that: there is some positive integer N such that for every $x \in \tilde{\mathcal{X}}_n$

$$\tilde{\phi}(x) = \frac{1}{N} \sum_{n=1}^N \delta_{u_n(x)} \quad (43)$$

for some set of points $\{u_n(x) : n = 1, \dots, N, x \in \tilde{\mathcal{X}}_n\}$ in \mathbb{R}^3 . Moreover, a careful reading of the construction show that the positive integer N is such bounded above by¹⁶

$$N \leq n(n-1)/2(5K+2) \quad (44)$$

for some hyperparameter $K \in \mathbb{N}_+$ to be decided upon shortly. The quantitative version of the result in question, namely (Andoni et al., 2018, Lemma 15), shows that for every distortion $D > 1$, taking $\frac{5n^4}{2(D-1)}$ $\text{aspect}(\tilde{\mathcal{X}}_n, d)^2 \leq K$ implies that: there is some scale $s > 0$ such that for every $x, \tilde{x} \in \tilde{\mathcal{X}}_n$ it holds that

$$s \|x - \tilde{x}\|_\infty^{1/2} \leq \mathcal{W}_2(\tilde{\phi}(x), \tilde{\phi}(\tilde{x})) \leq D s \|x - \tilde{x}\|_\infty^{1/2}. \quad (45)$$

Since ι is an isometry and every pointmass on \mathbb{R}^3 is a Gaussian measure with mean at that point and zero variance then the map $\phi \stackrel{\text{def}}{=} \tilde{\phi} \circ \iota^{-1}$ sends (\mathcal{X}_n, d) to $(\mathcal{GM}_2(\mathbb{R}^3), \mathcal{MW}_2)$ and it has representation

$$\phi(x) = \sum_{n=1}^N \frac{1}{N} N_3(u_n(x), \mathbf{0}),$$

16. As a brief reference, we note that, in the notation of (Andoni et al., 2018), $N \leq \#\mathcal{C}$, $\mathcal{C} \subseteq \cup_{i,j=1,\dots,n; i < j} \mathcal{B}_{i,j}$ and $\#\mathcal{B}_{i,j} = 5K + 2$; for some hyperparameter $K \in \mathbb{N}_+$ independent of the construction thus far.

where $\mathbf{0}$ is the 3×3 -dimensional zero matrix. Moreover, since ι is an isometry and by (Delon and Desolneux, 2020, Proposition 6) the inclusion of the set of finitely supported probability measures in $(\mathcal{P}_2(\mathbb{R}^3), \mathcal{W}_2)$ into $(\mathcal{GM}_2(\mathbb{R}^3), \mathcal{MW}_2)$ is an isometry then (45) implies that: there is a scale $s > 0$ such that for every $x, \tilde{x} \in \mathcal{X}_n$ it holds that

$$s d(x, \tilde{x})^{1/2} \leq \mathcal{MW}_2(\phi(x), \phi(\tilde{x})) \leq D s d(x, \tilde{x})^{1/2}. \quad (46)$$

That is, ϕ is a $\frac{1}{2}$ -bi-Hölder embedding of (\mathcal{X}_n, d) into $(\mathcal{GM}_2(\mathbb{R}^3), \mathcal{MW}_2)$. Applying the Key Lemma, Lemma 17, and to the map ϕ yields the conclusion; upon noting that the probabilistic transformer T is defined on all of \mathcal{X} . \blacksquare

B.3 Proof of Embedding Lemmata

Proof [Proof of Lemma 18] Since (\mathcal{X}, d) is a doubling metric space then, by (Naor and Neiman, 2012, Theorem 1.2) there exist constant $c, C > 0$ such that for any $0 < \epsilon < \frac{1}{2}$ there exists an injective function $\phi : (\mathcal{X}, d) \hookrightarrow (\mathbb{R}^D, \ell_2)$ satisfying

$$d_{\mathcal{X}}^{1-\epsilon}(x, \tilde{x}) \leq \|\phi(x) - \phi(\tilde{x})\|_{\ell_2} \leq \left[\frac{C \log(C_d)^2}{(1 - (1 - \epsilon))^2} \right] d_{\mathcal{X}}^{1-\epsilon}(x, \tilde{x}); \quad (47)$$

where $D \leq c \log(C_d)$. Since $\mathcal{X}_n \subseteq \mathcal{X}$ is finite, then \mathcal{X}_n is compact. By (47), ϕ is $(\frac{1}{2}, 1) \ni (1 - \epsilon)$ -Hölder continuous and therefore, by (Munkres, 2000, Theorem 26.5) $\phi(\mathcal{X}_n)$ must be a compact subset of (\mathbb{R}^D, ℓ_2) . As before, by appealing to the Heine-Borel theorem ((Munkres, 2000, Theorem 27.3)) we conclude that there must exist some $r > 0$ such that $\phi(\mathcal{X}_n) \subseteq \overline{\text{Ball}_{\mathbb{R}^D, \ell_2}(0, r)}$; where $\overline{\text{Ball}_{\mathbb{R}^D, \ell_2}(0, r)} \stackrel{\text{def.}}{=} \{u \in \mathbb{R}^D : \|u\|_{\ell_2} \leq r\}$. Thus, (Kloeckner, 2010, Proposition 7.4) implies that there exists some $b \in \mathbb{R}^D$ (which can be constructed using Algorithm 1) such that the map

$$\iota : (\overline{\text{Ball}_{\mathbb{R}^D, \ell_2}(0, r)}, \ell_2) \ni x \mapsto \frac{1}{D} \sum_{k=1}^D \delta_{\sqrt{D}(x+b)_k} \in (\mathcal{P}_2(\mathbb{R}), \mathcal{W}_2),$$

is an isometry. As in the proof of Lemma 24, since $\iota \circ \phi$ has range in the set of “degenerate Gaussian mixtures”; i.e. belonging to $\mathcal{Z} \stackrel{\text{def.}}{=} \{\mathbb{P} \in \mathcal{GM}_2(\mathbb{R}^d) : \exists N \in \mathbb{N}_+ \exists w \in \Delta_N \exists \mu_1, \dots, \mu_N \in \mathbb{R}^d \text{ s.t. } \mathbb{P} = \sum_{n=1}^N w_n \delta_{\mu_n}\}$ then, we obtain the conclusion by applying the comparison result in (Delon and Desolneux, 2020, Proposition 6) which states that the identity map on $(\mathcal{GM}_2(\mathbb{R}^d), \mathcal{W}_2) \ni \mathbb{P} \mapsto \mathbb{P} \in (\mathcal{GM}_2(\mathbb{R}^d), \mathcal{MW}_2)$ is an isometry on the collection of degenerate mixtures of in \mathcal{Z} . Setting $\phi \stackrel{\text{def.}}{=} \iota \circ \phi$ and $\alpha \stackrel{\text{def.}}{=} 1 - \epsilon$ concludes the proof. \blacksquare

B.4 Proof of Main Results

The proof of Theorem 4 relies on the following lemma relating the rate at which the *volume of a ball* doubles as a function of its radius to \mathcal{X}_n 's metric capacity. To formulate the required lemma, we must first revisit the notion of a *doubling measure* (Heinonen, 2001, Page 3 and Chapter 13).

Let us begin by recalling that, given any $x \in \mathcal{X}_n$ and any $r > 0$, the metric ball about x of radius r is defined by $B(x, r) \stackrel{\text{def.}}{=} \{u \in \mathcal{X}_n : d_{\mathcal{X}}(u, x) < r\}$. Briefly, we say that a (finite Borel) measure μ on (\mathcal{X}_n, d_n) is a *doubling measure* with doubling constant c_μ if for every $x \in \mathcal{X}_n$ and every $r > 0$ it holds that

$$\mu(B(x, 2r)) \leq c_\mu \mu(B(x, r)). \quad (48)$$

In (Vol'berg and Konyagin, 1987) it was shown that the existence of a doubling measure is equivalent to d_n being *doubling as a metric space*; this means that there is a (metric) doubling constant $c_{d_n} > 0$ such that: for each $r > 0$ every metric ball in \mathcal{X}_n of radius $2r$ can be covered by at most c_{d_n} metric balls of radius r . In (Bruè et al., 2021, Proposition 1.7), it is shown that the existence of a bound on $\text{cap}(\mathcal{X}_n)$ is equivalent

to the existence and a bound for c_{d_n} . Thus, a bound on the constant of a doubling measure c_μ should, in principle, imply a bound on the $\text{cap}(\mathcal{X}_n)$. The next explicitly spells out the relationship between these “metric quantities”.

Lemma 21 (Bounds of Metric Capacity via (Metric and Measure) Doubling Constants) *Let μ be a doubling measure on (\mathcal{X}_n, d_n) with doubling constant c_μ , as in (48). Then, the metric capacity of (\mathcal{X}_n, d_n) is bounded above by*

$$\log(c_{d_n}) \leq \log(\text{cap}(\mathcal{X}_n)) \leq 12 \log(c_\mu).$$

Proof [Proof of Lemma 21] Suppose that N is a positive integer for which there does not exist points $x_1, \dots, x_N \in B(x, r)$ satisfying

$$\text{Ball}_g(x, r) \subseteq \bigcup_{i=1}^N \text{Ball}_g(x_i, r/2),$$

then, N is an upper-bound for c_{d_n} . Moreover, there must exist N points x_1, \dots, x_N in $B(x, r)$ for which

$$0 < \min_{i,j=1,\dots,N; i \neq j} \min_{u \in \mathcal{X}_n, d_n(u, x_i) \leq 2^{-2}r} \min_{\tilde{u} \in \mathcal{X}_n, d_n(\tilde{u}, x_j) \leq 2^{-2}r} d_n(u, \tilde{u}).$$

Pick $N^* \in \underset{i=1,\dots,N}{\text{argmin}} \mu(B(x_i, r))$. Thus, we have that

$$\begin{aligned} \mu(B(x_{N^*}, 4r)) &\geq \sum_{i=1}^N \mu(B(x_i, 2^{-2}r)) \\ &\geq n \min_{i=1,\dots,N} \mu(B(x_i, r)) = n\mu(B(x_{N^*}, r)). \end{aligned}$$

Therefore, $c_d \leq c_\mu^4$. Now, by (Bruè et al., 2021, Proposition 1.7 (i)), we have that

$$\text{cap}(\mathcal{X}_n) \leq C_d^3. \tag{49}$$

Combining the estimate (49) with the estimate $c_d \leq c_\mu^4$ yields the desired upper-bound on $\text{cap}(\mathcal{X}_n)$. The lower-bound on $\text{cap}(\mathcal{X}_n)$ is given in (Bruè et al., 2021, Proposition 1.7 (i)). ■

Proof [Proof of Theorem 4] Fix an $n \geq 2$ -point metric subspace (\mathcal{X}_n, d_n) of $(\mathcal{X}, d_{\mathcal{X}})$. Restricting to (\mathcal{X}_n, d_n) , the result for $\mathcal{X} = \mathcal{X}_n$ follows from Lemma 17 applied to Lemma 18 and then using Lemma 21 to upper-bound $\log_2(c_{d_n})$ by $12 \log_2(\text{cap}(\mathcal{X}_n, d_n))$. The general statement, follows by simply noting that T constructed this way is defined on all of \mathcal{X} . ■

Remark 22 *The proof of Theorem 4 implies that the result holds with $\text{cap}(\mathcal{X}_n, d_n)$ replaced by \mathcal{X}_n 's doubling constant. Nevertheless, one can still derive the result with C_d^3 in place of $\text{cap}(\mathcal{X}_n)$ as a consequence of the present formulation of Theorem 4 and of (Bruè et al., 2021, Proposition 1.7 (ii)). Thus, we only make the remark to state that the result can be sharpened if one prefers to consider doubling constants rather than metric capacity.*

B.5 Proofs of Secondary Results

B.5.1 PROOFS OF DETERMINISTIC EMBEDDINGS

Let us consider the case where (\mathcal{X}_n, d_n) is a tree.

Lemma 23 (Bi-Lipschitz Embedding of (Finite) Combinatorial Trees into $(\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$) Let $G = (V, E)$ be an n -vertex tree and let (V, d_G) be its associated metric space (as in Example 2.1). There is an absolute constant $C > 0$ such that, for every $M \in \mathbb{N}_+$ there exists a map of the form

$$\varphi(x) \stackrel{\text{def.}}{=} \frac{1}{M} \sum_{k=1}^M N_d(\phi(x)_k, 0),$$

such that the following holds: for every $x, \tilde{x} \in \mathcal{X}_n$ we have that

$$d_G(x, \tilde{x}) \leq \mathcal{W}_2(\varphi(x), \varphi(\tilde{x})) \leq \mathcal{MW}_2(\varphi(x), \varphi(\tilde{x})) \leq Cn^{1/(d-1)}d_G(x, \tilde{x}),$$

where $\phi : V \rightarrow \mathbb{R}^M$ and $C > 0$ is an absolute constant independent of n , d , and of (\mathcal{X}_n, d_n) .

Proof [Proof of Lemma 23] Fix $\epsilon > 0$. Since d_n is a tree metric on \mathcal{X}_n then (Gupta, 2000, Theorem 1) implies there is an absolute constant $C > 0$ (independent of n and of (\mathcal{X}_n, d_n)) and bi-Lipschitz embedding $\phi : V \rightarrow \mathbb{R}^d$ satisfying: for every $x, \tilde{x} \in V$

$$d_n(x, \tilde{x}) \leq \|\phi(x) - \phi(\tilde{x})\| \leq Cn^{1/(d-1)}d_n(x, \tilde{x}). \quad (50)$$

Since \mathcal{X}_n is finite, then (\mathcal{X}_n, d_n) is compact. Since ϕ is a bi-Lipschitz map then, it is continuous and therefore by (Munkres, 2000, Theorem 26.5), $\phi(\mathcal{X}_n)$ is a compact subset of (\mathbb{R}^d, ℓ_2) . By the Heine-Borel theorem ((Munkres, 2000, Theorem 27.3)), there exists an $r > 0$ such that $\phi(\mathcal{X}_n) \subseteq \overline{\text{Ball}_{\mathbb{R}^d, \ell_2}(0, r)}$; where $\overline{\text{Ball}_{\mathbb{R}^d, \ell_2}(0, r)} \stackrel{\text{def.}}{=} \{u \in \mathbb{R}^d : \|u\|_{\ell_2} \leq r\}$. Since $\phi(\mathcal{X}_n) \subseteq \overline{\text{Ball}_{\mathbb{R}^d, \ell_2}(0, r)}$ then (Kloeckner, 2010, Proposition 7.4) applies. More specifically, the proof of (Kloeckner, 2010, Proposition 7.4) (see the comment in (Bertrand and Kloeckner, 2012, Example 5.5)) guarantee that there exists some $b \in \mathbb{R}^d$ (which can be constructed using Algorithm 1 in Section C) such that the map

$$\iota : \overline{\text{Ball}_{\mathbb{R}^d, \ell_2}(0, r)} \ni x \mapsto \frac{1}{d} \sum_{k=1}^d \delta_{\sqrt{d}(x+b)_k} \in (\mathcal{P}_2(\mathbb{R}), \mathcal{W}_2),$$

is an isometry. Since the composition of isometries is itself an isometry, then the map $\varphi \stackrel{\text{def.}}{=} \iota \circ \phi : (\mathcal{X}_n, d_n) \hookrightarrow (\mathcal{P}_2(\mathbb{R}), \mathcal{W}_2)$ is an isometry and is of the required form. Lastly, since $\iota \circ \phi$ has range in the set of “degenerate Gaussian mixtures” $\mathcal{Z} \stackrel{\text{def.}}{=} \{\mathbb{P} \in \mathcal{GM}_2(\mathbb{R}^d) : \exists N \in \mathbb{N}_+ \exists w \in \Delta_N \exists \mu_1, \dots, \mu_N \in \mathbb{R}^d \text{ s.t. } \mathbb{P} = \sum_{n=1}^N w_n \delta_{\mu_n}\}$ then, we obtain the conclusion by applying the comparison result in (Delon and Desolneux, 2020, Proposition 6) which states that the “identity map” on $(\mathcal{W}_2(\mathbb{R}), \mathcal{W}_2) \ni \mathbb{P} \mapsto \mathbb{P} \in (\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ is an isometry on the collection of finitely supported probability measures (i.e. “degenerate Gaussian mixtures”).

■

Proof [Proof of Proposition 11] Since (\mathcal{X}_n, d_n) is a combinatorial (metric) graph then $\text{aspect}(\mathcal{X}_n, d_n) \leq \text{diam}(\mathcal{X}_n, d_n)$. The result now follows directly from Lemma 17 applied to Lemma 23. ■

Our next case concerns the structure of a low-distortion embedding of a graph whose points are drawn from a Riemannian manifold. In this case, we see that the number of pointmasses required scales quadratically in the dimension of the latent Riemannian manifold.

Lemma 24 (Bi-Lipschitz Embeddings of Riemannian Datasets - For Compact Manifolds)

Let $n, d \in \mathbb{N}_+$. Let (M, g) be a d -dimensional Riemannian manifold with geodesic distance d_g and let $K \subseteq M$ be compact. There exists a constant C_K (depending only on K) such that for any $\mathcal{X}_n \subseteq M$ if $d_n \stackrel{\text{def.}}{=} d_g|_{\mathcal{X}_n}$, then exists a map $\varphi : (\mathcal{X}_n, d_n) \rightarrow (\mathcal{GM}_2(\mathbb{R}), \mathcal{MW}_2)$ of the form

$$\varphi(x) = \frac{1}{2d} \sum_{k=1}^{2d} N_1(\phi(x)_k, 0);$$

satisfying: for every $x, \tilde{x} \in \mathcal{X}$ the following holds

$$d_{\mathcal{X}}(x, \tilde{x}) \leq \mathcal{W}_2(\varphi(x), \varphi(\tilde{x})) \leq \mathcal{M}\mathcal{W}_2(\varphi(x), \varphi(\tilde{x})) \leq C C_K d_{\mathcal{X}}(x, \tilde{x});$$

where $\phi : (\mathcal{X}_n, d_n) \rightarrow (\mathbb{R}^{2d}, \ell_2)$ is a class- C^1 -(Riemannian) isometric embedding.

Proof [Proof of Lemma 24] Since (M, g) is a Riemannian manifold then, (Nash, 1954, Theorem 2) implies that there exists a class- C^1 (Riemannian isometric embedding) $\phi : (M, d_g) \hookrightarrow (\mathbb{R}^{2d}, \ell_2)$. Since ϕ is a class C^1 Riemannian isometric embedding, it is a diffeomorphism onto its image and therefore, ϕ has a class C^1 inverse on $\phi(M)$ which we denote by ϕ^{-1} . Since K is compact and ϕ is continuous then $\phi(K)$ is compact and therefore the following is finite

$$C_K \stackrel{\text{def}}{=} \max\{\max_{x \in K} \|\nabla \phi(x)\|, \max_{x \in \phi(K)} \|\nabla \phi^{-1}(x)\|\}.$$

Therefore, by the Rademacher-Stepanov Theorem (Federer, 1978, Theorem 3.1.6) we have that: for every $x, \tilde{x} \in K$

$$d_g(x, \tilde{x}) \leq \|\phi(x) - \phi(\tilde{x})\| \leq C_K d_g(x, \tilde{x}). \quad (51)$$

Since ϕ is class- C^1 , it is continuous and since \mathcal{X}_n is a compact subset of M (since it is finite) then by (Munkres, 2000, Theorem 26.5) $\phi(\mathcal{X}_n)$ is a non-empty compact subset of $(\mathbb{R}^{2d}, \ell_2)$. By the Heine-Borel Theorem (Munkres, 2000, Theorem 27.3), there exists some $r > 0$ such that $\phi(\mathcal{X}_n) \subseteq \overline{\text{Ball}_{(\mathbb{R}^{2d}, \ell_2)}(0, r)}$ where $\overline{\text{Ball}_{(\mathbb{R}^{2d}, \ell_2)}(0, r)} \stackrel{\text{def}}{=} \{z \in \mathbb{R}^{2d} : \|z\|_{\ell_2} \leq r\}$.

Since $\phi(\mathcal{X}_n) \subseteq \overline{\text{Ball}_{(\mathbb{R}^{2d}, \ell_2)}(0, r)}$ then (Kloeckner, 2010, Proposition 7.4) applies. More specifically, the proof of (Kloeckner, 2010, Proposition 7.4) (see the comment in (Bertrand and Kloeckner, 2012, Example 5.5)) guarantee that there exists some $b \in \mathbb{R}^{2d}$ (that can be built explicitly via Algorithm 1 in Section C) such that the map

$$\iota : (\overline{\text{Ball}_{(\mathbb{R}^{2d}, \ell_2)}(0, r)}, \ell_2) \ni x \mapsto \frac{1}{2d} \sum_{k=1}^{2d} \delta_{\sqrt{2d}(x+b)_k} \in (\mathcal{P}_2(\mathbb{R}), \mathcal{W}_2),$$

is an isometry. Combining the two maps, we find that the composite map

$$\varphi : (\mathcal{X}_n, d_g) \ni x \mapsto \frac{1}{2d} \sum_{k=1}^{2d} \delta_{\sqrt{2d}(\phi(x)+b)_k} \in (\mathcal{P}_2(\mathbb{R}), \mathcal{W}_2),$$

is an isometry. As in the proof of Lemma 23, since $\iota \circ \phi$ is has range in the set of “degenerate Gaussian mixtures” $\mathcal{Z} \stackrel{\text{def}}{=} \{\mathbb{P} \in \mathcal{G}\mathcal{M}_2(\mathbb{R}^d) : \exists N \in \mathbb{N}_+ \exists w \in \Delta_N \exists \mu_1, \dots, \mu_N \in \mathbb{R}^d \text{ s.t. } \mathbb{P} = \sum_{n=1}^N w_n \delta_{\mu_n}\}$ then, we obtain the conclusion by applying the comparison result in (Delon and Desolneux, 2020, Proposition 6) which states that the “identity map” on $(\mathcal{W}_2(\mathbb{R}), \mathcal{W}_2) \ni \mathbb{P} \mapsto \mathbb{P} \in (\mathcal{G}\mathcal{M}_2(\mathbb{R}), \mathcal{M}\mathcal{W}_2)$ is an isometry on the collection of finitely supported probability measures in \mathcal{Z} . Relabeling ϕ as $\phi(\cdot) + b$ yields the conclusion. ■

Proof [Proof of Proposition 8] The result is directly follows from Lemma 17 applied to the map φ of Lemma 24 and relabeling $C C_K$ as C_K . ■

B.5.2 PROOF OF COROLLARIES TO MAIN RESULTS

Proof [Proof of Corollary 12] We only upper-bound $\text{cap}(\mathcal{X}_n)$ in order to apply Theorem 4. By Lemma 21 $\log(\text{cap}(\mathcal{X}_n, d_n)) \leq 12 \log(c_\mu)$ for any doubling measure μ on (\mathcal{X}_n, d_n) . Since (\mathcal{X}_n, d_n) is the metric space associated to a (finite) combinatorial graph in which every two vertices are connected by at most 2 edges then, (Durand-Cartagena et al., 2021, Theorem 5, Proposition 18, and Proposition 19) we know that $c_\mu = 1 + \rho(A_V)$.

Note also that $\text{aspect}(\mathcal{X}_n, d_n) = \text{diam}(\mathcal{X}_n, d_n)$ since the shortest distance between any two points in a combinatorial (metric) graph is 1. Thus, $\log(\text{cap}(\mathcal{X}_n, d_n)) \leq 12 \log(1 + \rho(A_V))$ and the result now follows from Theorem 4. \blacksquare

Proof [Proof of Corollary 9] As argued in (Baudoïn and Garofalo, 2011, pages 1121-1122), the complete d -dimensional Riemannian manifold (M, g) satisfies the *curvature-dimension inequality* of (2) for some $r \in \mathbb{R}$ if and only if (M, g) 's Ricci curvature tensor ((Jost, 2017, Definition 4.3.3)) is uniformly lower-bounded by r . As demonstrated in (Baudoïn and Garofalo, 2011, Equation (1.1)) the lower Ricci-cvature bound on (M, g) together with the Bishop-Gromov Volume Comparison Theorem (e.g. see (Chavel, 1993, Theorem 3.10)) implies that $c_{d_n} \leq d$. Therefore, (Bruè et al., 2021, Proposition 1.7 (ii)) implies that $\log(\text{cap}(\mathcal{X}_n, d_n)) \leq 3 \log(c_{d_n})$. The result therefore follows from Theorem 4. \blacksquare

B.5.3 PROOF OF IMPOSSIBILITY RESULTS

Proof [Proof of Proposition 13]

Step 1 - A lower Bound when Embedding into Spaces of Markov type $p \geq 2$

On (Naor et al., 2006, page 173) the authors discuss that the proof of a result in (Linial et al., 2002) implies the following for any metric space \mathcal{R} of Markov type 2 (see (Naor et al., 2006, Definition 1.1)). If $\Gamma_{g,\delta} \stackrel{\text{def}}{=} (X_{g,\delta}, d_{g,\delta})$ is a graph (equipped with the usual graph geodesic distance) in which the shortest loop has length at least $g > 0$ (called $\Gamma_{g,\delta}$'s *girth*) and its average degree is $\delta > 2$ then, any bi-Lipschitz embedding $f : \Gamma_{g,\delta} \rightarrow (\mathcal{R}, d_{\mathcal{R}})$ (here $s = 1$) has distortion $\text{Dist}(f)$ at-least

$$\text{Dist}(f) \geq \frac{\delta - 2}{2 M_p(\mathcal{R})} g^{(p-1)/p}, \quad (52)$$

where p is the Markov-type of $(\mathcal{R}, d_{\mathcal{R}})$ as defined by in (Ball, 1992, Definition 1.1) and $M_p(\mathcal{R}) > 0$ is a constant depending only on \mathcal{R} 's geometry.

Step 2 - Markov Type of (M, g)

Let $(\mathcal{R}, d_{\mathcal{R}}) \stackrel{\text{def}}{=} (M, d_g)$ where (M, g) is a complete Riemannian manifold with sectional curvature pinched between $[-C, -c]$ for some $0 < c \leq C < \infty$. Therefore, (Naor et al., 2006, Corollary 6.5) applies implying that (M, d_g) has Markov type $p = 2$.

Step 3 - Realizing the Diverging Lower-Bound via Sextet Graphs

The main result of (Weiss, 1984) shows that, for every prime integer $n > 2$ satisfying $n \bmod 16 \in \{\pm 3, \pm 5, \pm 7\}$, the sextet graph with n vertices is a *cubic graph* (i.e. each vertex has exactly 3 edges connected to it and therefore $\delta = 3 > 2$) and its girth g is at-least $\frac{4}{3} \log_2(n) - 2$. Therefore, setting $\Gamma_{g,\delta}$ to be the sextet graph with n -vertices we deduce from (52) and step 2 that any bi-Lipschitz embedding $f : \Gamma_{g,\delta} \rightarrow \mathcal{R}$ satisfies

$$\text{Dist}(f) \geq \frac{1}{2 M_p(\mathcal{R})} \sqrt{\frac{4}{3} \log_2(n) - 2}.$$

Setting $c_{\mathcal{R}} \stackrel{\text{def}}{=} \frac{1}{2 M_p(\mathcal{R})}$ yields the conclusion. \blacksquare

Proof [Proof of Proposition 15] We argue by contradiction. Assume that \mathcal{R} is a Riemannian manifold which is bi-Lipschitz embedded in the Hilbert space ℓ^2 with distortion D_1 ; denote this bi-Lipschitz embedding by $\psi : (\mathcal{R}, g) \hookrightarrow (\mathbb{R}^{2d}, \ell_2)$ and with the property that: for every $n \in \mathbb{N}_+$ there is no n -point metric space (\mathcal{X}_n, d_n) which can be bi-Lipschitz embedded in (\mathcal{R}, d_g) with distortion $D_2 \left(\frac{\log(n)}{\log(\log(n))} \right)$ (here D_2 is a constant depending only on \mathcal{R}).

Therefore, for every n -point metric space (\mathcal{X}_n, d_n) there exists a bi-Lipschitz embedding $\phi : (\mathcal{X}_n, d_n) \hookrightarrow (\mathcal{R}, d_g)$ whose distortion D is less than $D_2 \frac{\log(n)}{\log(\log(n))}$. Thus $\psi \circ \phi : \mathcal{X}_n \rightarrow \ell^2$, is a bi-Lipschitz embedding with

at-least distortion $D_1 D_2$. Thus, for n large enough there is an n -point metric space which bi-Lipschitz embeds into (ℓ^2, ℓ_2) with distortion less than $\mathcal{O}\left(\frac{\log(n)}{\log(\log(n))}\right)$. This is a contradiction of a main result of (Bourgain, 1985). Whence, (\mathcal{R}, g) does not exist. ■

C. Algorithmic Construction of the “Bias Term” in our Embeddings’ Proofs

This short appendix contains an explicit algorithmic construction of the shift (or “bias term”), denoted by b , used when embedding a bounded Euclidean ball into the Wasserstein space. The term $b \in \mathbb{R}^d$ maps the Euclidean ball into the space $\mathbb{R}_{<}^d \stackrel{\text{def.}}{=} \{x \in \mathbb{R}^d : x_1 < \dots < x_d\}$ used in (Bertrand and Kloeckner, 2012, Example 5.5) and initially proposed in (Kloeckner, 2010, Proposition 7.4).

Algorithm 1 Initialize Bias

Require: Set of n -vectors $\mathbb{X} \stackrel{\text{def.}}{=} \{x^{(1)}, \dots, x^{(n)}\}$ in \mathbb{R}^K ,

$b_0 \stackrel{\text{def.}}{=} 0$ ▷ Initialize first shift

for $n = 1, \dots, N$ **do** $\tilde{x}_1^n \stackrel{\text{def.}}{=} x_1^{(n)} + b_1$ ▷ Dummy vectors

end for

for $k = 1, \dots, K$ **do** ▷ Iteratively build bias components

$b_k \stackrel{\text{def.}}{=} \max_{n \leq N} \text{ReLU}(x_{k-1}^{(n)} - x_{k-1}^{(n)})$

for $n \leq N$ **do** $\tilde{x}_k^n \stackrel{\text{def.}}{=} x_k^{(n)} + b_k$ ▷ Dummy vectors

end for

end for

return $b \stackrel{\text{def.}}{=} (b_1, \dots, b_K)$ ▷ Return Bias)

References

- B. Acciaio, J. Backhoff-Veraguas, and A. Zalashko. Causal optimal transport and its links to enlargement of filtrations and continuous-time stochastic optimization. *Stochastic Process. Appl.*, 130(5):2918–2953, 2020. URL <https://doi.org/10.1016/j.spa.2019.08.009>.
- Beatrice Acciaio, Anastasis Kratsios, and Gudmund Pammer. Designing universal causal deep learning models: The geometric (hyper)transformer. *Mathematical Finance (Forthcoming)*, page TBA, 2023. Special Issue: Machine Learning in Finance.
- Shivani Agarwal and Partha Niyogi. Generalization bounds for ranking algorithms via algorithmic stability. *J. Mach. Learn. Res.*, 10:441–474, 2009. ISSN 1532-4435.
- Shun-ichi Amari. *Information geometry and its applications*, volume 194 of *Applied Mathematical Sciences*. Springer, 2016.
- Alexandr Andoni, Assaf Naor, and Ofer Neiman. Snowflake universality of Wasserstein spaces. *Ann. Sci. Éc. Norm. Supér. (4)*, 51(3):657–700, 2018.
- Patrice Assouad. Étude d’une dimension métrique liée à la possibilité de plongements dans \mathbb{R}^n . *C. R. Acad. Sci. Paris Sér. A-B*, 288(15):A731–A734, 1979.
- Julio Backhoff, Daniel Bartl, Mathias Beiglböck, and Johannes Wiesel. Estimating processes in adapted Wasserstein distance. *Ann. Appl. Probab.*, 32(1):529–550, 2022.

- J. Backhoff-Veraguas and G. Pammer. Stability of martingale optimal transport and weak optimal transport. *Ann. Appl. Probab.*, 32(1):721–752, 2022a. URL <https://doi.org/10.1214/21-aap1694>.
- J. Backhoff-Veraguas and G. Pammer. Stability of martingale optimal transport and weak optimal transport. *Ann. Appl. Probab.*, 32(1):721–752, 2022b.
- Julio Backhoff-Veraguas, Daniel Bartl, Mathias Beiglböck, and Manu Eder. All adapted topologies are equal. *Probab. Theory Related Fields*, 178(3-4):1125–1172, 2020a. ISSN 0178-8051. URL <https://doi.org/10.1007/s00440-020-00993-8>.
- Julio Backhoff-Veraguas, Daniel Bartl, Mathias Beiglböck, and Manu Eder. All adapted topologies are equal. *Probab. Theory Related Fields*, 178(3-4):1125–1172, 2020b.
- Keith Ball. Markov chains, riesz transforms and lipschitz maps. *Geometric & Functional Analysis GAFA*, 2(2):137–172, 1992.
- Yair Bartal, Nathan Linial, Manor Mendel, and Assaf Naor. On metric Ramsey-type phenomena. *Ann. of Math. (2)*, 162(2):643–709, 2005. URL <https://doi.org/10.4007/annals.2005.162.643>.
- Peter L. Bartlett and Philip M. Long. Failures of model-dependent generalization bounds for least-norm interpolation. *J. Mach. Learn. Res.*, 22:Paper No. 204, 15, 2021. ISSN 1532-4435.
- Peter L Bartlett, Dylan J Foster, and Matus J Telgarsky. Spectrally-normalized margin bounds for neural networks. *Advances in neural information processing systems*, 30, 2017.
- Peter L. Bartlett, Nick Harvey, Christopher Liaw, and Abbas Mehrabian. Nearly-tight vc-dimension and pseudodimension bounds for piecewise linear neural networks. *Journal of Machine Learning Research*, 20(63):1–17, 2019. URL <http://jmlr.org/papers/v20/17-612.html>.
- Fabrice Baudoin and Nicola Garofalo. Perelman’s entropy and doubling property on Riemannian manifolds. *J. Geom. Anal.*, 21(4):1119–1131, 2011. URL <https://doi.org/10.1007/s12220-010-9180-x>.
- Fabrice Baudoin, Michel Bonnefont, and Nicola Garofalo. A sub-Riemannian curvature-dimension inequality, volume doubling property and the Poincaré inequality. *Math. Ann.*, 358(3-4):833–860, 2014. URL <https://doi.org/10.1007/s00208-013-0961-y>.
- Mathias Beiglböck, Pierre Henry-Labordère, and Nizar Touzi. Monotone martingale transport plans and Skorokhod embedding. *Stochastic Process. Appl.*, 127(9):3005–3013, 2017a. URL <https://doi.org/10.1016/j.spa.2017.01.004>.
- Mathias Beiglböck, Marcel Nutz, and Nizar Touzi. Complete duality for martingale optimal transport on the line. *Ann. Probab.*, 45(5):3038–3074, 2017b.
- Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.
- Mikhail Belkin and Partha Niyogi. Convergence of laplacian eigenmaps. *Advances in neural information processing systems*, 19, 2006.
- Jérôme Bertrand and Benoît Kloeckner. A geometric study of Wasserstein spaces: Hadamard spaces. *J. Topol. Anal.*, 4(4):515–542, 2012. URL <https://doi.org/10.1142/S1793525312500227>.
- Norman L Biggs and MJ Hoare. The sextet construction for cubic graphs. *Combinatorica*, 3(2):153–165, 1983.
- Jérémie Bigot, Raúl Gouet, Thierry Klein, and Alfredo López. Geodesic PCA in the Wasserstein space by convex PCA. *Ann. Inst. Henri Poincaré Probab. Stat.*, 53(1):1–26, 2017. URL <https://doi.org/10.1214/15-AIHP706>.

- Eugene Bilokopytov and Anastasis Kratsios. Non-Euclidean Universal Approximation. *NeurIPS*, 33, 2020.
- Jocelyne Bion-Nadal and Denis Talay. On a Wasserstein-type distance between solutions to stochastic differential equations. *Ann. Appl. Probab.*, 29(3):1609–1639, 2019. URL <https://doi.org/10.1214/18-AAP1423>.
- Christopher M Bishop. Mixture density networks. *Unpublished*, 1994.
- J. Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel J. Math.*, 52(1-2):46–52, 1985.
- Elia Bruè, Simone Di Marino, and Federico Stra. Linear Lipschitz and C^1 extension operators through random projection. *J. Funct. Anal.*, 280(4):Paper No. 108868, 21, 2021. URL <https://doi.org/10.1016/j.jfa.2020.108868>.
- Sebastien Bubeck, Ronen Eldan, Yin Tat Lee, and Dan Mikulincer. Network size and size of the weights in memorization with two-layers neural networks. *Advances in Neural Information Processing Systems*, 33: 4977–4986, 2020.
- Fred Buckley and Frank Harary. On the Euclidean dimension of a wheel. *Graphs Combin.*, 4(1):23–30, 1988. URL <https://doi.org/10.1007/BF01864150>.
- Martin Burger and Andreas Neubauer. Error bounds for approximation with neural networks. *J. Approx. Theory*, 112(2):235–250, 2001.
- Ines Chami, Albert Gu, Vaggos Chatziafratis, and Christopher Ré. From trees to continuous embeddings and back: Hyperbolic hierarchical clustering. *Advances in Neural Information Processing Systems*, 33: 15065–15076, 2020. URL <https://dl.acm.org/doi/abs/10.5555/3495724.3496987>.
- Ines Chami, Albert Gu, Dat P Nguyen, and Christopher Ré. Horopca: Hyperbolic dimensionality reduction via horospherical projections. In *International Conference on Machine Learning*, pages 1419–1429. PMLR, 2021.
- Isaac Chavel. *Riemannian geometry-a modern introduction*, volume 108 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 1993.
- Yaqing Chen, Zhenhua Lin, and Hans-Georg Müller. Wasserstein regression. *Journal of the American Statistical Association*, pages 1–14, 2021.
- Patrick Cheridito, Arnulf Jentzen, and Florian Rossmannek. Efficient approximation of high-dimensional functions with neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–15, 2021.
- Julien Chevallier. Uniform decomposition of probability measures: quantization, clustering and rate of convergence. *Journal of Applied Probability*, 55(4):1037–1045, 2018.
- Sueli I.R. Costa, Sandra A. Santos, and João E. Strapasson. Fisher information distance: A geometrical reading. *Discrete Applied Mathematics*, 197:59–69, 2015. Distance Geometry and Applications.
- Calin Cruceru, Gary Becigneul, and Octavian-Eugen Ganea. Computationally tractable riemannian manifolds for graph embeddings. *arXiv preprint arXiv:2002.08665*, 2020.
- Calin Cruceru, Gary Becigneul, and Octavian-Eugen Ganea. Computationally tractable riemannian manifolds for graph embeddings. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(8):7133–7141, May 2021. URL <https://ojs.aaai.org/index.php/AAAI/article/view/16877>.

- Juan Antonio Cuesta and Carlos Matran. Notes on the wasserstein metric in hilbert spaces. *The Annals of Probability*, 17(3):1264–1276, 1989. ISSN 00911798. URL <http://www.jstor.org/stable/2244406>.
- Xiongtao Dai and Hans-Georg Müller. Principal component analysis for functional data on Riemannian manifolds and spheres. *Ann. Statist.*, 46(6B):3334–3361, 2018.
- Amit Daniely. Neural networks learning and memorization with (almost) no over-parameterization. *Advances in Neural Information Processing Systems*, 33:9007–9016, 2020.
- Kinkar Ch. Das and Pawan Kumar. Some new bounds on the spectral radius of graphs. *Discrete Math.*, 281(1-3):149–161, 2004. ISSN 0012-365X.
- Valentin Debarnot and Pierre Weiss. Deep-blur : Blind identification and deblurring with convolutional neural networks. June 2022.
- Julie Delon and Agnès Desolneux. A Wasserstein-type distance in the space of Gaussian mixture models. *SIAM Journal on Imaging Sciences*, 13(2):936–970, 2020.
- Julie Delon and Agnès Desolneux. Optimal transport between gaussian mixture models. <https://github.com/judelo/gmmot>, 2021.
- Bhuwan Dhingra, Christopher J. Shallue, Mohammad Norouzi, Andrew M. Dai, and George E. Dahl. Embedding text in hyperbolic spaces. *TextGraphs@NAACL-HLT*, pages 59–69, 2018.
- Yan Dolinsky and H. Mete Soner. Martingale optimal transport and robust hedging in continuous time. *Probab. Theory Related Fields*, 160(1-2):391–427, 2014. ISSN 0178-8051.
- James G. Dowty. Chentsov’s theorem for exponential families. *Inf. Geom.*, 1(1):117–135, 2018.
- Estibalitz Durand-Cartagena, Javier Soria, and Pedro Tradacete. Doubling constants and spectral theory on graphs. *arXiv preprint arXiv:2111.09199*, 2021.
- Aryeh Dvoretzky. Some results on convex bodies and Banach spaces. *Proc. Internat. Sympos. Linear Spaces (Jerusalem, 1960)*, pages 123–160, 1961.
- Weinan E. and Stephan Wojtowytsch. Representation formulas and pointwise properties for Barron functions. *Calc. Var. Partial Differential Equations*, 61(2):Paper No. 46, 37, 2022.
- Marzieh Eidi and Jürgen Jost. Ollivier Ricci curvature of directed hypergraphs. *Scientific Reports*, 10(1):12466, 2020.
- P. Erdős. On an extremal problem in graph theory. *Colloq. Math.*, 13:251–254, 1964/65. ISSN 0010-1354.
- P. Erdős, A. Rényi, and V. T. Sós. On a problem of graph theory. *Studia Sci. Math. Hungar.*, 1:215–235, 1966. ISSN 0081-6906.
- Paul Erdős and Horst Sachs. Regukre graphen gegebener taillenweite mit minimaler knotenzahl. *Wissenschaftliche Zeitschrift der Martin-Luther-Universität Halle*, pages 251–258, 1963.
- Sylvester Eriksson-Bique. Quantitative bi-Lipschitz embeddings of bounded-curvature manifolds and orbifolds. *Geometry & Topology*, 22(4):1961–2026, 2018.
- K. J. Falconer. *The geometry of fractal sets*, volume 85 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 1986.
- Herbert Federer. Colloquium lectures on geometric measure theory. *Bull. Amer. Math. Soc.*, 84(3):291–338, 1978.

- Vitaly Feldman and Chiyuan Zhang. What neural networks memorize and why: Discovering the long tail via influence estimation. *Advances in Neural Information Processing Systems*, 33:2881–2891, 2020.
- Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z. Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, Léo Gautheron, Nathalie T.H. Gayraud, Hicham Janati, Alain Rakotomamonjy, Ievgen Redko, Antoine Rolet, Antony Schutz, Vivien Seguy, Danica J. Sutherland, Romain Tavenard, Alexander Tong, and Titouan Vayer. POT: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021. URL <http://jmlr.org/papers/v22/20-451.html>.
- P. Thomas Fletcher, Suresh Venkatasubramanian, and Sarang Joshi. The geometric median on riemannian manifolds with application to robust atlas estimation. *NeuroImage*, 45(1):S143 – S152, 2009. Mathematics in Brain Imaging.
- M Maurice Fréchet. Sur quelques points du calcul fonctionnel. *Rendiconti del Circolo Matematico di Palermo (1884-1940)*, 22(1):1–72, 1906.
- Octavian Ganea, Gary Becigneul, and Thomas Hofmann. Hyperbolic neural networks. *Advances in Neural Information Processing Systems*, 31:5345–5355, 2018. URL <https://proceedings.neurips.cc/paper/2018/file/dbab2adc8f9d078009ee3fa810bea142-Paper.pdf>.
- Francesco Di Giovanni, Giulia Luise, and Michael M. Bronstein. Heterogeneous manifolds for curvature-aware graph embedding. In *ICLR 2022 Workshop on Geometrical and Topological Representation Learning*, 2022. URL <https://openreview.net/forum?id=rtUxsN-kaxc>.
- Siegfried Graf and Harald Luschgy. *Foundations of quantization for probability distributions*, volume 1730 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 2000.
- Gaoyue Guo and Jan Obłóć. Computational methods for martingale optimal transport problems. *Ann. Appl. Probab.*, 29(6):3311–3347, 2019.
- A. Gupta. Embedding tree metrics into low-dimensional Euclidean spaces. *Discrete Comput. Geom.*, 24(1):105–116, 2000.
- Juha Heinonen. *Lectures on analysis on metric spaces*. Universitext. Springer-Verlag, New York, 2001.
- Juha Heinonen. *Geometric embeddings of metric spaces*, volume 90 of *Report. University of Jyväskylä Department of Mathematics and Statistics*. University of Jyväskylä, Jyväskylä, 2003.
- Nhat Ho, XuanLong Nguyen, Mikhail Yurochkin, Hung Hai Bui, Viet Huynh, and Dinh Phung. Multilevel clustering via wasserstein means. In *International Conference on Machine Learning*, pages 1501–1509. PMLR, 2017.
- H. Hopf and W. Rinow. Ueber den Begriff der vollständigen differentialgeometrischen Fläche. *Comment. Math. Helv.*, 3(1):209–225, 1931.
- Jürgen Jost. *Riemannian geometry and geometric analysis*. Universitext. Springer, Heidelberg, seventh edition, 2017.
- Heinrich Jung. Ueber die kleinste kugel, die eine räumliche figur einschliesst. *Journal für die reine und angewandte Mathematik*, 123:241–257, 1901. URL <http://eudml.org/doc/149122>.
- Mark Jungerman and Gerhard Ringel. The genus of the n -octahedron: regular cases. *J. Graph Theory*, 2(1):69–75, 1978. ISSN 0364-9024.
- Karin Usadi Katz and Mikhail G. Katz. Bi-Lipschitz approximation by finite-dimensional imbeddings. *Geom. Dedicata*, 150:131–136, 2011.

- Patrick Kidger and Terry Lyons. Universal approximation with deep narrow networks. *Proceedings of Thirty Third Conference on Learning Theory*, 125:2306–2327, 09–12 Jul 2020.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- Benoît Kloeckner. A geometric study of Wasserstein spaces: Euclidean spaces. *Ann. Sc. Norm. Super. Pisa Cl. Sci. (5)*, 9(2):297–323, 2010. URL <https://doi.org/10.2422/2036-2145.2010.2.03>.
- Benoît Kloeckner. Approximation by finitely supported measures. *ESAIM Control Optim. Calc. Var.*, 18(2): 343–359, 2012.
- Benoît R Kloeckner and Jérôme Bertrand. A geometric study of wasserstein spaces: isometric rigidity in negative curvature. *International Mathematics Research Notices*, 2016(5):1368–1386, 2016.
- Max Kochurov, Rasul Karimov, and Serge Kozlukov. Geoopt: Riemannian optimization in pytorch. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, Virtual, 13–18 Jul 2020. PMLR.
- Prodromos Kolyvakis, Alexandros Kalousis, and Dimitris Kiritsis. Hyperbolic knowledge graph embeddings for knowledge base completion. In *European Semantic Web Conference*, pages 199–214. Springer, 2020.
- Anastasis Kratsios. Universal Regular Conditional Distributions. *Constructive Approximation (Forthcoming)*, 2023.
- Anastasis Kratsios and Léonie Papon. Universal approximation theorems for differentiable geometric deep learning. *Journal of Machine Learning Research*, 23:1–73, 2022.
- Anastasis Kratsios, Behnoosh Zamanlooy, Tianlin Liu, and Ivan Dokmanić. Universal approximation under constraints is possible with transformers. In *International Conference on Learning Representations*, 2022.
- R. Krauthgamer, J. R. Lee, M. Mendel, and A. Naor. Measured descent: a new embedding method for finite metrics. *Geom. Funct. Anal.*, 15(4):839–858, 2005.
- Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguná. Hyperbolic geometry of complex networks. *Physical Review E*, 82(3):036106, 2010.
- Matt Le, Stephen Roller, Laetitia Papaxanthos, Douwe Kiela, and Maximilian Nickel. Inferring concept hierarchies from text corpora via hyperbolic embeddings. *ArXiv*, 2019.
- Enrico Le Donne and Tapio Rajala. Assouad dimension, nagata dimension, and uniformly close metric tangents. *Indiana University Mathematics Journal*, pages 21–54, 2015.
- N. Linial, A. Magen, and A. Naor. Girth and Euclidean distortion. *Geom. Funct. Anal.*, 12(2):380–394, 2002.
- Chong Liu and Ariel Neufeld. Compactness criterion for semimartingale laws and semimartingale optimal transport. *Trans. Amer. Math. Soc.*, 372(1):187–231, 2019.
- Yating Liu and Gilles Pagès. Convergence rate of optimal quantization and application to the clustering performance of the empirical measure. *Journal of Machine Learning Research*, 21(86):1–36, 2020. URL <http://jmlr.org/papers/v21/18-804.html>.
- Federico López and Michael Strube. A fully hyperbolic neural model for hierarchical multi-class classification. *ArXiv*, 2020.

- Federico Lopez, Beatrice Pozzetti, Steve Trettel, Michael Strube, and Anna Wienhard. Symmetric spaces for graph embeddings: A finsler-riemannian approach. *Proceedings of the 38th International Conference on Machine Learning*, 139:7090–7101, 18–24 Jul 2021.
- Grigorii A Margulis. Explicit constructions of graphs without short cycles and low density codes. *Combinatorica*, 2(1):71–78, 1982.
- Grigorii Aleksandrovich Margulis. Explicit group-theoretical constructions of combinatorial schemes and their application to the design of expanders and concentrators. *Problemy peredachi informatsii*, 24(1):51–60, 1988.
- Jiri Matoušek. Lecture notes on metric embeddings. Technical report, Technical report, ETH Zürich, 2013.
- Jiří Matoušek. On the distortion required for embedding finite metric spaces into normed spaces. *Israel J. Math.*, 93:333–344, 1996. ISSN 0021-2172.
- Jiří Matoušek. *Lectures on discrete geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002.
- Song Mei and Andrea Montanari. The generalization error of random features regression: Precise asymptotics and the double descent curve. *Comm. Pure Appl. Math.*, 75(4):667–766, 2022.
- Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Generalization error of random feature and kernel methods: Hypercontractivity and kernel matrix concentration. *Appl. Comput. Harmon. Anal.*, 59:3–84, 2022.
- Liang Mi, Wen Zhang, Xianfeng Gu, and Yalin Wang. Variational wasserstein clustering. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 322–337, 2018.
- Nina Miolane, Nicolas Guigui, Alice Le Brigant, Johan Mathe, Benjamin Hou, Yann Thanwerdas, Stefan Heyder, Olivier Peltre, Niklas Koep, Hadi Zaatiti, Hatem Hajri, Yann Cabanes, Thomas Gerald, Paul Chauchat, Christian Shewmake, Daniel Brooks, Bernhard Kainz, Claire Donnat, Susan Holmes, and Xavier Pennec. Geomstats: A python package for Riemannian geometry in machine learning. *Journal of Machine Learning Research*, 21(223):1–9, 2020. URL <http://jmlr.org/papers/v21/19-027.html>.
- James R. Munkres. *Topology*. Prentice Hall, Inc., Upper Saddle River, NJ, 2000. Second edition.
- Tamara Munzner. H3: Laying out large directed graphs in 3d hyperbolic space. In *Proceedings of VIZ'97: Visualization Conference, Information Visualization Symposium and Parallel Rendering Symposium*, pages 2–10. IEEE, 1997.
- Alessandro Muscoloni, Josephine Maria Thomas, Sara Ciucci, Ginestra Bianconi, and Carlo Vittorio Cannistraci. Machine learning meets complex networks via coalescent embedding in the hyperbolic space. *Nature communications*, 8(1):1–19, 2017.
- Assaf Naor. Metric embeddings and lipschitz extensions. *Princeton University, Department of Mathematics*, 2015.
- Assaf Naor. Metric dimension reduction: a snapshot of the Ribe program. *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. I. Plenary lectures*, pages 759–837, 2018.
- Assaf Naor and Ofer Neiman. Assouad’s theorem with dimension independent of the snowflaking. *Revista Matemática Iberoamericana*, 28(4):1123–1142, 2012.
- Assaf Naor and Terence Tao. Scale-oblivious metric fragmentation and the nonlinear Dvoretzky theorem. *Israel J. Math.*, 192(1):489–504, 2012. ISSN 0021-2172.

- Assaf Naor, Yuval Peres, Oded Schramm, and Scott Sheffield. Markov chains in smooth banach spaces and gromov-hyperbolic metric spaces. *Duke Mathematical Journal*, 134(1):165–197, 2006.
- John Nash. C^1 isometric imbeddings. *Ann. of Math. (2)*, 60:383–396, 1954.
- Behnam Neyshabur, Russ R Salakhutdinov, and Nati Srebro. Path-sgd: Path-normalized optimization in deep neural networks. *Advances in neural information processing systems*, 28, 2015.
- Maximillian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. *Advances in neural information processing systems*, 30, 2017.
- James Orlin. A faster strongly polynomial minimum cost flow algorithm. In *Proceedings of the Twentieth annual ACM symposium on Theory of Computing*, pages 377–387, 1988.
- Mikhail I. Ostrovskii. *Metric embeddings*, volume 49 of *De Gruyter Studies in Mathematics*. De Gruyter, Berlin, 2013. Bilipschitz and coarse embeddings into Banach spaces.
- Pierre Pansu. Métriques de Carnot-Carathéodory et quasiisométries des espaces symétriques de rang un. *Ann. of Math. (2)*, 129(1):1–60, 1989. ISSN 0003-486X.
- S. Park, C. Yun, J. Lee, and J. Shin. Minimum width for universal approximation, 2020.
- Sejun Park, Jaeho Lee, Chulhee Yun, and Jinwoo Shin. Provable memorization via deep neural networks using sub-linear parameters. *Proceedings of Thirty Fourth Conference on Learning Theory*, 134:3627–3661, 15–19 Aug 2021.
- Georg Ch. Pflug and Alois Pichler. Dynamic generation of scenario trees. *Comput. Optim. Appl.*, 62(3): 641–668, 2015.
- Michael Puthawala, Konik Kothari, Matti Lassas, Ivan Dokmanić, and Maarten de Hoop. Globally injective ReLU networks. *Journal of Machine Learning Research*, 23:1–55, 2022.
- Erzsébet Ravasz and Albert-László Barabási. Hierarchical organization in complex networks. *Physical review E*, 67(2):026112, 2003.
- P. L. Robinson. The sphere is not flat. *The American Mathematical Monthly*, 113(2):171–173, 2006a. URL <http://www.jstor.org/stable/27641870>.
- PL Robinson. The sphere is not flat. *The American Mathematical Monthly*, 113(2):171–173, 2006b.
- Areejit Samal, R. P. Sreejith, Jiao Gu, Shiping Liu, Emil Saucan, and Jürgen Jost. Comparative analysis of two discretizations of Ricci curvature for complex networks. *Scientific Reports*, 8(1):8650, 2018.
- Rik Sarkar. Low distortion delaunay embedding of trees in hyperbolic plane. In *International Symposium on Graph Drawing*, pages 355–366. Springer, 2011.
- Zuowei Shen, Haizhao Yang, and Shijun Zhang. Deep network with approximation error being reciprocal of width to power of square root of depth. *Neural Comput.*, 33(4):1005–1036, 2021. ISSN 0899-7667.
- Ryohei Shimizu, YUSUKE Mukuta, and Tatsuya Harada. Hyperbolic neural networks++. In *International Conference on Learning Representations*, 2021.
- SA Shkarin. Isometric embedding of finite ultrametric spaces in banach spaces. *Topology and its Applications*, 142(1-3):13–17, 2004.
- Eduardo D. Sontag. Shattering all sets of ‘ k ’ points in “general position” requires $(k|1)/2$ parameters. *Neural Computation*, 9(2):337–348, 1997a.

- Eduardo D Sontag. Shattering all sets of ‘k’ points in “general position” requires $(k-1)/2$ parameters. *Neural Computation*, 9(2):337–348, 1997b.
- Taiji Suzuki. Adaptivity of deep ReLU network for learning in besov and mixed smooth besov spaces: optimal rate and curse of dimensionality. *International Conference on Learning Representations*, 2019.
- Puoya Tabaghi and Ivan Dokmanić. Hyperbolic distance matrices. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1728–1738, 2020.
- Puoya Tabaghi and Ivan Dokmanić. On procrustes analysis in hyperbolic space. *IEEE Signal Processing Letters*, 28:1120–1124, 2021.
- Puoya Tabaghi, Eli Chien, Chao Pan, Jianhao Peng, and Olgica Milenković. Linear classifiers in product space forms. *arXiv preprint arXiv:2102.10204*, 2021.
- Gal Vardi, Gilad Yehudai, and Ohad Shamir. On the optimal memorization power of ReLU neural networks. *International Conference on Learning Representations*, 2022.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. *International Conference on Learning Representations*, 2018.
- Roman Vershynin. Memory capacity of neural networks with threshold and rectified linear unit activations. *SIAM Journal on Mathematics of Data Science*, 2(4):1004–1033, 2020.
- I. A. Vestfrid and A. F. Timan. A universality property of Hilbert spaces. *Dokl. Akad. Nauk SSSR*, 246(3): 528–530, 1979. ISSN 0002-3264.
- A. L. Vol’berg and S. V. Konyagin. On measures with the doubling condition. *Izv. Akad. Nauk SSSR Ser. Mat.*, 51(3):666–675, 1987.
- Leonid Nisonovich Wasserstein. Markov processes over denumerable products of spaces, describing large systems of automata. *Problemy Peredachi Informatsii*, 5(3):64–72, 1969.
- Alfred Weiss. Girths of bipartite sextet graphs. *Combinatorica*, 4(2):241–245, 1984.
- Dmitry Yarotsky. Error bounds for approximations with deep ReLU networks. *Neural Networks*, 94:103 – 114, 2017. ISSN 0893-6080.
- Jianbo Ye, Panruo Wu, James Z Wang, and Jia Li. Fast discrete distribution clustering using wasserstein barycenter with sparse support. *IEEE Transactions on Signal Processing*, 65(9):2317–2332, 2017.
- Shuxin Zheng, Qi Meng, Huishuai Zhang, Wei Chen, Nenghai Yu, and Tie-Yan Liu. Capacity control of ReLU neural networks by basis-path norm. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 5925–5932, Jul. 2019.
- Ding-Xuan Zhou. Universality of deep convolutional neural networks. *Applied and computational harmonic analysis*, 48(2):787–794, 2020.
- George Kingsley Zipf. Human behavior and the principle of least effort. pages 1–334, 1949.