Geometric Deep Learning with Quasiconformal Neural Networks: An Introduction

Nico Alvarado
nfalvarado@mat.uc.cl

Hans Lobel
halobel@ing.puc.cl

Abstract

We introduce Quasiconformal Neural Networks (QNNs), a novel framework that integrates quasiconformal maps into neural architectures, providing a rigorous mathematical basis for handling non-Euclidean data. QNNs control geometric distortions using bounded maximal dilatation across network layers, preserving essential data structures. We present theoretical results that guarantee the stability and geometric consistency of QNNs. This work opens new avenues in geometric deep learning, particularly for applications involving complex topologies, with significant implications for fields such as image registration and medical imaging.

1 Introduction

The recent advancements in deep learning have been largely fueled by the development of new architectures and optimization techniques. However, many of these models operate in Euclidean spaces, which limits their ability to capture and process data with intrinsic geometric structures. In an attempt to bridge this gap, *quasiconformal maps* have emerged as a powerful tool. Originating from complex analysis and Teichmüller theory, quasiconformal maps allow the transformation of domains in a manner that distorts angles but controls the distortion of shapes [18]. This property makes them particularly suitable for modeling non-Euclidean geometries, which arise in various real-world data types, such as image processing, medical imaging, and computer vision.

Quasiconformal neural networks (QNNs) are a novel class of models that leverage quasiconformal maps to introduce geometric flexibility into deep learning architectures. By integrating these maps, QNNs provide a principled way to handle datasets with complex topologies, enabling transformations that are both flexible and mathematically well-behaved. This new perspective not only expands the applicability of neural networks to previously challenging domains but also enhances their robustness when applied to data that is highly structured, such as surfaces and manifolds.

The connection between quasiconformal maps and Teichmüller theory is particularly noteworthy in this context. Teichmüller theory, which studies the moduli spaces of Riemann surfaces and quasiconformal deformations, provides a rich mathematical framework for understanding the geometric structure of data. By incorporating ideas from Teichmüller theory into deep learning, QNNs offer a novel approach to learning in non-Euclidean spaces, opening the door to applications in areas where traditional neural networks struggle.

The applications of QNNs span several fields. In computer vision, for example, quasiconformal maps enable shape-preserving transformations that can be critical for tasks such as image registration and object recognition. In medical imaging, these maps allow for more accurate modeling of anatomical structures, improving both diagnostic accuracy and treatment planning. Furthermore, in fields like 3D modeling and graphics, the ability to manipulate complex geometries with controlled distortion offers significant advantages over traditional techniques [17].

In this paper, we explore the theoretical foundations of quasiconformal neural networks and highlight their potential across various domains. We begin by reviewing the mathematical background of

Mathematics of Modern Machine Learning Workshop at NeurIPS 2024.

quasiconformal maps and their connection to deep learning, followed by an analysis of how these maps can be effectively integrated into modern neural network architectures.

1.1 Related work

Geometric Deep Learning Geometric deep learning is a broad term encompassing techniques that extend traditional deep learning models to non-Euclidean domains, such as graphs, manifolds, and other geometric spaces [6, 5]. These methods have been particularly successful in applications where data has an inherent non-Euclidean structure, such as social networks, 3D shape analysis, and molecular modeling [16, 22, 20, 2]. However, many of these methods rely on discrete representations of geometry and lack explicit control over the continuous geometric transformations applied to the data.

Theoretical Developments in Neural Networks with Geometric Constraints There has been significant interest in developing neural network architectures that can process data with built-in geometric constraints. For example, works on equivariant neural networks [12, 11, 13, 3, 23], have explored the idea of making neural networks invariant to certain geometric transformations like rotations and translations. These approaches have shown great promise in improving the generalization and robustness of models in tasks involving highly structured data. However, while these models ensure invariance to predefined geometric transformations, they often lack the flexibility to learn more general, task-specific geometric deformations.

Quasiconformal Neural Networks (QNNs) offer a different perspective by incorporating the flexibility of quasiconformal maps into the network structure. This allows the network to learn deformations that are both flexible and geometrically controlled. Unlike equivariant networks that focus on predefined transformations, QNNs provide the ability to adaptively learn geometrically consistent transformations that are tailored to the data and the task at hand. This novel framework builds upon existing work in geometric deep learning, but expands it with the rigorous mathematical properties of quasiconformal mappings.

Hyperbolic Neural Networks Hyperbolic neural networks (HNNs) and quasiconformal neural networks share common ground in their use of non-Euclidean geometries to model complex data structures [19, 9, 15, 21]. HNNs operate within hyperbolic space, which is particularly well-suited for representing hierarchical or tree-like data, where the natural curvature of the space enables efficient embeddings of such structures. This is similar to the goal of QNNs, which leverage quasiconformal mappings to control data transformations while preserving important structural properties, such as local angles and shapes.

In both cases, the networks aim to provide a more flexible and accurate representation of complex, structured data that traditional Euclidean-based networks struggle to capture. HNNs exploit the exponential scaling of distances in hyperbolic space, making them particularly effective for capturing hierarchical relationships [1, 7, 8]. Similarly, QNNs, through quasiconformal maps, allow for controlled warping of data that can preserve key relationships under transformation, enabling better generalization on irregular data.

The adaptability of both approaches to non-Euclidean domains, such as graphs and manifolds, highlights their potential for tasks involving data with inherent structure. While HNNs are specifically designed to capture global properties such as hierarchy [4, 14], QNNs focus on controlling local distortions in the feature space. Both methods, however, underscore the importance of geometry in enhancing the expressiveness of neural networks for complex data, making them relevant to each other in the broader context of advancing geometric deep learning methods.

1.2 Paper contributions

1. We show that QNNs are stable under small geometric changes in the input data (Theorem 3.2).
2. We show that QNNs allow for flexible geometric transformations while maintaining control over how much the data is distorted (Proposition 3.3).

3. We also proved that QNNs preserve important topological properties of the data, such as Betti numbers and homology groups, even when the data undergoes smooth geometric deformations (Proposition 3.4).

2 Preliminaries

Now we need to review some basics related to quasiconformal maps.

Quasiconformal maps.

Definition 2.1. Let $f : \mathbb{R}^n \to \mathbb{R}^n$ be a map. The map f is called *quasiconformal* if it satisfies the following conditions: (1) f is a homeomorphism, (2) f is differentiable almost everywhere in $\mathbb{R}^n$, (3) The differential of f, $Df(x)$, satisfies

$$\frac{\|Df(x)\|}{\inf_{|v|=1} |Df(x)v|} \leq K \quad \text{for some constant } K \geq 1,$$

where $\|Df(x)\|$ is the operator norm of the differential, representing the maximal stretching factor, and $\inf_{|v|=1} |Df(x)v|$ is the minimal stretching factor at x, and (4) The Jacobian determinant $J_f(x)$ is nonzero almost everywhere.

Definition 2.2. Let $\Omega \subset \mathbb{C}$ be a domain. A homeomorphism $f : \Omega \to f(\Omega) \subset \mathbb{C}$ is said to be *quasiconformal* if it satisfies the following properties: (1) f is differentiable almost everywhere (for a Lebesgue measure), (2) the partial derivatives of f satisfy the Beltrami equation:

$$\frac{\partial f}{\partial \bar{z}} = \mu(z)\frac{\partial f}{\partial z},$$

where $\mu : \Omega \to \mathbb{C}$ is a measurable function called the *Beltrami coefficient* and satisfies $\|\mu\|_\infty < 1$ almost everywhere in Ω, and (3) the *dilatation* of f is bounded, that is, there exists a constant $K \geq 1$ such that:

$$\frac{|f_z| + |f_{\bar{z}}|}{|f_z| - |f_{\bar{z}}|} \leq K \quad \text{almost everywhere},$$

where $f_z = \frac{\partial f}{\partial z}$ and $f_{\bar{z}} = \frac{\partial f}{\partial \bar{z}}$ are the complex derivatives of f. The constant K is called the *maximal dilatation* of f, and a map with $K = 1$ is conformal.

By controlling the magnitude of the Beltrami coefficients within the network, the diffeomorphic property of mappings can be maintained, ensuring that the mappings remain invertible and free of topological inconsistencies ([10]).

Let $T \colon \mathbb{C} \to \mathbb{C}$ be a linear transformation defined as $T(u) = au + b\overline{u}$. Then we can get a quasiconformal map from an ellipse to the unit circle (See Figure 1).

In higher dimensions, the analog of the Beltrami equation controls how the map's differential deviates from being conformal, similarly to how Definition 2.1 bounds the distortion using the operator norm $\|Df(x)\|$. So, in the rest of the paper we will refer to the Beltrami coefficients instead of $\|Df(x)\|$.

Solving the Beltrami equation provides a quasiconformal map for a given Beltrami coefficient. This equation forms the foundation for the development of quasiconformal deformations in Teichmüller theory and plays a critical role in the construction of the maps used in QNNs.

Quasiconformal maps generalize conformal maps by allowing controlled distortion. The function μ measures how much f deviates from being conformal. If $\mu = 0$ everywhere, the map is conformal, and it preserves angles locally. If $\mu \neq 0$, the map distorts angles, but the amount of distortion is limited by the value of K.

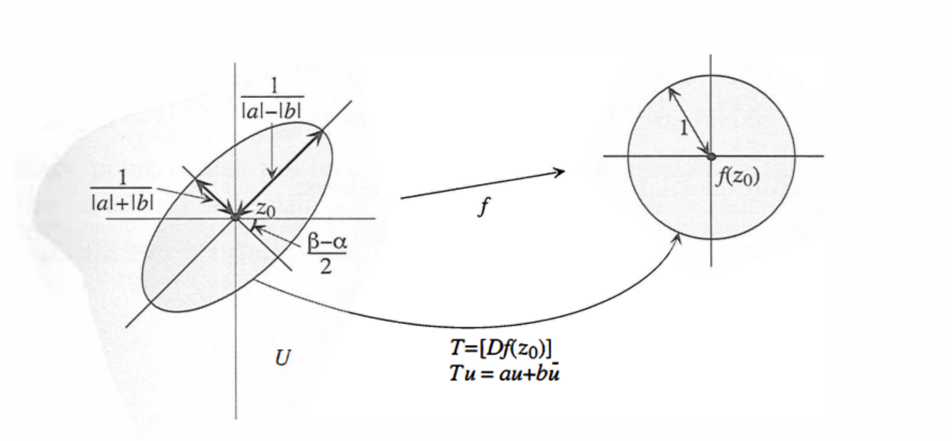

Figure 1: If f is quasiconformal and of class $\mathcal{C}^1$ then its derivative at z_0 takes the the ellipse on the left to the unit circle on the right (for more details see [18]).

Example 2.3. *A simple example of a quasiconformal map is the* affine stretch map*, which stretches the complex plane by different factors along the real and imaginary axes. Let* $f : \mathbb{C} \to \mathbb{C}$ *be defined as*

$$f(z) = \alpha x + i\beta y, \quad \textit{where} \quad z = x + iy, \quad \alpha > 0, \quad \beta > 0.$$

In this map, the real part is scaled by α *and the imaginary part is scaled by* β*. The Beltrami coefficient* μ *is given by*

$$\mu = \frac{\alpha - \beta}{\alpha + \beta}.$$

The map is quasiconformal as long as $\mu \in (-1, 1)$*, meaning that* α *and* β *must be positive and cannot differ too much. The dilatation* K *of the map is given by*

$$K = \frac{1 + |\mu|}{1 - |\mu|}.$$

Example 2.4. *Another example is the* logarithmic spiral map*. Consider the map* $f : \mathbb{C}\backslash\{0\} \to \mathbb{C}\backslash\{0\}$ *given by*

$$f(z) = z^{\lambda}, \quad \lambda \in \mathbb{R}, \lambda > 0.$$

In polar coordinates $z = re^{i\theta}$*, the map becomes*

$$f(re^{i\theta}) = r^{\lambda} e^{i\lambda\theta}.$$

This map preserves angles but stretches distances by a factor of $r^{\lambda-1}$*, making it quasiconformal for* $\lambda \neq 1$*. The Beltrami coefficient for this map is:*

$$\mu = \frac{\lambda - 1}{\lambda + 1}.$$

For $\lambda \in (0, \infty)$*, the map is quasiconformal as long as* $\mu \in (-1, 1)$*. The dilatation* K *is similarly given by*

$$K = \frac{1 + |\mu|}{1 - |\mu|}.$$

Distortion and Maximal Dilatation. The key measure of how much a quasiconformal map distorts the local geometry is its *maximal dilatation*, K. This quantity describes the ratio of the maximal and minimal stretching that occurs under the mapping.

Definition 2.5. For a quasiconformal map $f : \Omega \to \mathbb{C}$, the dilatation at a point is given by:

$$K(z) = \frac{|f_z| + |f_{\bar{z}}|}{|f_z| - |f_{\bar{z}}|}, \quad \text{for } f_z \neq 0.$$

The maximal dilatation K is the essential supremum of $K(z)$ over all $z \in \Omega$. A smaller K implies less distortion, with $K = 1$ corresponding to conformal (angle-preserving) maps.

Teichmüller Theory and Quasiconformal Maps. Teichmüller theory is the study of the deformation of complex structures on Riemann surfaces, particularly through quasiconformal maps. It is deeply intertwined with the study of moduli spaces, which parameterize distinct geometric structures. Central to Teichmüller theory is the analysis of quasiconformal maps that provide a controlled way to deform Riemann surfaces, offering a rich geometric and analytic framework.

Definition 2.6. Let S be a compact Riemann surface. The *Teichmüller space* $\mathcal{T}(S)$ of S is the space of equivalence classes of marked Riemann surfaces (X, f), where X is a Riemann surface and $f : S \to X$ is a quasiconformal map, with two such maps considered equivalent if they are homotopic by a conformal map.

A central result in Teichmüller theory is the existence of extremal quasiconformal maps. Given two Riemann surfaces X and Y, an extremal quasiconformal map is a map that minimizes the maximal dilatation $K(f)$ among all quasiconformal maps homotopic to a given boundary condition.

This result guarantees that for every pair of points in Teichmüller space, there exists a unique quasiconformal map that realizes the shortest "distance" between them in terms of dilatation. This is fundamental in applications where minimizing distortion is critical.

In the context of neural networks, Teichmüller theory provides a structured way to model and manipulate data with complex geometric structures, such as images or surfaces with non-trivial topology. By using quasiconformal maps to deform data, QNNs can handle a wide range of tasks that require robustness to geometric variations while minimizing distortion. This ability to operate in Teichmüller space offers a powerful framework for learning representations that preserve the essential structure of data while allowing flexibility in its geometric form.

3 Quasiconformal Neural Networks

A QNN is a neural network architecture that incorporates quasiconformal mappings as a means to deform and process data in non-Euclidean or geometrically structured spaces. Formally, let us define a QNN with the following components.

Definition 3.1. A QNN is a function $f : \Omega \to \mathbb{R}^n$ composed of layers, where each layer represents a transformation of the input data. Each transformation is designed to respect the geometric properties of the data, using quasiconformal maps as one of the key transformation mechanisms. More formally, let $x \in \mathbb{R}^n$ be the input data, and let $f_k : \mathbb{R}^n \to \mathbb{R}^n$ represent the transformation at the k-th layer. The network is defined as:

$$f(x) = f_L \circ f_{L-1} \circ \cdots \circ f_1(x),$$

where L is the total number of layers, and each f_k involves a quasiconformal transformation.

At the core of a QNN is the *quasiconformal layer*, which applies a quasiconformal map to the data. The quasiconformal transformation is then applied to the input data, allowing the network to process it in a way that preserves geometric structures while permitting controlled deformations.

In a QNN, the Beltrami coefficient μ is parameterized and learned during the training process. Let $\theta \in \mathbb{R}^p$ be the set of trainable parameters for the network, including the parameters defining the quasiconformal maps. The network learns the optimal Beltrami coefficients μ_θ for each layer that minimize the loss function, subject to the quasiconformal constraint $\|\mu_\theta\|_\infty < 1$. Formally, this involves solving an optimization problem of the form

$$\theta^* = \arg\min_\theta \mathcal{L}(f_\theta(x), y) \quad \text{subject to } \|\mu_\theta\|_\infty < 1,$$

where $\mathcal{L}$ is the loss function, and y represents the target output.

To ensure that the network learns quasiconformal maps, an additional regularization term is often added to the loss function, penalizing deviations from quasiconformality. Specifically, the regularization term encourages the Beltrami coefficient μ_θ to remain bounded in norm, ensuring that the learned maps maintain a controlled distortion. The total loss function can thus be written as

$$\mathcal{L}_{\text{total}} = \mathcal{L}(f_\theta(x), y) + \lambda \|\mu_\theta\|_\infty,$$

where λ is a regularization parameter controlling the strength of the quasiconformal constraint.

3.1 Theoretical results

Now we state the main results of the paper.

Theorem 3.2. *Let $f_\theta : \mathbb{R}^n \to \mathbb{R}^n$ be a QNN, and let $x \in \mathbb{R}^n$ be the input. Suppose the network's quasiconformal maps are parameterized by a Beltrami coefficient μ_θ satisfying $\|\mu_\theta\|_\infty < 1$. For any small perturbation δx of the input, the output perturbation $\delta f_\theta(x)$ is controlled by the maximal dilatation $K(f_\theta)$, i.e.,*

$$\|\delta f_\theta(x)\| \leq K(f_\theta)\|\delta x\|.$$

Thus, QNNs are stable under small geometric perturbations of the input data.

Proof. See Appendix A □

The previous result shows that the output perturbation $\|\delta f_\theta(x)\|$ is linearly bounded by the input perturbation $\|\delta x\|$, with the proportionality constant being the maximal dilatation $K(f_\theta)$. Therefore, the QNN is stable under small geometric perturbations of the input data, as long as the maximal dilatation remains bounded.

Proposition 3.3. *Let $f_\theta : \mathbb{R}^n \to \mathbb{R}^n$ be a QNN. Suppose the Beltrami coefficient μ_θ for each quasiconformal map satisfies $\|\mu_\theta\|_\infty \leq \mu_{\max} < 1$. Then the maximal dilatation $K(f_\theta)$ of the overall transformation is bounded by*

$$K(f_\theta) \leq \frac{1 + \mu_{\max}}{1 - \mu_{\max}}.$$

Thus, the distortion of the network's transformations is controlled, ensuring that the network preserves the overall structure of the input data while allowing for flexible geometric deformations.

Proof. See Appendix B □

Proposition 3.4. *Let X be a topological space equipped with a Riemannian metric and let f_θ be a Quasiconformal Neural Network acting on data embedded in X. If the data undergoes a smooth topological deformation represented by a quasiconformal map f, then f_θ remains robust in terms of preserving essential topological properties of the data. Specifically, the Betti numbers and homology groups of the deformed data $f(X)$ are preserved under f_θ.*

Proof. See Appendix C □

This results shows that the QNN f_θ preserves the topological invariants of the data, including the Betti numbers and homology groups, after a smooth topological deformation represented by a quasiconformal map. This proves that QNNs are robust to topological changes in the input data.

4 Conclusions and Future work

In this paper, we introduced the concept of QNNs, a novel approach that integrates the mathematical theory of quasiconformal maps into deep learning architectures. By leveraging the unique properties of quasiconformal maps, QNNs offer a flexible yet geometrically constrained framework for processing complex data that resides in non-Euclidean spaces. Our theoretical analysis demonstrated that QNNs can effectively control geometric distortion, as the maximal dilatation is bounded by the parameters of the underlying quasiconformal maps. This ability to manage distortion while preserving important data geometry makes QNNs a promising tool for applications in areas such as image registration, medical imaging, and 3D surface modeling.

It is important to emphasize that this work represents a basic, foundational study, and a great deal of further theoretical and empirical research is required to fully explore and validate the potential of QNNs. Several key areas remain unexplored. From a theoretical perspective, while we have established basic bounds on the dilatation of QNNs, a more in-depth analysis is needed to understand their capacity, convergence properties, and stability in high-dimensional and complex settings. Future work should delve deeper into the connections between QNNs and more advanced mathematical tools from quasiconformal geometry and Teichmüller theory, which may help us better understand the behavior of these networks in practical scenarios.

Acknowledgments and Disclosure of Funding

The authors acknowledge funding support from the National Center of Artificial Intelligence, Chile (CENIA FB210017, Basal ANID) and Fondecyt Regular grant number 1241772 entitled "Embed the city: An artificial intelligence-based approach to learn transferable representations from multi-source and multi-modal urban data."

References

[1] Nico Alvarado and Hans Lobel. Hyperbolic optimizer as a dynamical system. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 1243–1260. PMLR, 21–27 Jul 2024.

[2] Kenneth Atz, Fabio Grisoni, and Gisbert Schneider. Geometric deep learning on molecular representations. *Nature Reviews Chemistry*, 5(9):625–638, 2021.

[3] Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess Smidt, and Boris Kozinsky. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature Communications*, 13(1):2453, 2022.

[4] Michele Borassi, Alessandro Chessa, and Guido Caldarelli. Hyperbolicity measures democracy in real-world networks. *Physical Review E*, 92(3):032812, 2015.

[5] Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Velickovic. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.

[6] Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.

[7] Ines Chami. *Representation Learning and Algorithms in Hyperbolic Spaces*. Stanford University, 2021.

[8] 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.

[9] Ines Chami, Rex Ying, Christopher Ré, and Jure Leskovec. Hyperbolic graph convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

[10] Qiguang Chen, Zhiwen Li, and Lok Ming Lui. A deep learning framework for diffeomorphic mapping problems via quasi-conformal geometry applied to imaging. *SIAM Journal on Imaging Sciences*, 17(1):501–539, 2024.

[11] Taco Cohen, Mario Geiger, Jonas Koehler, and Max Welling. Gauge equivariant convolutional networks and the icosahedral cnn. *arXiv preprint arXiv:1902.04615*, 2019.

[12] Taco S Cohen and Max Welling. Group equivariance in convolutional neural networks. In *International conference on machine learning (ICML)*, pages 2990–2999, 2016.

[13] Carlos Esteves, Claire Allen-Blanchette, Ameesh Makadia, and Kostas Daniilidis. Learning so(3) equivariant representations with spherical cnns. In *European Conference on Computer Vision (ECCV)*, pages 52–68, 2018.

[14] Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. Hyperbolic entailment cones for learning hierarchical embeddings. In *International Conference on Machine Learning*, pages 1646–1655. PMLR, 2018.

[15] Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. Hyperbolic neural networks. *Advances in neural information processing systems*, 31, 2018.

[16] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pages 1263–1272, 2017.

[17] Yuchen Guo, Hanqun Cao, and Lok Ming Lui. Deep learning-based quasi-conformal surface registration for partial 3d faces applied to facial recognition. *Journal of Machine Learning Research*, 20:1–31, 2019.

[18] John H. Hubbard. *Teichmüller Theory and Applications to Geometry, Topology, and Dynamics Volume 1: Teichmüller Theory*. Matrix Editions, 2006.

[19] Qi Liu, Maximilian Nickel, and Douwe Kiela. Hyperbolic graph neural networks. *Advances in neural information processing systems*, 32, 2019.

[20] Hongwei Pei, Bingbing Wei, Kevin Chang, Yuncheng Lei, and Bo Yang. Geom-gcn: Geometric graph convolutional networks. *arXiv preprint arXiv:2002.05287*, 2020.

[21] Wei Peng, Tuomas Varanka, Abdelrahman Mostafa, Henglin Shi, and Guoying Zhao. Hyperbolic deep neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):10023–10044, 2021.

[22] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 652–660, 2017.

[23] Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation- and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*, 2018.

A Proof of Theorem 3.2

Let $f_\theta = f_L \circ f_{L-1} \circ \cdots \circ f_1$ be a QNN, where each layer transformation $f_k : \mathbb{R}^n \to \mathbb{R}^n$ is quasiconformal. Each map f_k satisfies the Beltrami equation in local coordinates:

$$\frac{\partial f_k}{\partial \bar{z}} = \mu_k(z) \frac{\partial f_k}{\partial z}, \quad \|\mu_k\|_\infty < 1.$$

The map f_k is therefore quasiconformal, with maximal dilatation $K(f_k)$ given by:

$$K(f_k) = \frac{1 + \|\mu_k\|_\infty}{1 - \|\mu_k\|_\infty}.$$

Now, consider a perturbation δx in the input data at layer k, i.e., $x \mapsto x + \delta x$. The perturbation in the output at layer k is denoted by $\delta f_k(x)$.

For each quasiconformal map f_k, we know that locally the map behaves as a linear transformation with bounded distortion. More precisely, for small perturbations δx, we have the linear approximation:

$$f_k(x + \delta x) \approx f_k(x) + \mathbf{J}_{f_k}(x) \cdot \delta x,$$

where $\mathbf{J}_{f_k}(x)$ is the Jacobian matrix of f_k at point x. The Jacobian matrix $\mathbf{J}_{f_k}(x)$ satisfies the following bounds due to the quasiconformality of f_k:

$$\lambda_{\min}(x)\|\delta x\| \leq \|\mathbf{J}_{f_k}(x) \cdot \delta x\| \leq \lambda_{\max}(x)\|\delta x\|,$$

where $\lambda_{\min}(x)$ and $\lambda_{\max}(x)$ are the minimum and maximum singular values of $\mathbf{J}_{f_k}(x)$, respectively.

For a quasiconformal map, the ratio of the maximal and minimal singular values at any point is bounded by the maximal dilatation $K(f_k)$, i.e.,

$$\frac{\lambda_{\max}(x)}{\lambda_{\min}(x)} \leq K(f_k).$$

Thus, for small perturbations δx, we have:

$$\|\delta f_k(x)\| = \|f_k(x + \delta x) - f_k(x)\| \leq K(f_k)\|\delta x\|.$$

Since the QNN is composed of multiple layers of quasiconformal maps, each layer contributes to the overall distortion. Let δx_k represent the perturbation at the input of layer k. The output perturbation at the k-th layer is:

$$\|\delta f_k(x_k)\| \le K(f_k)\|\delta x_k\|.$$

Now, passing this perturbation to the next layer, the perturbation at the next layer's input is $\delta x_{k+1} = \delta f_k(x_k)$, so

$$\|\delta x_{k+1}\| = \|\delta f_k(x_k)\| \le K(f_k)\|\delta x_k\|.$$

By repeating this for each layer, we obtain the total perturbation at the final output of the network as

$$\|\delta f_\theta(x)\| \le \prod_{k=1}^{L} K(f_k)\|\delta x_1\|.$$

Since the total maximal dilatation of the QNN is bounded by the product of the dilatations of individual layers, we define

$$K(f_\theta) = \prod_{k=1}^{L} K(f_k).$$

Thus, we obtain the final bound on the perturbation of the output

$$\|\delta f_\theta(x)\| \le K(f_\theta)\|\delta x\|.$$

B Proof of Proposition 3.3

Let $f_i : \mathbb{R}^n \to \mathbb{R}^n$ be a quasiconformal map in the i-th layer of a QNN. Since $\|\mu_i(z)\|_\infty \le \mu_{\max}$, we have

$$K(f_i) \le \frac{1+\mu_{\max}}{1-\mu_{\max}}.$$

Thus, the maximal dilatation of each quasiconformal map in the QNN is bounded by $K_{\max}$.

Now consider the QNN f_θ, which is composed of multiple quasiconformal layers. Let $f_1, f_2, \ldots, f_L$ represent the quasiconformal maps in the L layers of the network. The maximal dilatation of the composition is bounded by the product the L distorted layers in the composition but since each map corrects and smooths out distortions from previous layers, we are preventing an exponential accumulation of distortion. Thus, we have

$$K(f_\theta) \le \frac{1+\mu_{\max}}{1-\mu_{\max}}.$$

C Proof of Proposition 3.4

By definition, a quasiconformal map $f : X \to X'$ is a homeomorphism, meaning that it is both continuous and has a continuous inverse. Homeomorphisms preserve the Betti numbers $b_i(X)$, which are the ranks of the homology groups $H_i(X, \mathbb{Z})$. The Betti numbers represent the number of i-dimensional holes in the space. Since f is a homeomorphism, it induces an isomorphism on the homology groups:

$$f_* : H_i(X, \mathbb{Z}) \to H_i(X', \mathbb{Z}),$$

for each $i \ge 0$. Therefore, the Betti numbers of X are preserved under the quasiconformal map f, i.e.,

$$b_i(X) = b_i(X'), \quad \text{for all } i \ge 0.$$

Next, consider the action of the QNN f_θ on the data. Each layer of f_θ involves a transformation $f_k : \mathbb{R}^n \to \mathbb{R}^n$, which is either a quasiconformal map or a standard neural network layer (such as an affine transformation followed by a non-linear activation function).

If the data embedded in X undergoes a topological deformation via $f : X \to X'$, the composition of this map with the transformations of the QNN preserves the topological properties as well. Let f_θ be represented as the composition of layer transformations f_k:

$$f_\theta = f_L \circ f_{L-1} \circ \cdots \circ f_1.$$

Since each f_k is either quasiconformal or a topologically trivial map (such as an affine transformation), the composition remains a quasiconformal transformation, and thus a homeomorphism. This ensures that the QNN does not alter the topological invariants of the data.

Given that f_θ is composed of quasiconformal maps, it induces an isomorphism on the homology groups of the space. Specifically, for each $i \geq 0$, the map f_θ induces a homomorphism on the homology groups $H_i(X, \mathbb{Z})$:

$$(f_\theta)_* : H_i(X, \mathbb{Z}) \to H_i(X', \mathbb{Z}),$$

which is an isomorphism due to the homeomorphic nature of the quasiconformal maps. As a result, the Betti numbers, which are the ranks of the homology groups, are preserved under the action of the QNN. Therefore, the topological structure of the data, as measured by the Betti numbers $b_i(X)$, is invariant under the action of the QNN.