
Vector Quantile Regression on Manifolds

Marco Pegoraro

Sapienza, University of Rome
Technion

Sanketh Vedula

Technion
Sibylla

Aviv A. Rosenberg

Technion
Sibylla

Irene Tallini

Sapienza, University of Rome
Technion

Emanuele Rodolà

Sapienza, University of Rome

Alex M. Bronstein

Technion
Sibylla

Abstract

Quantile regression (QR) is a statistical tool for distribution-free estimation of conditional quantiles of a target variable given explanatory features. QR is limited by the assumption that the target distribution is univariate and defined on an Euclidean domain. Although the notion of quantiles was recently extended to multi-variate distributions, QR for multi-variate distributions on manifolds remains underexplored, even though many important applications inherently involve data distributed on, e.g., spheres (climate and geological phenomena), and tori (dihedral angles in proteins). By leveraging optimal transport theory and c -concave functions, we meaningfully define conditional vector quantile functions of high-dimensional variables on manifolds (M-CVQFs). Our approach allows for quantile estimation, regression, and computation of conditional confidence sets and likelihoods. We demonstrate the approach’s efficacy and provide insights regarding the meaning of non-Euclidean quantiles through synthetic and real data experiments.

1 INTRODUCTION

Quantile regression (QR) (Koenker and Bassett, 1978) is a powerful statistical tool that estimates *conditional quantiles* of a target variable Y , given covariates \mathbf{X} . QR

is usually formulated as a regression problem minimizing the *pinball loss*, the definition of which inherently assumes both scalar and Euclidean data. QR has been extensively employed in applications with scalar target variables in the Euclidean domain, where the notion of a quantile is both well-defined and widely understood. However, there exist many real-world applications in which the target data are distributed on a *manifold*, such as a sphere or a cylinder in case of climate measurements (Lagona, 2018; Scher and Messori, 2020), a torus in case of protein dihedral angles (Rosenberg et al., 2022), or more general manifolds for medical images (Pennec et al., 2019). These datasets are better represented as points lying on manifolds, which are geometric spaces with nontrivial topological and geometric properties. A meaningful extension of QR to manifolds would, therefore, unlock the power of this statistical tool for a broader range of applications.

Recently, new perspectives on quantile functions have emerged, allowing for their extension into high-dimensional data. Notably, Carlier et al. (2016) proposed *vector quantile regression* (VQR), defining the *vector quantile function* (VQF) as a co-monotonic map—which can be obtained by solving a suitable optimal transport (OT) problem—between a multivariate uniform distribution and the target distribution of interest. They further showed that quantile regression can be framed as a *conditional* OT problem where a family of OT maps, parametrized by covariates \mathbf{X} , are solved for simultaneously. The resulting maps are the *conditional vector quantile functions* (CVQFs) of the multivariate target variable $\mathbf{Y}|\mathbf{X}$. This approach is, however, limited by the assumption of linearity and the use of the primal OT formulation, which is hard to scale. More recently, Rosenberg et al. (2023) extended VQR by proposing a non-linear extension and introducing a more scalable solver which exploits the entropic-regularized dual of the conditional OT problem. Although this work intro-

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).




Figure 1: **Sampling and confidence sets for the ‘Scaled Heart’ and ‘Scaled Star’ distributions $\mathbf{Y}|\mathbf{X}$ on \mathcal{S}^2 and \mathcal{T}_2 , under different conditioning values.** The conditioning variable x controls the scale of the distribution. τ -contours shown as colored lines. The probability of $\mathbf{Y}|\mathbf{X}$ falling inside a τ -contour is τ .

duced a learned embedding, to incorporate inductive bias through the structure of \mathbf{X} , the approach is unable to exploit any intrinsic structure present in \mathbf{Y} .

Leveraging the OT-based formulation of quantile functions, Hallin et al. (2022) proposed quantile estimation for spherical data, by solving OT between the base and target distributions on the n -sphere. However, their approach has the following limitations: (i) it is limited to quantile *estimation*, as opposed to regression (only supports *unconditional* quantile functions); (ii) it is defined only for the sphere, and (iii) it requires solving the primal OT formulation, resulting in a large-scale linear program which becomes impractical even for moderately-sized problems. Despite these limitations, Hallin et al. provide a motivating example for the potential of OT methods for estimating distributions on manifolds, and is, to our knowledge, the only approach proposed so far for defining non-Euclidean quantiles.

Contributions. We propose a novel, *scalable* approach for estimating *multi-dimensional conditional* quantile functions on *manifolds*. We build upon a dual formulation of the Riemannian conditional OT problem to solve for a family of OT maps parametrized in \mathbf{X} . We model the conditional OT maps as gradients of c -convex potentials, which are represented by partially input c -convex neural networks. Thus, we address the major limitations of both previous works (Hallin et al., 2022; Rosenberg et al., 2023), in a holistic framework supporting multivariate distributions on any manifold for which the exponential map is known. To the best of our knowledge, this is the first work to estimate *conditional* VQFs on manifolds, and, more generally, to address the problem of conditional OT on manifolds. Furthermore, our method extends beyond the capabilities of existing approaches by not only computing quantiles but also addressing additional aspects, including sampling and likelihood computation. This ability distinguishes our work and aligns with recent advancements in generative modeling on Riemannian geometries (Brehmer and Cranmer, 2020; Chen and Lipman, 2023). We demonstrate the effectiveness of our approach through extensive experiments on synthetic

and real datasets on the sphere and torus. The code is available at <https://github.com/Marco-Peg/mvqr>.

2 BACKGROUND

In the upcoming section, we offer a concise introduction to fundamental concepts: quantile regression, c -convexity, and Riemannian Optimal Transport (OT). For a more extensive overview of Riemannian manifolds and all the pertinent formulas employed throughout this paper, we refer to Section A of the Appendix and Do Carmo and Flaherty Francis (1992).

Notation. Throughout, \mathbf{Y} , \mathbf{X} denote random variables and vectors, respectively; deterministic scalars, vectors and matrices are denoted as y , \mathbf{x} , and \mathbf{X} . $p_{\mathbf{Y}}$ denotes the density of the random vector \mathbf{Y} and $p_{(\mathbf{Y}, \mathbf{X})}$ denotes the joint density of \mathbf{X} and \mathbf{Y} . $\mathbf{1}_N$ denotes an N -dimensional vector of ones, \odot denotes the elementwise (Hadamard) product, and $\mathbb{1}_A$ is the indicator function of a set A . $Q_{\mathbf{Y}|\mathbf{X}}(\mathbf{u}; \mathbf{x})$ is the *manifold conditional vector quantile function* M-CVQF (defined in the sequel) of the variable $\mathbf{Y}|\mathbf{X}$, evaluated at the vector quantile level \mathbf{u} , for $\mathbf{X} = \mathbf{x}$. $\mathcal{P}(\mathcal{M})$ denotes the set of probability measures over \mathcal{M} .

Quantile functions. The *quantile function* Q_Y of a scalar-valued random variable $Y \in \mathbb{R}$ is commonly defined as $Q_Y(u) = \inf\{y \in \mathbb{R} : \mathbb{P}[Y \leq y] \geq u\}$. Of note are two important facts concerning Q_Y : (i) it is a unique monotonic map that maps a uniform random variable $U \sim [0, 1]$ to Y ; and (ii) it is the optimal transport map between U and Y where the ground cost is the negative inner product.

Quantile regression. Given a scalar-valued response Y and covariates \mathbf{X} , *quantile regression* (QR) aims to estimate the quantile of a variable $Y|\mathbf{X}$. The standard approach for solving QR is via the minimization of the *pinball loss*. Carlier et al. (2016) showed that QR can be equivalently written as an optimal transport problem between U and Y , with the ground cost as the negative inner product, and subject to additional

mean-independence constraint on the transport plan which depends on \mathbf{X} . While the pinball loss, being a pointwise loss, is not well-defined for vector-valued targets, the optimal transport formulation of QR can be generalized to vector-valued targets by simply modifying the ground cost. This gave rise to *vector quantile regression* (VQR), a multivariate analogue of QR. The goal of the current work is to further generalize VQR to densities defined on manifolds.

c -convexity. In this paper, we consider Riemannian d -dimensional manifolds (\mathcal{M}, g) with the Riemannian metric g , embedded in \mathbb{R}^D . A function $\varphi : \mathcal{M} \rightarrow \mathbb{R} \cup \{+\infty\}$ is c -convex with respect to cost c if and only if it is not identically $+\infty$ and there exists an $\alpha : \mathcal{M} \rightarrow \mathbb{R} \cup \{\pm\infty\}$ such that,

$$\varphi(\mathbf{y}) = \sup_{\mathbf{z} \in \mathcal{M}} \{-c(\mathbf{z}, \mathbf{y}) + \alpha(\mathbf{z})\}.$$

Similarly, a c -concave function $\tilde{\varphi} : \mathcal{M} \rightarrow \mathbb{R} \cup \{-\infty\}$ is defined as,

$$\tilde{\varphi}(\mathbf{y}) = \inf_{\mathbf{z} \in \mathcal{M}} \{c(\mathbf{z}, \mathbf{y}) + \alpha(\mathbf{z})\}.$$

Furthermore, if φ is c -convex, then $-\varphi$ is c -concave, and a convex-combination of c -convex functions is c -convex (Villani, 2021). The c -transform of a function $\varphi : \mathcal{M} \rightarrow \mathbb{R}$ is defined via infimal convolution,

$$\varphi^c(\mathbf{u}) = \inf_{\mathbf{z} \in \mathcal{M}} \{c(\mathbf{z}, \mathbf{u}) - \varphi(\mathbf{z})\}.$$

Lastly, a function φ is c -convex if and only if it satisfies the involution property: $\varphi^{cc} = \varphi$. In the case of squared Euclidean cost, c -convexity reduces to standard convexity, and c -transform becomes the well-known *Legendre-Fenchel transform*.

Riemannian optimal transport. Given two probability distributions and a cost function, the optimal transport (OT) problem consists of finding the mapping that pushes one distribution into the other while minimizing the overall cost. Here we are interested in distributions defined on \mathcal{M} , and a cost function given by $c(\mathbf{y}, \mathbf{z}) = \frac{1}{2}d_{\mathcal{M}}(\mathbf{y}, \mathbf{z})^2$, where $d_{\mathcal{M}}$ is a geodesic distance on \mathcal{M} . We refer to the associated optimal transport problem as *Riemannian OT*. Formally, we define $\mathcal{S}(\mu, \nu)$ as the set of maps from \mathcal{M} to \mathcal{M} , pushing a base probability measure $\mu \in \mathcal{P}(\mathcal{M})$ to a target measure $\nu \in \mathcal{P}(\mathcal{M})$, i.e. $\mathcal{S}(\mu, \nu) = \{\mathbf{s} : \mathcal{M} \rightarrow \mathcal{M} | \mathbf{s}_{\#}\mu = \nu\}$. The Monge problem on \mathcal{M} with cost c consists of finding

$$\inf_{\mathbf{s} \in \mathcal{S}(\mu, \nu)} \int_{\mathcal{M}} c(\mathbf{u}, \mathbf{s}(\mathbf{u})) d\mu(\mathbf{u}). \quad (1)$$

The optimal $\mathbf{s}(\mathbf{u})$ in (1) is the optimal transport map. To avoid the non-convex optimization in Monge's formulation, we consider the Kantorovich relaxation, in

which we search for a joint distribution of the measures instead of a map between them. Denoting by $\Gamma(\mu, \nu)$ the set of joint distributions on $\mathcal{M} \times \mathcal{M}$ that admit μ and ν as marginals, the Kantorovich formulation is given by

$$\inf_{\gamma \in \Gamma(\mu, \nu)} \int_{\mathcal{M} \times \mathcal{M}} c(\mathbf{u}, \mathbf{y}) d\gamma(\mathbf{u}, \mathbf{y}), \quad (2)$$

and its dual formulation is given by,

$$\begin{aligned} \sup_{\varphi, \psi} \int_{\mathcal{M}} \varphi(\mathbf{u}) d\mu(\mathbf{u}) + \int_{\mathcal{M}} \psi(\mathbf{y}) d\nu(\mathbf{y}) \\ \text{s.t. } \varphi(\mathbf{u}) + \psi(\mathbf{y}) \leq c(\mathbf{u}, \mathbf{y}), \end{aligned} \quad (3)$$

where $\varphi, \psi : \mathcal{M} \rightarrow \mathbb{R}$ are bounded and continuous c -concave functions, referred to as the *potential functions*. As shown in Theorem 9 of McCann (2001), the optimal transport map $\mathbf{s} : \mathcal{M} \rightarrow \mathcal{M}$ is a unique minimizer of the Monge problem (1) and can be obtained from the potential φ as,

$$\mathbf{s}(\mathbf{u}) = \exp_{\mathbf{u}}[-\nabla_{\mathbf{u}}\varphi(\mathbf{u})] \quad (4)$$

where ∇ is the intrinsic gradient on \mathcal{M} , and \exp is the exponential map. The optimal dual potentials, referred to as φ^* and ψ^* , are each other's c -transform:

$$\begin{aligned} \varphi^*(\mathbf{y}) &= \inf_{\mathbf{z} \in \mathcal{M}} \{c(\mathbf{z}, \mathbf{y}) - \psi^*(\mathbf{z})\} \\ \psi^*(\mathbf{y}) &= \inf_{\mathbf{z} \in \mathcal{M}} \{c(\mathbf{z}, \mathbf{y}) - \varphi^*(\mathbf{z})\} \end{aligned}$$

and, therefore, both c -concave. For a more extensive overview of Optimal Transport and c -convexity, we refer to Villani et al. (2009) and Santambrogio (2015).

3 QUANTILE REGRESSION ON MANIFOLDS

3.1 Vector quantile functions on manifolds

Let \mathbf{Y} be a random variable supported on the manifold \mathcal{M} , with distribution ν . We define the *manifold uniform distribution* $\mathcal{U}_{\mathcal{M}}$ on $A \subseteq \mathcal{M}$ as the distribution having the density $p_{\mathbf{Y}}(\mathbf{y}) = \mathbb{I}_A(\mathbf{y})/V(A)$, where $V(A)$ denotes the volume of A . The manifold vector quantile function (M-VQF) of \mathbf{Y} , $Q_{\mathbf{Y}} : \mathcal{M} \rightarrow \mathcal{M}$, is defined by the OT map obtained as a solution to the Riemannian OT problem (1), where the base distribution is $\mu = \mathcal{U}_{\mathcal{M}}$ and the target distribution ν is that of \mathbf{Y} (Hallin et al., 2022). We refer to this problem as manifold vector quantile estimation (M-VQE). The resulting M-VQF can therefore be written as,

$$Q_{\mathbf{Y}}(\mathbf{u}) = \exp_{\mathbf{u}}[-\nabla_{\mathbf{u}}\varphi(\mathbf{u})].$$

In the Euclidean case, the exponential map is identity and c -convexity is simply convexity, so the M-VQF

reduces to the Euclidean vector quantile function as defined in Carlier et al. (2016); Chernozhukov et al. (2017); Rosenberg et al. (2023). Moreover, when the random variable is Euclidean and scalar-valued, this definition naturally recovers the one-dimensional quantile function.

3.2 Extension to quantile regression

In the regression case, we estimate the conditional quantile function of $\mathbf{Y}|\mathbf{X} = \mathbf{x}$, denoted by $Q_{\mathbf{Y}|\mathbf{X}}$. This requires solving a family of OT problems parameterized by \mathbf{x} . Following the approach of Carlier et al. (2016) for the Euclidean domain, we average the OT losses over \mathbf{x} . Denoting by ξ the distribution of the joint variable (\mathbf{X}, \mathbf{Y}) and by $\mu = \mathcal{U}_{\mathcal{M}}$, the uniform base distribution, the quantile regression problem on manifolds becomes:

$$\begin{aligned} \sup_{\varphi, \psi} \int_{(\mathcal{X} \times \mathcal{M}) \times \mathcal{M}} (\varphi(\mathbf{u}; \mathbf{x}) + \psi(\mathbf{y}; \mathbf{x})) d\xi(\mathbf{x}, \mathbf{y}) d\mu(\mathbf{u}) \\ \text{s.t. } \forall \mathbf{x}, \mathbf{y}, \mathbf{u} \quad \varphi(\mathbf{u}; \mathbf{x}) + \psi(\mathbf{y}; \mathbf{x}) \leq c(\mathbf{u}, \mathbf{y}) \end{aligned} \quad (5)$$

where $\varphi, \psi : \mathcal{M} \times \mathcal{X} \rightarrow \mathbb{R}$ are bounded and continuous c -concave functions in \mathbf{u} . The manifold conditional vector quantile function (M-CVQF) will thus be a map $\mathcal{M} \times \mathcal{X} \rightarrow \mathcal{M}$ such that:

$$Q_{\mathbf{Y}|\mathbf{X}}(\mathbf{u}, \mathbf{x}) = \exp_{\mathbf{u}}[-\nabla_{\mathbf{u}}\varphi(\mathbf{u}, \mathbf{x})].$$

Discretization. We can discretize Equation 5 by sampling $\{\mathbf{u}_i\}_{i=1}^T \sim \mathcal{U}_{\mathcal{M}}$ from μ and $\{(\mathbf{y}_j, \mathbf{x}_j)\}_{j=1}^N$ from ξ , getting:

$$\begin{aligned} \max_{\varphi, \psi} \sum_{i=1}^T \mu_i \sum_{j=1}^N \xi_j \varphi(\mathbf{u}_i; \mathbf{x}_j) + \sum_{j=1}^N \xi_j \psi(\mathbf{y}_j; \mathbf{x}_j) \\ \text{s.t. } \forall i, j : \varphi(\mathbf{u}_i; \mathbf{x}_j) + \psi(\mathbf{y}_j; \mathbf{x}_j) \leq c(\mathbf{u}_i, \mathbf{y}_j). \end{aligned}$$

By writing one of the potentials as the c -transform of the other we obtain the following optimization problem,

$$\begin{aligned} \max_{\varphi} \sum_{i=1}^T \mu_i \sum_{j=1}^N \xi_j \varphi(\mathbf{u}_i; \mathbf{x}_j) \\ + \sum_{j=1}^N \xi_j \min_{\mathbf{u} \in \mathcal{M}} \{c(\mathbf{u}, \mathbf{y}_j) - \varphi(\mathbf{u}; \mathbf{x}_j)\}, \end{aligned} \quad (6)$$

with $\boldsymbol{\mu} = \frac{1}{T}\mathbf{1}_T$, $\boldsymbol{\xi} = \frac{1}{N}\mathbf{1}_N$ and where $\varphi(\mathbf{u}; \mathbf{x}) : \mathcal{M} \times \mathcal{X} \rightarrow \mathbb{R}$ is a partial c -concave function in \mathbf{u} . We refer to this problem as manifold vector quantile regression (M-VQR). The other potential can be retrieved via the c -transform,

$$\psi(\mathbf{y}; \mathbf{x}_j) = \min_{\mathbf{u} \in \mathcal{M}} \{c(\mathbf{u}, \mathbf{y}) - \varphi(\mathbf{u}; \mathbf{x}_j)\}. \quad (7)$$

3.3 Confidence sets

With the estimated quantile function $Q_{\mathbf{Y}|\mathbf{X}}$, we can compute *confidence sets* on the target distribution. To define a set of nested regions with a μ -probability content of $\tau \in [0, 1]$ on \mathcal{M} , we must first choose a central point, or *pole*, $\boldsymbol{\omega} \in \mathcal{M}$. This point will play the role of the median for ν , around which the contours will be nested. We opted to compute the pole using the *Fréchet mean* of ν , defined as

$$\boldsymbol{\omega} = \arg \min_{\mathbf{y} \in \mathcal{M}} \mathbb{E}_{\mathbf{y} \sim \nu} [c(\mathbf{y}, \mathbf{y})]; \quad (8)$$

other choices are also possible. Under the base distribution, the τ -contour, containing μ -probability of τ and centered at $\boldsymbol{\omega}$, can be defined as

$$\mathcal{C}_{\tau}^{\mathbf{U}} = \{\mathbf{u} \in \mathcal{M} : C_{\boldsymbol{\omega}}^*(\mathbf{u}) = \tau\} \quad (9)$$

where $C_{\boldsymbol{\omega}}^*$ is a function mapping the geodesic distance of \mathbf{u} from the pole $\boldsymbol{\omega}$ to the probability τ . A more detailed explanation on how we define $C_{\boldsymbol{\omega}}^*$ is provided in Section D.2 of the Appendix. The τ conditional quantile contour, under the target distribution of $\mathbf{Y}|\mathbf{X}$, can then be obtained via the image $\mathcal{C}_{\tau}^{\mathbf{Y}|\mathbf{X}} := Q_{\mathbf{Y}|\mathbf{X}}(\mathcal{C}_{\tau}^{\mathbf{U}}; \mathbf{x})$. Figure 1 shows $\mathcal{C}_{\tau}^{\mathbf{Y}|\mathbf{X}}$ for multiple distributions and values of τ .

3.4 Likelihood

Once the conditional quantile function has been estimated, we can use it to also obtain the conditional likelihood function $p_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}; \mathbf{x})$:

$$p_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}; \mathbf{x}) = \frac{1}{V(\mathcal{M})} \cdot \left| \nabla_{\mathbf{y}} Q_{\mathbf{Y}|\mathbf{X}}^{-1}(\mathbf{y}; \mathbf{x}) \right|$$

where ∇ denotes the Jacobian on \mathcal{M} and $\mathbf{U} \sim \mathcal{U}_{\mathcal{M}}$ (with a slight abuse of notation).

4 IMPLEMENTATION DETAILS

Below we discuss details related to the implementation of c -concave functions, partial input c -concave neural networks, and regularization techniques employed.

Discrete c -concave functions. The fundamental building block for our input c -concave potentials are c -concave functions. Following Cohen et al. (2021), we parameterize each c -concave function φ with $\{(\mathbf{z}_i, \alpha_i)\}_{i=1}^M \subset \mathcal{M} \times \mathbb{R}$, where $\alpha_i = \alpha(\mathbf{z}_i)$ are the (learned) values of an implicitly defined function on \mathcal{M} , and \mathbf{z}_i are the (learned) points on which it is sampled. A c -concave function is thus obtained by applying the c -transform to the implicit function $\alpha(z)$:

$$\varphi(\mathbf{u}) = \min_{i=1, \dots, M} \{c(\mathbf{z}_i, \mathbf{u}) + \alpha(\mathbf{z}_i)\}. \quad (10)$$

This formulation results in a piecewise smooth *approximation* of the c -concave potential. Cohen et al. (2021) prove that this discrete approximation of c -concave potential has the expressive power to represent arbitrary c -concave potentials on compact manifolds.

Input c -concave networks. We use two properties of c -concave functions to build input c -concave neural networks: (i) a convex combination of c -concave functions is c -concave, and (ii) applying a concave and monotone function to a c -concave function retains its c -concavity. This is similar in spirit to input-convex neural networks (Amos et al., 2017). An L -layered c -concave function β can be obtained by combining $L + 1$ c -concave functions φ_i , as follows:

$$\begin{aligned}\beta_0(\mathbf{u}) &= \varphi_0(\mathbf{u}) \\ \beta_l(\mathbf{u}) &= (1 - \omega_l)\varphi_l(\mathbf{u}) + \omega_l\sigma(\beta_{l-1}(\mathbf{u})),\end{aligned}\quad (11)$$

where $l \in \{1, \dots, L\}$, $\omega_l \in [0, 1]$ are learnable weights, and $\sigma(s) = \min\{0, s\}$ is a concave monotone function.

Partial input c -concave networks. In the regression setting, we model the potentials as a partially input c -convex neural networks: c -convex in \mathbf{u} but not in \mathbf{x} . A similar approach was also used in Bunne et al. (2022) where they model conditional OT maps on Euclidean domain using partially input convex neural networks. Assuming $\mathcal{X} = \mathbb{R}^k$, we implement a partial c -concave potential $\varphi(\mathbf{u}; \mathbf{x})$ as a non-negative sum of $n + 1$ functions which are c -concave in \mathbf{u} , and apply non-linear trainable transformations on \mathbf{x} :

$$\begin{aligned}\varphi_0(\mathbf{u}; \mathbf{x}) &= \beta_0(\mathbf{u}), & \beta_0 : \mathcal{M} &\rightarrow \mathbb{R} \\ \varphi_i(\mathbf{u}; \mathbf{x}) &= \varphi_{i-1}(\mathbf{u}; \mathbf{x}) + \beta_i(\mathbf{u})^\top \mathbf{g}_i \circ \dots \circ \mathbf{g}_1(\mathbf{x}), & (12) \\ \varphi(\mathbf{u}; \mathbf{x}) &= \varphi_n(\mathbf{u}; \mathbf{x})\end{aligned}$$

with $i \in \{1, \dots, n\}$, $\beta_i : \mathcal{M} \rightarrow \mathbb{R}^{k_i}$ and $\mathbf{g}_i(\mathbf{x}) : \mathbb{R}^{k_{i-1}} \rightarrow \mathbb{R}_+^{k_i}$ ($k_0 = k$) parametrized using neural networks. In particular, each β_i is implemented as a k_i -stack of c -concave functions (10): $\beta_i = (\beta_{i1}, \dots, \beta_{ik_i})$ where $\beta_{ij} : \mathcal{M} \rightarrow \mathbb{R}$ with $j \in \{1, \dots, k_i\}$. Conversely, the \mathbf{g}_i 's can be arbitrary learned embedding functions. For example if \mathbf{X} are images, a CNN-based $\mathbf{g}(\mathbf{x})$ can be used to leverage translation equivariance and the hierarchical nature of image features. In our experiments, we implement the \mathbf{g}_i 's as a multilayer perceptron (MLP) with rectified linear unit (ReLU) activation as the last layer to ensure non-negative values.

Computing convex conjugates. Given a c -concave potential $\varphi(\mathbf{u}, \mathbf{x})$, its convex conjugate is defined by the c -transform,

$$\varphi^c(\mathbf{y}; \mathbf{x}) = \inf_{\mathbf{u} \in \mathcal{M}} \{c(\mathbf{u}, \mathbf{y}) - \varphi(\mathbf{u}; \mathbf{x})\}.$$

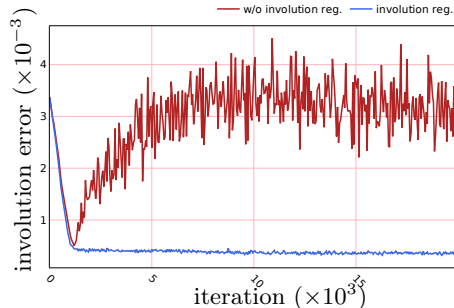


Figure 2: **Impact of involution regularization.** The results are from a M-VQE trained on \mathcal{S}^2 , where the target distribution is a von-Mises distribution and the c -concave potential consists of 3 layers and $\gamma = 0.1$. Involution error is dramatically reduced when training with the involution regularization.

In practice, we compute this conjugate numerically by sampling several points on \mathcal{M} , and evaluating the conjugate explicitly, i.e.,

$$\varphi^c(\mathbf{y}; \mathbf{x}) = \min_{i=1 \dots T} \{c(\mathbf{u}_i, \mathbf{y}) - \varphi(\mathbf{u}_i; \mathbf{x})\}, \quad (13)$$

where $\mathbf{u}_1 \dots \mathbf{u}_T \sim \mathcal{U}_{\mathcal{M}}$.

Both in Equation 10 and 13, we replace the minimum with a soft-minimum to maintain differentiability:

$$\min_{\gamma} (a_1, \dots, a_n) = -\gamma \log \sum_{i=1}^n \exp\left(-\frac{a_i}{\gamma}\right)$$

with $\min_{\gamma} \rightarrow \min$ as $\gamma \rightarrow 0$. Notice that using the soft-minimum is equivalent to performing entropic OT (see Appendix A.4, A.5 in Rosenberg et al. (2023)).

Involution regularization. The use of the soft-minimum, while improving gradient computation over the c -concave function parameters, may compromise the c -concavity of the network. This compromise can lead to a violation of the involution property during training, as already noted in Cohen et al. (2021) and demonstrated in Figure 2. To address this issue, we introduce an extra term in the VQR loss (6) which penalizes deviations from the involution property. Given $\{\mathbf{u}_i\}_{i=1}^T \sim \mathcal{U}_{\mathcal{M}}$ and $\{\mathbf{y}_j, \mathbf{x}_j\}_{i=1}^N \sim P(\mathbf{Y}, \mathbf{X})$, we compute the *involution regularization* as:

$$\sum_{j=1}^N \sum_{i=1}^T \|\varphi^{cc}(\mathbf{u}_i; \mathbf{x}_j) - \varphi(\mathbf{u}_i; \mathbf{x}_j)\|, \quad (14)$$

where φ^{cc} is obtained by computing the numerical conjugate via equation (13) twice on φ . This regularization is conceptually similar to the cycle-consistency regularization proposed by Korotin et al. (2019). However,




Figure 3: **M-VQE approximation of the quantile function $Q_{\mathbf{Y}}$ of the ‘Multimodal von-Mises’ distribution produces nested, smooth, and valid contours and correctly estimates the likelihood function.** Subfigure (a) shows ground truth and estimated likelihood functions. We use the Mollweide projection to plot the whole sphere surface. Subfigure (b) shows τ -contours overlaid on the ground truth samples, for different values of τ . Graph (c) plots the requested coverage level on the horizontal axis and the coverage achieved by the model on the vertical axis.

our approaches differ in two ways: (i) cycle consistency implies involution but only up to a constant, and (ii) we train a single c -convex potential and involution regularizer is used to ensure the c -convexity of the potential, whereas Korotin et al. (2019) train two potentials and cycle consistency is employed for promoting convex-conjugacy between them. Figure 2 demonstrates that this regularizer effectively reduces the involution error, and thus strongly promotes c -concavity of the learned potentials.

Identity initialization. Similarly to Korotin et al. (2019) and Cohen et al. (2021), we pre-train our model to represent the identity map. In practice, we observe that this serves as a good initialization for training.

5 EXPERIMENTS

We validate our method on synthetic and real-world datasets, and on two different manifolds: the sphere \mathcal{S}^2 and the 2-dimensional torus $\mathcal{T}^2 = \mathcal{S}^1 \times \mathcal{S}^1$. Both manifolds have a closed-form expression for distances, exp/log-maps, and the ground cost is the squared geodesic distance. We report all the formulas in Section A of the Appendix. For the synthetic datasets, where we have access to the data generating process, we measure the quality of sampling, likelihood estimation, and confidence sets constructed from the estimated M-CVQFs. In real-world scenarios, where ground truth likelihood are unavailable, we quantify sampling quality and confidence set validity, and visually present the estimated likelihoods.

5.1 Evaluation Metrics

Below we describe the quantitative metrics used to evaluate $\hat{Q}_{\mathbf{Y}|\mathbf{X}}$, the estimated M-CVQF.

Sampling. We assess sampling quality by computing L_1 distance between kernel density estimates (KDE) obtained using samples drawn from the groundtruth and estimated distributions. We employ manifold-specific kernels for the KDEs. The KDE- L_1 distance is measured as:

$$\text{KDE-}L_1(\mathcal{Y}, \mathcal{Y}^{gt}) = |p_{\mathcal{Y}}^{\text{KDE}} - p_{\mathcal{Y}^{gt}}^{\text{KDE}}|,$$

where $\mathcal{Y}, \mathcal{Y}^{gt}$ are equally-sized sets of samples drawn from the estimated and true distribution, respectively. In the conditional setting, we first sample $\{\mathbf{x}_i\}_{i=1}^M \sim \mathbf{X}$. For each \mathbf{x}_i , we sample a set of points $\mathcal{Y}_{\mathbf{x}_i}$ from $\hat{Q}_{\mathbf{Y}|\mathbf{X}}$, a set of points $\mathcal{Y}_{\mathbf{x}_i}^{gt}$ from the true $p_{\mathbf{Y}|\mathbf{X}=\mathbf{x}_i}$, and report the mean KDE- L_1 distance over all \mathbf{x}_i .

Likelihood. When groundtruth likelihood is available, we assess the quality of the model’s likelihood using estimated sample size (ESS), originally proposed by Kish (1968) as also used in Cohen et al. (2021). It is computed as follows:

$$\text{ESS}\% = 100 \times \frac{(\sum_j^N w_j)^2}{N \cdot \sum_j^N w_j^2} \quad \text{with} \quad w_j = \frac{p_{\mathbf{Y}|\mathbf{X}}^{gt}(\mathbf{y}_j, \mathbf{x}_j)}{p_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}_j, \mathbf{x}_j)},$$

where $\{\mathbf{y}_j, \mathbf{x}_j\}_{j=1}^N \sim p(\mathbf{Y}, \mathbf{X})$, $p_{\mathbf{Y}|\mathbf{X}}^{gt}$ and $p_{\mathbf{Y}|\mathbf{X}}$ are the groundtruth and estimated conditional likelihoods, respectively.

Confidence sets. Estimating the validity of a confidence set $\mathcal{C}_{\tau}^{\mathbf{Y}|\mathbf{X}}$ requires computing $\mathbb{P}[\mathbf{y} \in \mathcal{C}_{\tau}^{\mathbf{Y}|\mathbf{X}}]$, which is non-trivial because confidence sets on the target distribution assume an arbitrary shape on the manifold. However, checking whether a point is inside a τ confidence set for \mathbf{U} , i.e., $\mathcal{C}_{\tau}^{\mathbf{U}}$, is straightforward: it can be verified by checking that $C_{\omega}^*(\mathbf{u}) \leq \tau$. To exploit this property we use the inverse M-CVQF, which maps $\mathbf{Y}|\mathbf{X}$ back to $\mathbf{U} \sim \mathcal{U}_{\mathcal{M}}$, computed as

Table 1: **M-VQR approximation of $Q_{\mathbf{Y}|\mathbf{X}}$ allows building confidence sets with good marginal coverage, to perform good quality conditional sampling, and to approximate the likelihood accurately.** The table reports on the left marginal coverage values for different confidence sets built using the estimated quantile function $Q_{\mathbf{Y}|\mathbf{X}}$, averaged over 50 values of the continuous conditioning \mathbf{X} , with relative standard deviation. KDE- L_1 values in the are also averages over 50 values of \mathbf{X} . The last column reports the $ESS\%$ computed over samples drawn from the whole distribution $\mathbf{Y}|\mathbf{X}$.

\mathcal{M}	$\mathbf{Y} \mathbf{X}$	Coverage (%)			Mean error	KDE- L_1 ($\times 10^{-4}$)	$ESS\%$
		20	30	60			
\mathcal{S}^2	Cond. Multimodal	20.22 ± 0.44	40.28 ± 0.56	60.22 ± 0.49	0.44 ± 0.34	14.2 ± 2.09	94.93%
	Scaled Star	19.50 ± 0.51	38.40 ± 0.63	59.35 ± 0.70	0.86 ± 0.64	2.25 ± 0.99	89.30%
	Scaled Heart	20.57 ± 0.46	40.82 ± 0.60	60.02 ± 0.53	0.60 ± 0.42	4.46 ± 1.18	89.84%
\mathcal{T}^2	Cond. Multimodal	20.06 ± 0.43	40.68 ± 0.61	60.46 ± 0.68	0.54 ± 0.43	17.5 ± 2.92	96.82%
	Scaled Star	19.81 ± 0.50	39.86 ± 0.55	60.39 ± 0.61	0.47 ± 0.35	5.85 ± 1.92	84.79%
	Scaled Heart	19.95 ± 0.47	40.50 ± 0.65	61.18 ± 0.60	0.61 ± 0.49	7.93 ± 2.14	88.05%

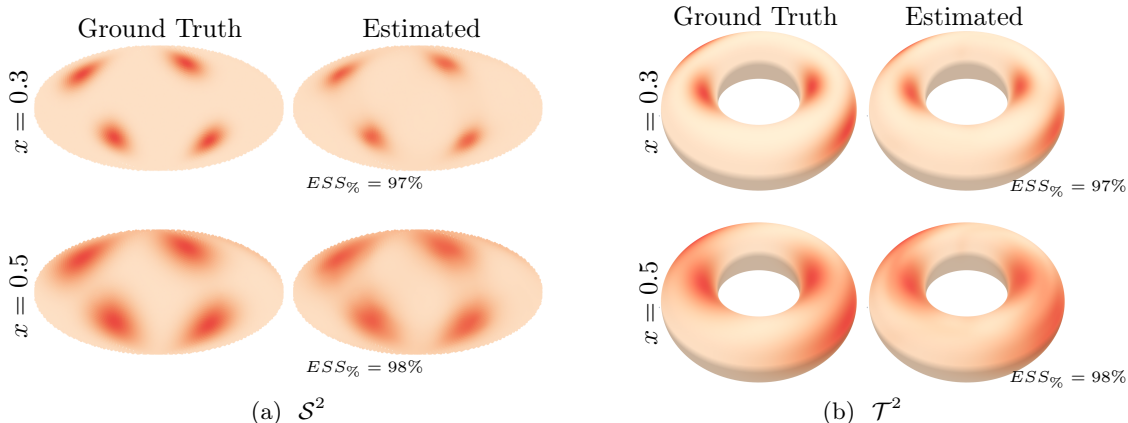


Figure 4: **Likelihood function $p_{\mathbf{Y}|\mathbf{X}}$ for the ‘Conditional Multimodal’ distribution.** The covariate \mathbf{X} controls the scale of the distribution. $ESS\%$ values are also reported. We use the mollweide projection to plot the whole sphere surface.

$Q_{\mathbf{Y}|\mathbf{X}}^{-1}(\mathbf{y}; \mathbf{x}) = \exp_{\mathbf{y}}[-\nabla_{\mathbf{y}}\psi(\mathbf{y}; \mathbf{x})]$. The coverage of the confidence set $\mathcal{C}_{\tau}^{\mathbf{Y}|\mathbf{X}}$ is then estimated by:

$$\begin{aligned} \mathbb{P}[\mathbf{y} \in \mathcal{C}_{\tau}^{\mathbf{Y}|\mathbf{X}}] &= \mathbb{P}[Q_{\mathbf{Y}|\mathbf{X}}^{-1}(\mathbf{y}; \mathbf{x}) \in \mathcal{C}_{\tau}^{\mathbf{U}}] \\ &= \mathbb{E}_{\mathbf{U}}[\mathbb{I}_{\mathcal{C}_{\tau}^{\mathbf{U}}}(\mathbf{U}) \leq \tau] \approx \frac{1}{N} \sum_{i=1}^N \mathbb{I}_{\mathcal{C}_{\tau}^{\mathbf{U}}}(\mathbf{u}_i) \leq \tau \end{aligned}$$

where $\mathbf{u}_i = Q_{\mathbf{Y}|\mathbf{X}}^{-1}(\mathbf{y}_i; \mathbf{x}_i)$.

5.2 Datasets

Synthetic datasets. For synthetic data experiments, we defined three data generating processes. The ‘Conditional Multimodal’ distribution is a mixture of four von-Mises distributions on \mathcal{M} , conditioned on $\mathbf{X} \in \mathbb{R}$ which controls the covariance of the distributions. ‘Scaled Heart’ and ‘Scaled Star’ are uniform distributions on a subset $A \subset \mathcal{M}$ which is heart-shaped and star-shaped, respectively. The conditioning variable $\mathbf{X} \in \mathbb{R}$ controls

the scale of A . All aforementioned distributions are defined both on \mathcal{S}^2 and \mathcal{T}^2 .

Real-world datasets. As a candidate for distributions defined on \mathcal{S}^2 , we consider the ‘Continental drift’ dataset, consisting of continental drift measured over 26 eras. We consider the era as the conditioning variable, and continental drift as the target variable. The source maps are taken from the 2020 Colorado Plateau Geosystems ¹. We convert each Mollweide earth image into a spherical point cloud and into Euclidean coordinates.

As a real-world use case on the torus $\mathcal{T} = \mathcal{S}^1 \times \mathcal{S}^1$, we consider a dataset of local dihedral angles ($\varphi, \psi \in [0, 2\pi]$) measured in the backbone of folded protein structures. Pairs of angles can be naturally represented as points of on a torus, because each angle is represented by a point on \mathcal{S}^1 . Protein structures are composed of

¹<https://deeptimemaps.com/>

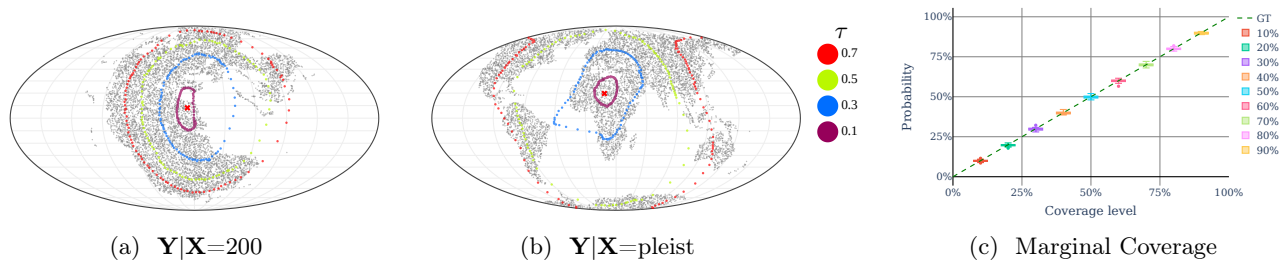


Figure 5: τ -confidence sets constructed with M-VQR on the ‘Continental Drift’ dataset are smooth, nested, and valid. Subfigures (a) and (b) report τ -contours overlayed on the ground truth samples, for different values of τ ; each subfigure represents conditioning on a different era. Mollweide projection is used to visualize the whole sphere. Graph (c) shows the coverage achieved by the model as a function of the requested coverage level, averaged over the different conditionings with relative confidence bars.




Figure 6: τ -contours constructed with M-VQR on the ‘Dihedral Angles’ dataset are smooth, nested, and valid. Subfigures (a), (b) and (c) report τ -contours overlayed on the ground truth samples, for different values of τ ; each subfigure depicts conditioning by a different amino acid. The dihedral angles are plotted on a “flat torus”, with the axes representing the dihedral angles in degrees. The wrap-around, due to the toroidal domain, is clearly visible in the contours. Graph (d) shows the the coverage achieved by the model as a function of the requested coverage level, averaged over the different conditionings with relative confidence bars.

chains of common amino acids (AAs), with each AA in the chain assuming different dihedral angles depending on its location, chemistry, context, and other biological factors. We use the AA identity as the conditioning variable, and the corresponding dihedral angles as the target. Note that here X is categorical, with 20 possible values indicating the AA identity. We use the dihedral angle dataset curated by Rosenberg et al. (2022) and group the dihedral angles by their AA identity.

5.3 Synthetic data experiments

Below we report results both for unconditional estimation (approximating Q_Y , without covariates) and for regression (approximating $Q_{Y|X}$).

Quantile estimation. Though it is not the primary focus of this paper, evaluating the performance of our approach in the simpler quantile estimation scenario is worthwhile, because until now the only existing approach for this problem was that of Hallin et al. (2022), which is limited to the sphere and neither scalable nor differentiable. Figure 3 visually and quantitatively

demonstrates that the estimated M-VQF results in smooth, nested, and valid contours. Both the likelihood visualization and the ESS value (93%) indicate a close approximation to the ground truth likelihood. We also report a KDE- L_1 distance of 3.9×10^{-3} ; compared to an upper bound of 0.7×10^{-3} when measured between two independent pairs of groundtruth samples. This indicates good sampling quality of the M-VQF.

With regard to scalability, in Section C.1 of the Appendix, we demonstrate that our approach, based on the relaxed dual OT problem, is substantially more scalable than the previous approach of Hallin et al. (2022), which involves solving a discrete OT problem as a linear-program.

Quantile regression. Quantitative results for coverage, sampling, and likelihood in the regression setting are reported in Table 1. We observe that the confidence sets constructed by estimated M-CVQFs result in a maximum coverage violation of 1.5% with respect to the nominal coverage. KDE- L_1 shows an order of magnitude of 10^{-3} at maximum against a magnitude

of 10^{-4} for KDE- L_1 between ground truth samples, indicating good sampling. ESS is between 84%-96%, demonstrating that the conditional likelihoods computed from the estimated M-CVQFs are reasonably accurate compared to the groundtruth.

Figure 1 provides a visualization of conditional contours for the ‘Scaled Heart’ and ‘Scaled Star’ distributions. Figures 13c and 13f depict the estimated conditional likelihoods for the ‘Conditional Multimodal’ distribution dataset on \mathcal{S}^2 and \mathcal{T}_2 , respectively. Additional visualizations are provided in Section E of the Appendix.

5.4 Real data experiments

Figure 5 presents results on the ‘Continental Drift’ dataset, reporting the coverage and confidence sets visualizations for two different conditioning values. Figure 6 reports the same results for the ‘Dihedral Angles’ dataset. In both cases, the resulting contours are nested and smooth, as desired. The coverage values are almost perfectly aligned with the ground truth, and demonstrate low variance over the conditioning variables. In Figure 6, the wrap-around of the confidence sets due to the manifold structure of the domain is clearly visible. The mean KDE- L_1 for the ‘Continental Drift’ dataset is $(1.72 \pm 0.26) \times 10^{-3}$ against $(9.45 \pm 1.05) \times 10^{-4}$ from the ground truth samples; while for the ‘Dihedral Angles’ dataset we obtain a mean KDE- L_1 of $(2.61 \pm 0.56) \times 10^{-3}$ against $(7.34 \pm 2.88) \times 10^{-4}$. The likelihood plots for both datasets are provided in Section E of the Appendix.

6 DISCUSSION AND CONCLUSIONS

Our work provides the first formulation of conditional vector quantile functions on manifolds, by extending non-linear VQR to non-Euclidean domains, together with conditional OT and quantile regression on manifolds. Our key contributions are the novel formulation of nonlinear VQR as a Riemannian OT problem, its parametrization with partial input c -concave neural networks, and the involution regularization approach for training.

One potential limitation of this work is that the estimated potential functions might not be c -concave, as this property is only promoted, but not enforced, through the regularized objective. In future studies, we hope to analyze it theoretically and provide ways to overcome this limitation. Another avenue for exploration is extending the proposed approach to domains possessing closed-form formulations for computing c -concave functions, such as Lie groups.

In summary, our approach enables the estimation of conditional quantiles and construction of confidence sets on general manifolds for which the exponential map is known by fitting data sampled directly from the joint distribution. These capabilities may open the door to exciting new applications in diverse domains such as pose estimation, weather modeling, and protein structure prediction. We believe that M-VQR is thus a powerful and useful new addition to the toolbox of directional statistics.

Acknowledgments

S.V., A.A.R., and A.M.B. were partially supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 863839), by the Council For Higher Education - Planning & Budgeting Committee, and by the Israeli Smart Transportation Research Center (ISTRC). I.T., M.P. and E.R. were partially supported by the ERC grant no.802554 (SPECGEO), PRIN 2020 project no.2020TA3K9N (LEGO.AI), and PNRR MUR project PE0000013-FAIR.

References

- Amos, B., Xu, L., and Kolter, J. Z. (2017). Input convex neural networks. In *International Conference on Machine Learning*, pages 146–155. PMLR.
- Brehmer, J. and Cranmer, K. (2020). Flows for simultaneous manifold learning and density estimation. *Advances in Neural Information Processing Systems*, 33:442–453.
- Bunne, C., Krause, A., and Cuturi, M. (2022). Supervised training of conditional monge maps. *Advances in Neural Information Processing Systems*, 35:6859–6872.
- Carlier, G., Chernozhukov, V., and Galichon, A. (2016). Vector quantile regression: An optimal transport approach. *Annals of Statistics*, 44(3):1165–1192.
- Chen, R. T. and Lipman, Y. (2023). Riemannian flow matching on general geometries. *arXiv preprint arXiv:2302.03660*.
- Chernozhukov, V., Galichon, A., Hallin, M., and Henry, M. (2017). Monge–Kantorovich depth, quantiles, ranks and signs. *The Annals of Statistics*, 45(1):223 – 256.
- Cohen, S., Amos, B., and Lipman, Y. (2021). Riemannian convex potential maps. In *International Conference on Machine Learning*, pages 2028–2038. PMLR.
- Do Carmo, M. P. and Flaherty Francis, J. (1992). *Riemannian geometry*, volume 6. Springer.

- Hallin, M., Liu, H., and Verdebout, T. (2022). Nonparametric measure-transportation-based methods for directional data. arXiv preprint arXiv:2212.10345.
- Kish, L. (1968). Survey sampling. Biometrische Zeitschrift, 10(1):88–89.
- Koenker, R. and Bassett, G. (1978). Regression Quantiles. Econometrica, 46(1):33.
- Korotin, A., Egiazarian, V., Asadulaev, A., Safin, A., and Burnaev, E. (2019). Wasserstein-2 generative networks. arXiv preprint arXiv:1909.13082.
- Lagona, F. (2018). Correlated cylindrical data. In Applied Directional Statistics, pages 61–76. Chapman and Hall/CRC.
- McCann, R. J. (2001). Polar factorization of maps on riemannian manifolds. Geometric & Functional Analysis GAFA, 11(3):589–608.
- Pennec, X., Sommer, S., and Fletcher, T. (2019). Riemannian geometric statistics in medical image analysis. Academic Press.
- Rosenberg, A. A., Marx, A., and Bronstein, A. M. (2022). Codon-specific ramachandran plots show amino acid backbone conformation depends on identity of the translated codon. Nature communications, 13(1):2815.
- Rosenberg, A. A., Vedula, S., Romano, Y., and Bronstein, A. M. (2023). Fast nonlinear vector quantile regression. International Conference on Learning Representations (ICLR).
- Santambrogio, F. (2015). Optimal transport for applied mathematicians. Birkäuser, NY, 55(58-63):94.
- Scher, S. and Messori, G. (2020). Spherical convolution and other forms of informed machine learning for deep neural network based weather forecasts. arXiv preprint arXiv:2008.13524.
- Villani, C. (2021). Topics in optimal transportation, volume 58. American Mathematical Soc.
- Villani, C. et al. (2009). Optimal transport: old and new, volume 338. Springer.

APPENDIX

A DIFFERENTIAL GEOMETRY

In this Section, we provide an overview of the main concepts of differential geometry used in the main paper.

Manifolds. In this paper, we consider Riemannian d -dimensional manifolds (\mathcal{M}, g) with the Riemannian metric g , embedded in \mathbb{R}^D . Given a point $\mathbf{y} \in \mathcal{M}$, the *tangent space* $T_{\mathbf{y}}\mathcal{M}$ is defined as the linear subspace $T_{\mathbf{y}}\mathcal{M} = \{\mathbf{v} \in \mathbb{R}^D : \mathbf{v}^\top \mathbf{y} = 0\}$. The Riemannian metric defines an inner product $\langle \cdot, \cdot \rangle_g : T_{\mathbf{y}}\mathcal{M} \times T_{\mathbf{y}}\mathcal{M} \mapsto \mathbb{R}$ on the tangent space, which induces a geodesic distance $d_{\mathcal{M}}(\mathbf{y}, \mathbf{z})$ between every pair of points $\mathbf{y}, \mathbf{z} \in \mathcal{M}$, as the minimum length of a curve connecting the two points,

$$d_{\mathcal{M}}(\mathbf{y}, \mathbf{z}) = \inf_{\gamma} \int_0^1 \|\dot{\gamma}(t)\|_g dt,$$

where $\gamma : [0, 1] \rightarrow \mathcal{M}$ and $\gamma(0) = \mathbf{y}$, $\gamma(1) = \mathbf{z}$.

The *exponential map* of a manifold projects an infinitesimal displacement of \mathbf{y} along the tangent vector \mathbf{v} back to the manifold. Given a point $\mathbf{y} \in \mathcal{M}$ with a tangent vector $\mathbf{v} \in T_{\mathbf{y}}\mathcal{M}$, and given the unique geodesic $\gamma : [0, 1] \rightarrow \mathcal{M}$ such that $\gamma(0) = \mathbf{y}$ and $\dot{\gamma}(0) = \mathbf{v}$, the exponential map at \mathbf{y} is defined as $\exp_{\mathbf{y}}(\mathbf{v}) = \gamma(1)$. The tangent space inner product structure also allows one to define the intrinsic gradient and Jacobian over the manifold.

Manifold uniform distribution. Given the volume measure $d\mathcal{M}(\mathbf{y})$ representing the infinitesimal volume element at each point \mathbf{y} of the manifold, a random variable \mathbf{Y} follows a manifold uniform distribution \mathcal{U}_A on the bounded subset $A \subseteq \mathcal{M}$ if its probability density function (PDF) is constant within A :

$$p_{\mathbf{Y}}(\mathbf{y}) = \frac{\mathbb{I}_A(\mathbf{y})}{V(A)} \quad (15)$$

where $V(A)$ is the volume of the set A with respect to the volume measure $d\mathcal{M}$. The uniform distribution on the manifold, with respect to the volume measure, assigns probabilities to subsets of the manifold based on their intrinsic volumes.

Sphere. On the n -sphere \mathcal{S}^n , the exponential map and the intrinsic distance are provided as closed-form expressions. If $\mathbf{y}, \mathbf{u} \in \mathcal{S}^n$ and $\mathbf{v} \in T_{\mathbf{y}}\mathcal{S}^n$,

$$\exp_{\mathbf{y}}^{\mathcal{S}^n}(\mathbf{v}) = \mathbf{y} \cos(\|\mathbf{v}\|) + \frac{\mathbf{v}}{\|\mathbf{v}\|} \sin(\|\mathbf{v}\|) \quad (16)$$

$$d_{\mathcal{S}^n}(\mathbf{y}, \mathbf{u}) = \arccos(\mathbf{u}^\top \mathbf{v}), \quad (17)$$

where $\|\cdot\|$ is the standard Euclidean norm.

Torus. The torus \mathcal{T}^2 can be defined as a product manifold between two 1-sphere: $\mathcal{T}^2 = \mathcal{S}^1 \times \mathcal{S}^1$. On general product manifolds of the form $\mathcal{M} = \mathcal{M}_1 \times \dots \times \mathcal{M}_l$, the squared intrinsic distance is simply

$$d_{\mathcal{M}}^2(\mathbf{y}, \mathbf{u}) = d_{\mathcal{M}_1}^2(\mathbf{y}_1, \mathbf{u}_1) + \dots + d_{\mathcal{M}_l}^2(\mathbf{y}_l, \mathbf{u}_l). \quad (18)$$

where $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_l)$, and $\mathbf{u}_j \in \mathcal{M}_j$, $j \in [l]$ (and similarly for \mathbf{y}). The exponential map on the product manifold is the cartesian product of exponential maps on the individual manifolds. Therefore, the exponential map and intrinsic distance on the torus \mathcal{T}^2 is defined as:

$$\exp_{\mathbf{y}}^{\mathcal{T}^2}(\mathbf{v}) = \exp_{\mathbf{y}_1}^{\mathcal{S}^1}(\mathbf{v}_1) \times \exp_{\mathbf{y}_2}^{\mathcal{S}^1}(\mathbf{v}_2) \quad (19)$$

$$d_{\mathcal{T}^2}(\mathbf{y}, \mathbf{u}) = \sqrt{d_{\mathcal{S}^1}^2(\mathbf{y}_1, \mathbf{u}_1) + d_{\mathcal{S}^1}^2(\mathbf{y}_2, \mathbf{u}_2)}. \quad (20)$$

B DERIVATION OF MANIFOLD VQR LOSS FUNCTION

The dual formulation of the Kantorovich problem is given by

$$\begin{aligned} & \sup_{\varphi, \psi} \int_{\mathcal{M}} \varphi(\mathbf{u}) d\mu(\mathbf{u}) + \int_{\mathcal{M}} \psi(\mathbf{y}) d\nu(\mathbf{y}) \\ & \text{s.t.} \quad \varphi(\mathbf{u}) + \psi(\mathbf{y}) \leq c(\mathbf{u}, \mathbf{y}), \end{aligned} \quad (21)$$

due to the linearity of integration we have

$$\begin{aligned} & \sup_{\varphi, \psi} \int_{\mathcal{M} \times \mathcal{M}} (\varphi(\mathbf{u}) + \psi(\mathbf{y})) d\mu(\mathbf{u}) d\nu(\mathbf{y}) \\ & \text{s.t.} \quad \varphi(\mathbf{u}) + \psi(\mathbf{y}) \leq c(\mathbf{u}, \mathbf{y}) \end{aligned}$$

In the regression setting, φ and ψ are c -concave functions, parametric in \mathbf{x} and concave in \mathbf{u} , and we denote the measure corresponding to the joint distribution of \mathbf{X}, \mathbf{Y} as $\xi(\mathbf{x}, \mathbf{y})$. The manifold vector quantile regression loss is simply obtained by taking an expectation of Equation 21 with respect to \mathbf{x} ,

$$\begin{aligned} & \sup_{\varphi, \psi} \int_{\mathcal{X} \times \mathcal{M} \times \mathcal{M}} \varphi(\mathbf{u}; \mathbf{x}) d\xi(\mathbf{x}, \mathbf{y}) d\mu(\mathbf{u}) + \int_{\mathcal{X} \times \mathcal{M}} \psi(\mathbf{y}; \mathbf{x}) d\xi(\mathbf{x}, \mathbf{y}) \\ & \text{s.t.} \quad \varphi(\mathbf{u}; \mathbf{x}) + \psi(\mathbf{y}; \mathbf{x}) \leq c(\mathbf{u}, \mathbf{y}). \end{aligned}$$

In the finite sample setting where $\{\mathbf{u}\}_{i=1}^T \sim \mathcal{U}_{\mathcal{M}}$, $\{\mathbf{x}_j, \mathbf{y}_j\} \sim p_{\{\mathbf{X}, \mathbf{Y}\}}$, the above can be written as

$$\begin{aligned} & \max_{\varphi, \psi} \sum_{j=1}^N \xi_j \sum_{i=1}^T \mu_i \varphi(\mathbf{u}_i; \mathbf{x}_j) + \sum_{j=1}^N \xi_j \psi(\mathbf{y}_j; \mathbf{x}_j) \\ & \text{s.t.} \quad \forall i, j : \quad \varphi(\mathbf{u}_i; \mathbf{x}_j) + \psi(\mathbf{y}_j; \mathbf{x}_j) \leq c(\mathbf{u}_i, \mathbf{y}_j), \end{aligned} \quad (22)$$

where $\boldsymbol{\mu} = \frac{1}{T} \mathbf{1}_T$, $\boldsymbol{\xi} = \frac{1}{N} \mathbf{1}_N$, are measures corresponding to the respective sample densities.

Estimation. In the main paper, we evaluated the performance of example of manifold vector quantile estimation. For convenience, below we provide the finite-sample version of the loss function in the unconditional case. Given samples $\{\mathbf{u}_i\}_{i=1}^T \sim \mathcal{U}_{\mathcal{M}}$ and $\{\mathbf{y}_j\}_{j=1}^N \sim \mathbf{Y}$ from a random target variable \mathbf{Y} , the dual formulation of the Kantorovich problem can be discretized as

$$\begin{aligned} & \max_{\varphi, \psi} \sum_{i=1}^T \mu_i \varphi(\mathbf{u}_i) + \sum_{j=1}^N \nu_j \psi(\mathbf{y}_j) \\ & \text{s.t.} \quad \forall i, j : \quad \varphi(\mathbf{u}_i) + \psi(\mathbf{y}_j) \leq c(\mathbf{u}_i, \mathbf{y}_j), \end{aligned} \quad (23)$$

where $\boldsymbol{\mu} = \frac{1}{T} \mathbf{1}_T$, $\boldsymbol{\nu} = \frac{1}{N} \mathbf{1}_N$.

By writing one of the potentials in terms of the other leveraging the c -transform, we obtain the following max-min optimization problem,

$$\max_{\varphi} \sum_{i=1}^T \mu_i \varphi(\mathbf{u}_i) + \sum_{j=1}^N \nu_j \min_{\mathbf{u} \in \mathcal{M}} \{c(\mathbf{u}, \mathbf{y}_j) - \varphi(\mathbf{u})\}. \quad (24)$$

C COMPARISON TO PRIOR WORKS

In this section, we compare our method to two prior works, Hallin et al. (2022) and Cohen et al. (2021). We evaluate the scalability of our method in comparison to Hallin et al. (2022). We then study the impact of different training loss functions, to offer insights into how our approach compares to Cohen et al. (2021). Through this comparative exploration, we aim to provide a holistic view of our method's strengths and capabilities in relation to the existing state-of-the-art. It is important to note that these comparisons are carried out on manifold vector quantile estimation (M-VQE), as neither aforementioned works train conditional maps (M-VQR).




Figure 7: **Probability contained in the contours computed using the inner distance $d_{\mathcal{M}}(\omega, \mathbf{u})$ (left) and the mapping function $C_{\omega}^*(\mathbf{u})$ (right).** The distances are normalized to be in the range $[0, 1]$. The mapping function C_{ω}^* creates a linear dependency between the distance of the points in a contour and the amount of probability contained.

C.1 Comparison to Hallin et al. (2022)

Hallin et al. (2022) propose solving the OT problem resulting from manifold vector quantile estimation (M-VQE) as a linear assignment problem with the ground cost set to be the squared geodesic distance. Given $\{\mathbf{u}_i\}_{i=1}^T \sim \mathcal{U}_{\mathcal{M}}$ and $\{\mathbf{y}_j\}_{j=1}^N \sim P_{\mathbf{Y}}$, the primal OT formulation proposed by Hallin et al. (2022) solves the following optimization problem

$$\begin{aligned} \min_{\mathbf{\Pi} \geq 0} \quad & \sum_{j=1}^N \sum_{i=1}^T \pi_{ij} c(\mathbf{u}_i, \mathbf{y}_j) \\ \text{s.t.} \quad & \mathbf{\Pi} \mathbf{1}_N = \mathbf{1}_T, \mathbf{\Pi}^{\top} \mathbf{1}_T = \mathbf{1}_N, \end{aligned}$$

where $\mathbf{\Pi} \in \mathbb{R}^{T \times N}$ is the assignment matrix, and $c(\mathbf{u}_i, \mathbf{y}_j) = d_{\mathcal{M}}^2(\mathbf{u}_i, \mathbf{y}_j)$. This approach has several limitations. First, the aforementioned problem solves a *discrete* optimal transport problem, thus it recovers the transport map only at pre-specified points. This discrete representation of the map does not allow computation of the likelihood which is defined as the determinant of the Jacobian of the inverse map. In contrast, we solve a *continuous* optimal transport problem, we recover a continuous, differentiable, and invertible formulation of the transport map. Second, the primal OT formulation is solved using a linear program solver. As a result, their approach scales poorly with the number of samples, both in the number of optimization variables and run-time. Instead our approach solves the relaxed formulation of the dual optimal transport problem which is amenable to gradient-based optimization and scales to large sample sizes. For example, when $N, T = 10000$, our M-VQE solver converges in 8 minutes, whereas the linear program solver, which estimates over 100 million parameters, does not converge even in 3 hours. Finally, Hallin et al. (2022) solve only the estimation problem, whereas our approach solves the more general regression problem.

C.2 Comparison to Cohen et al. (2021)

While we adopt the parametrization strategy for c -concave function introduced by Cohen et al. (2021), our work differs from Cohen et al. (2021) in two crucial ways: (i) we employ a different training strategy; (ii) we retrieve a family of maps conditioned on covariates \mathbf{X} , whereas Cohen et al. (2021) retrieve a single map.

Cohen et al. (2021) propose two alternative strategies to train a normalizing flow between the base and target distributions: (i) training only the forward map by minimizing the KL-divergence w.r.to the groundtruth target samples, (ii) training only the inverse map by performing maximum likelihood in with respect to the base distribution. In contrast to these both training strategies, we retrieve both forward and inverse maps *simultaneously* by solving the optimal transport problem.

In Figure 8, we compare these three training strategies, both visually and quantitatively, by evaluating the likelihood and sampling quality on the ‘Multimodal von-Mises’ distribution. Across all strategies, to allow for a fair comparison, we train a single c -concave potential discretized at 200 points ($\{z_i, \alpha_i\}_{i=1}^{200}$). Likelihoods are computed by leveraging the Jacobian of the inverse of the quantile function $Q_{\mathbf{Y}}^{-1}$. For RCPM with KL-divergence (RCPM-KL), we train the forward potential φ and then calculate the backward potential using the c -transform to compute the likelihood. Conversely, for RCPM with maximum likelihood (RCPM-LH), we sample points using the c -transform of the trained backward potential. The results highlight that our method attains a comparable ESS% and the smallest KDE- L_1 error. Notably, RCPM-LH achieves a nearly perfect likelihood, possibly attributed to the chosen loss function. However, in terms of sampling quality, our model, M-VQE, remains superior.

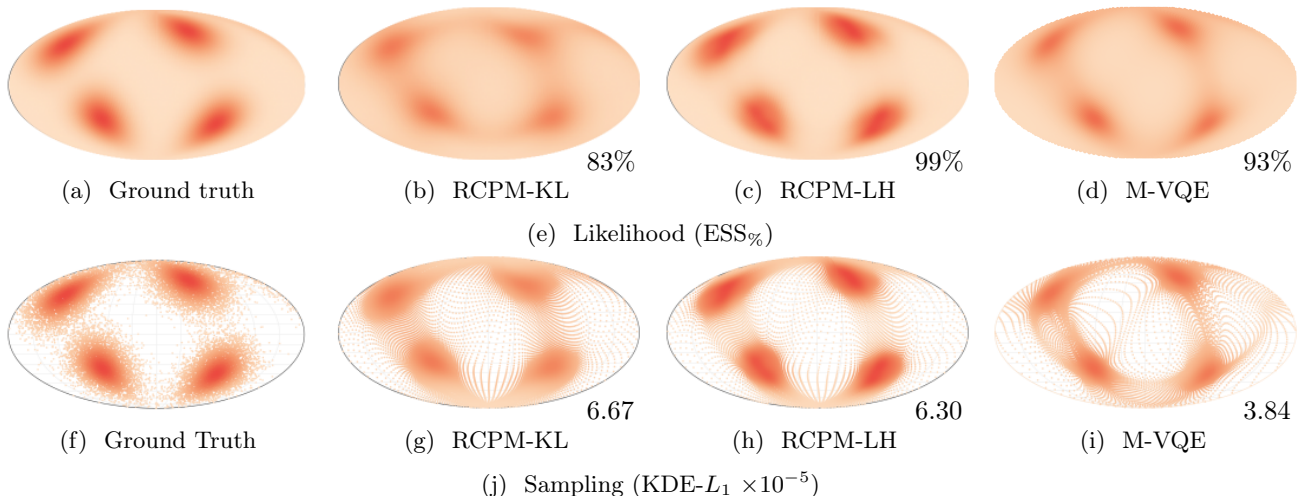


Figure 8: **Likelihood and Sampling on the ‘Multimodal von-Mises’ distribution.** We compare the results obtain with our M-VQE and the RCPM method (Cohen et al., 2021) trained with KL-divergence (RCPM-KL) and Maximum log-likelihood (RCPM-LH). We also report the ESS% and KDE- L_1 ($\times 10^{-5}$) values.

D EXPERIMENTAL DETAILS

D.1 Hyperparameters

We train all models for 5×10^4 iterations, using Adam optimizer with a learning rate set at 10^{-3} . We discretize the c-concave potentials using 200 points (α 's) and approximate the potentials with 4-6 layered c-concave functions.

Multimodal von-Mises. For the Multimodal von-Mises dataset, we train an M-CVQF with a learning rate of 0.001, incorporating a single c-concave β potential.

Synthetic datasets. In the case of the synthetic datasets, our approach involves training an M-CVQF with a learning rate of 0.001, implementing a c-concave β potential comprising 4 layers. The c-concave function is further augmented with inner sizes of [2, 4, 8].

Continental drift. For the Continental Drift dataset, we train an M-CVQF using a learning rate of 0.0005. The c-concave β potential is composed of 6 layers, while the c-concave function encompasses inner sizes of [4, 8, 12, 16, 20, 24]. The dataset contains 3.8M samples.

Dihedral angles. In the context of the Dihedral Angles dataset, our approach involves training an M-CVQF with a learning rate of 0.0005. Similar to the Continental Drift dataset, the c-concave β potential consists of 6 layers, and the c-concave function's inner sizes are specified as [4, 8, 12, 16, 20]. The dataset contains 350k samples.

Spherical uniform distribution \mathcal{U}_{S^n} . We compute the manifold uniform distribution on the n -sphere sampling each component of the vector $\mathbf{v} \in \mathbb{R}^n$ from a random uniform distribution and then normalizing \mathbf{v} such that $\|\mathbf{v}\| = 1$.

Torus uniform distribution \mathcal{U}_{T^2} . We compute the manifold uniform distribution on the torus sampling from the spherical uniform distributions \mathcal{U}_{S^1} on the S^1 and then concatenating the results: $\mathbf{u} = [\mathbf{u}_1, \mathbf{u}_2] \sim \mathcal{U}_{T^2}$ with $\mathbf{u}_1, \mathbf{u}_2 \sim \mathcal{U}_{S^1}$.

D.2 Creating confidence sets

Given a point $\omega \in \mathcal{M}$, we define a confidence set with a level of confidence $(1 - \tau)$ on a manifold uniform distribution as the set of points contained in a τ -contour with pole ω :

$$\mathbb{C}_{(1-\tau)}^{\mathbf{U}} = \{\mathbf{u} \in \mathcal{M} : C_{\omega}^*(\mathbf{u}) \leq \tau\} \quad (25)$$

C_{ω}^* is a function that maps distances $d_{\mathcal{M}}(\omega, \mathbf{u})$, with $\mathbf{u} \in \mathcal{M}$, to the probabilities $\tau \in [0, 1]$. We build this function empirically for each pole ω . We compute the set of points $\mathcal{C}_{\kappa}^* = \{\mathbf{u} \in \mathcal{M} : d_{\mathcal{M}}(\omega, \mathbf{u}) = \kappa\}$ and the amount of



 Figure 9: τ -contours of the Manifold Uniform Distributions $\mathcal{U}_{\mathcal{M}}$

 Table 2: **KDE- L_1** ($\times 10^{-4}$) **computed over samples of M-VQR and GT.** We report the values for all the conditional distributions considered in the paper.

\mathcal{M}	$\mathbf{Y} \mathbf{X}$	M-VQR	GT
\mathcal{S}^2	Cond. Multimodal	14.2 ± 2.09	8.21 ± 2.69
	Scaled Star	2.25 ± 0.99	2.95 ± 1.25
	Scaled Heart	4.46 ± 1.18	3.59 ± 1.21
	Continental Drift	17.2 ± 2.6	9.45 ± 1.06
\mathcal{T}^2	Cond. Multimodal	17.5 ± 2.92	11.4 ± 4.87
	Scaled Star	5.85 ± 1.92	3.44 ± 2.07
	Scaled Heart	7.93 ± 2.14	4.17 ± 2.1
	Dihedral angles	26.1 ± 5.6	7.34 ± 2.88

probability τ_κ contained in \mathcal{C}_κ^* as the percentage of points $\mathbf{u} \in \mathcal{M}$ with $d_{\mathcal{M}}(\boldsymbol{\omega}, \mathbf{u}) \leq \kappa$. The function \mathcal{C}_ω^* is then the interpolation of the pairs $\{(\kappa, \tau_\kappa)\}_{\kappa=1}^{N_\kappa}$. Figure 7 shows the amount of probability contained in the contours as their distance from the pole increases with and without using \mathcal{C}_ω^* .

D.2.1 τ -contours

The τ -contours are computed using the vector quantile function $\hat{Q}_{\mathbf{Y}}$. First of all, we compute the τ -contours $\mathcal{C}_\tau^{\mathbf{U}}$ on the base distribution μ (Figure 9). Then we map $\mathcal{C}_\tau^{\mathbf{U}}$ to the target distribution using the learned vector quantile function $\hat{Q}_{\mathbf{Y}}$: $\mathcal{C}_\tau^{\mathbf{Y}} := \hat{Q}_{\mathbf{Y}}(\mathcal{C}_\tau^{\mathbf{U}})$.

E ADDITIONAL EXPERIMENTS RESULTS

In this section, we present additional experimental results and accompanying plots that extend upon the findings reported in the main paper.

KDE- L_1 scores. Table 2 provides a detailed breakdown of the KDE- L_1 scores for all the conditional distributions featured in the primary paper. Specifically, we furnish both the mean and standard deviation of the KDE- L_1 values, calculated for two distinct scenarios: (1) between ground truth samples and samples generated by our method (M-VQR), and (2) solely among samples drawn from the ground truth distribution (GT). The latter value aids us in determining the extent to which any discrepancies are attributed to finite sampling, thus allowing us to measure the effectiveness of our method in accurately capturing the ground truth distribution through sampling.

Sample complexity. For M-VQE, we experimented with ablating the number of samples, and observed that the estimation accuracy (ESS, KDE- L_1) saturates at $N = 1k$. In contrast, M-VQR has much higher sample complexity. This is intuitive because modeling conditional distributions is significantly harder than estimating a single distribution (as in M-VQE). To demonstrate this, we performed an ablation study of M-VQR on conditional multi-modal distributions with increasing number of samples, as presented in Figure 10. The results suggest that

the Coverage error and ESS saturate at $N = 12.5k$, while KDE-L1 error drops as N increases.




Figure 10: **Performance variation as the number of training samples increases.** We report the Coverage error, KDE- L_1 and ESS(%) for the M-VQR trained on the Conditional Multimodal distribution on a sphere.

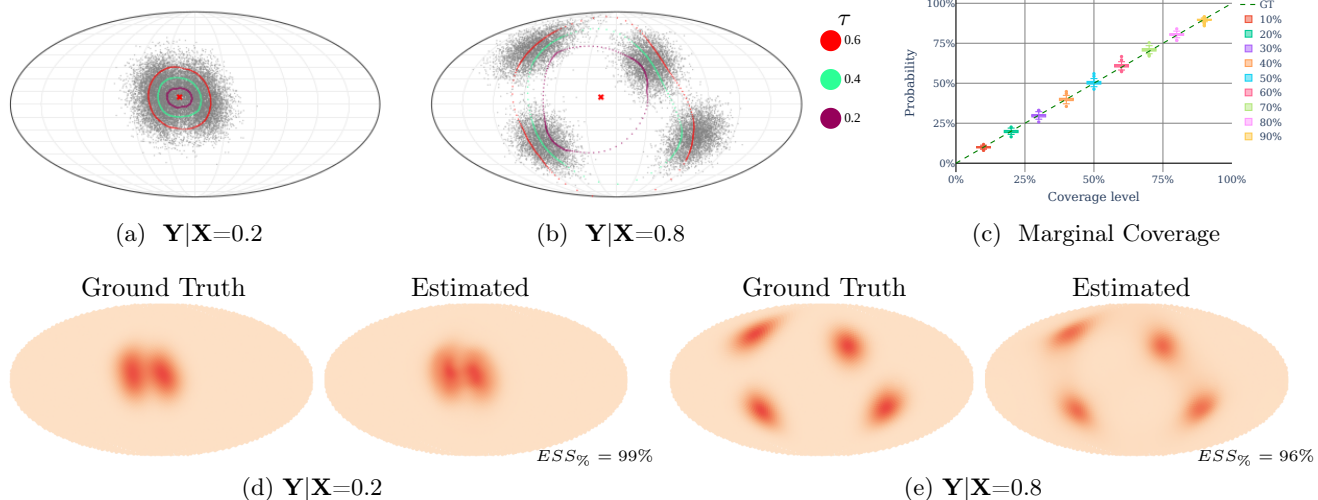


Figure 11: τ -confidence sets and Likelihood computed with M-VQR on the ‘Translated Multimodal’ dataset. Subfigures (a) and (b) report τ -contours overlayed on the ground truth samples, for different values of τ . Mollweide projection is used to visualize the whole sphere. Graph (c) shows the coverage achieved by the model as a function of the requested coverage level, averaged over the different conditionings with relative confidence bars. Subfigures (d) and (e) show the likelihood and the ESS%.

Translated Multimodal. In the main paper, we considered synthetic datasets where conditioning governed the scale of the distribution. Here, we consider a more intricate scenario: a mixture of four von-Mises distributions defined on the manifold \mathcal{M} , where the conditioning controls the positions of the modes. Figure 11 showcases the results on the S^2 sphere. We train an M-VQR model employing a c -concave β potential with 4 layers. The c -concave function is further enriched with inner sizes of $[2, 4, 8]$. Remarkably, our method consistently generates contours with all the desired characteristics: they are nested, smooth, and constitute valid contours, with a mean coverage error of $1.01\% \pm 0.97\%$. Additionally, we present the likelihood, which achieves a notable ESS% of 93.72% and a mean KDE- L_1 value of $(14.38 \pm 5.03) \times 10^{-4}$. The ground truth KDE- L_1 stands at $(7.759 \pm 1.52) \times 10^{-4}$. These results further demonstrate our approach accurately captures the conditional distribution.

Transformation of samples vs learned points of the c -concave function. We explore the impact of a transformation denoted as T on the learned c -concave function, focusing on the S^2 sphere and two von-Mises distributions, Y_1 and Y_2 . The second distribution, Y_2 , is derived by applying a 3D rotation to the first, signifying the effect of T . We train two separate models, M-VQE $_1$ and M-VQE $_2$, on these distributions and investigate the consequences of applying the inverse transformation T^{-1} to the second model. Specifically, we apply T^{-1} to both the contours estimated by M-VQE $_2$ ($T^{-1}Y_2$) and the learned support points, denoted as z_2 , from M-VQE $_2$ ($T^{-1}z_2$). The final two plots in Figure 12 reveal that the results are remarkably close to the contours estimated by M-VQE $_1$. To quantify this closeness, we compare the coverage error of M-VQE $_1$ with $T^{-1}Y_2$ and $T^{-1}z_2$, resulting in mean differences of $0.60\% \pm 0.43\%$ and $0.93\% \pm 0.65\%$, respectively. Furthermore, we compute the KDE- L_1 error for all four cases and observe that the error remains consistent across these scenarios. This analysis




Figure 12: **Impact of Transformation T on Estimated Contours.** In this visualization, the first two plots depict the contours derived from two distinct M-VQE models, each trained on different distributions. The second distribution is obtained by applying a transformation T to the first, i.e., $\mathbf{Y}_2 = T\mathbf{Y}_1$. The third plot showcases the contours produced using the second model and subsequently applying the inverse transformation to these contours, yielding $T^{-1}\mathbf{Y}_2$. The fourth plot illustrates the contours generated by applying the inverse transformation to the support points z_2 learned by the second model, resulting in $T^{-1}z_2$. Under each plot, we report the KDE- L_1 ($\times 10^{-3}$) error.

highlights the learned points where the c -concave function is discretized roughly corresponds to the location of the density that is being modeled.

M-VQR: synthetic data experiments. Figure 13 showcases the likelihood alongside the $\text{ESS}_{\%}$ (Effective Sample Size as a percentage) for two distinct conditioning examples within each synthetic distribution featured in the main paper. Notably, in all cases presented, the $\text{ESS}_{\%}$ exceeds 84%, demonstrating that our method efficiently computes the likelihood of the target conditional distribution. This effectiveness is further illustrated in the accompanying plots.

M-VQR: Real data experiments. Figure 14 and 15 present the results of likelihood and sampling calculations by our method, M-VQR, applied to the real data distributions examined in the primary paper. It is evident that, in all cases, M-VQR’s sampled data closely aligns with the ground truth, with a few outliers that may be attributed to discrete formulation approximations. Moreover, we illustrate the likelihoods $p_{\mathbf{Y}|\mathbf{X}}$ computed using M-VQR. In this context, we lack a ground truth likelihood for direct comparison, as these distributions are derived from finite sets of samples. However, by comparing the computed likelihoods with the ground truth sampling, we observe that higher values of $p_{\mathbf{Y}|\mathbf{X}}$ tend to correspond to regions with a denser concentration of samples.

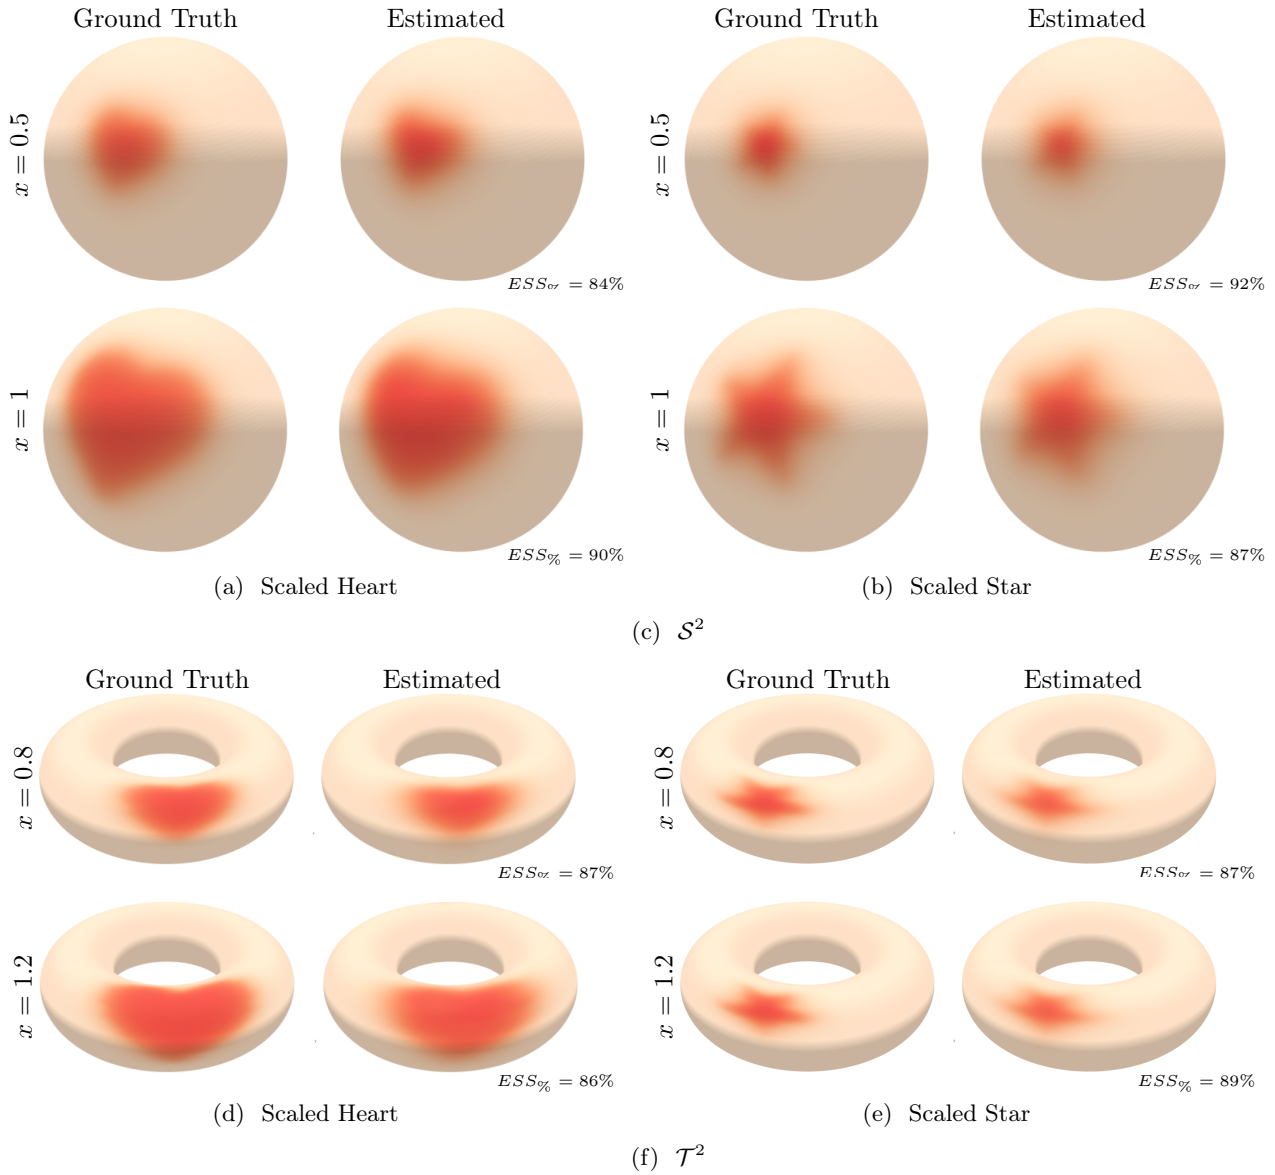


Figure 13: **Likelihood function $p_{Y|X}$ for the Synthetic distributions.** The covariate X controls the distribution scale. ESS% values are also reported.

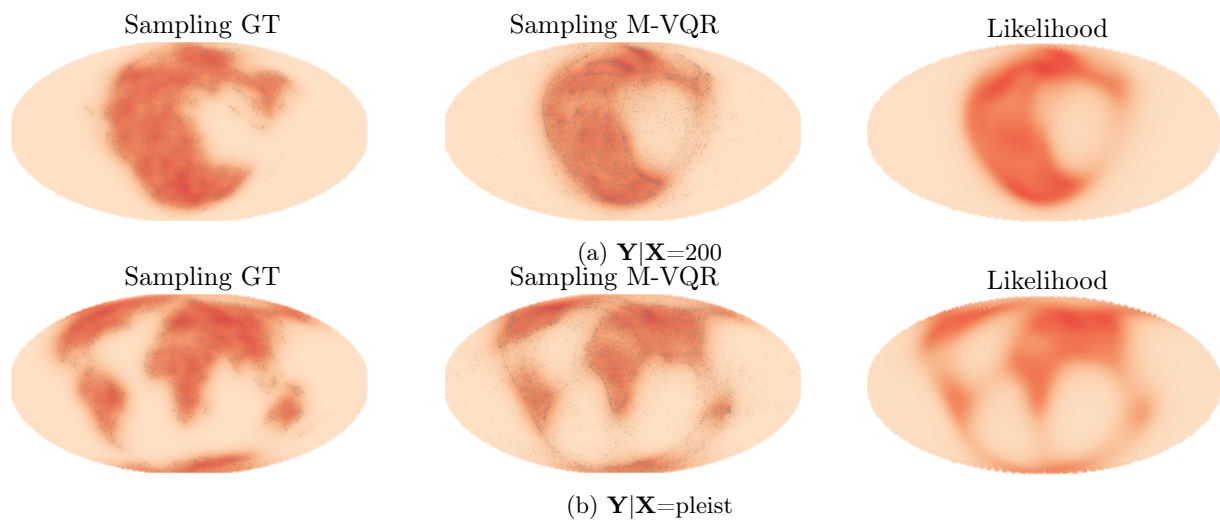


Figure 14: Sampling and Likelihood $p_{Y|X}$ from M-VQR on the 'Continental Drift' dataset.

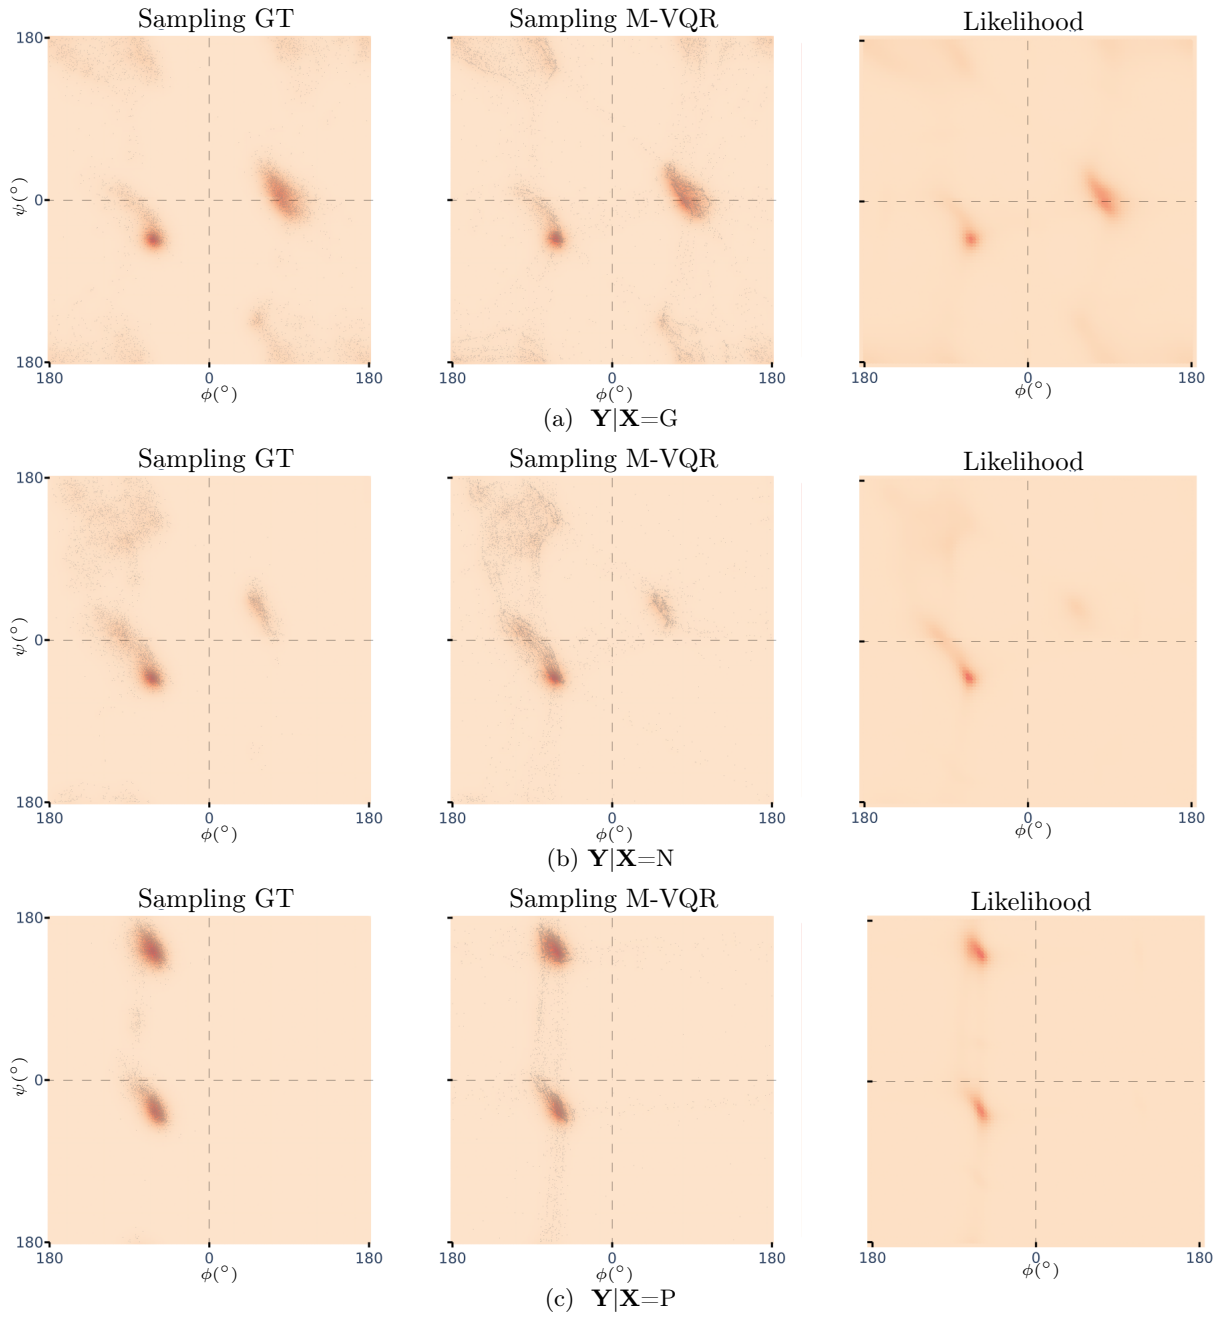


Figure 15: Sampling and Likelihood $p_{\mathbf{Y}|\mathbf{X}}$ from M-VQR on the ‘Dihedral Angles’ dataset.