

Calibrated Multiple-Output Quantile Regression with Representation Learning

Shai Feldman

SHAI.FELDMAN@CS.TECHNION.AC.IL

*Department of Computer Science
Technion—Israel Institute of Technology
Technion City, Haifa 32000, Israel*

Stephen Bates

STEPHENBATES@CS.BERKELEY.EDU

*Departments of Electrical Engineering and Computer Science and of Statistics
University of California, Berkeley
Berkeley, CA 94720, USA*

Yaniv Romano

YROMANO@TECHNION.AC.IL

*Departments of Electrical and Computer Engineering and of Computer Science
Technion—Israel Institute of Technology
Technion City, Haifa 32000, Israel*

Editor: Samuel Kaski

Abstract

We develop a method to generate predictive regions that cover a multivariate response variable with a user-specified probability. Our work is composed of two components. First, we use a deep generative model to learn a representation of the response that has a unimodal distribution. Existing multiple-output quantile regression approaches are effective in such cases, so we apply them on the learned representation, and then transform the solution to the original space of the response. This process results in a flexible and informative region that can have an arbitrary shape, a property that existing methods lack. Second, we propose an extension of conformal prediction to the multivariate response setting that modifies any method to return sets with a pre-specified coverage level. The desired coverage is theoretically guaranteed in the finite-sample case for any distribution. Experiments conducted on both real and synthetic data show that our method constructs regions that are significantly smaller compared to existing techniques.

Keywords: conformal prediction, uncertainty quantification, quantile regression, multiple regression, variational auto-encoder

1. Introduction

In real-world applications, it is often required to estimate more than one response variable. Consider, for example, estimating the effects and side effects of a drug given the patient's demographic information and medical measurements. These two responses may be correlated in the way that when the drug is effective the side effects are more severe (Schuell et al., 2005), and this relation might not be linear. In such high-stakes settings, giving point predictions for the drug's effects and side effects is insufficient; the decision-maker must know the plausible effects for an individual. The plausible effects can be represented as a region in the multidimensional space that covers a pre-specified proportion (e.g., 90%) of the drug's

possible outcomes. In the one-dimensional case, the region reduces to an interval, determined by lower and upper bounds for the response variable. The problem of constructing such a prediction interval is extensively investigated in the literature (Koenker and Bassett, 1978; Izbicki et al., 2020; Meinshausen, 2006; Guan, 2019; Gupta et al., 2021). This approach can be naïvely extended to the multivariate case by estimating a prediction interval for each response separately. However, this process will result in a rectangle-shaped region, whereas the shape of the true distribution of the response variables can be arbitrary, not a rectangle, and even not convex. In that case, the predicted region is likely to be over-conservative: it would not reflect the true underlying uncertainty. A better approach is to predict both outcome variables (drug effect and side effect) jointly. This strategy encourages the model to exclude unlikely combinations of the two from the predicted region. In this work, we show how to construct a region that reflects the true distribution of the response variables while attaining the pre-specified coverage level.

1.1 Problem Formulation

This paper studies the problem of constructing reliable uncertainty estimates in multivariate regression problems. Suppose we are given n training samples $\{(X_i, Y_i)\}_{i=1}^n$, where $X \in \mathbb{R}^p$ is a feature vector, and $Y \in \mathbb{R}^d$ is a response vector. Given a new test point X_{n+1} , our goal is to construct a region of values in which the unknown test response Y_{n+1} falls with high probability. Formally, we seek to build a *marginal distribution-free quantile region* $\hat{R}(X_{n+1}) \subseteq \mathbb{R}^d$ that is likely to contain the response Y_{n+1} with a user-specified coverage probability $1 - \alpha$:

$$\mathbb{P}[Y_{n+1} \in \hat{R}(X_{n+1})] \geq 1 - \alpha, \tag{1}$$

for any joint distribution P_{XY} and any sample size n . This property is called *marginal coverage*, and in order to guarantee it, we assume that all samples $\{(X_i, Y_i)\}_{i=1}^{n+1}$ are drawn exchangeably. That is, we assume that the training samples and test samples follow the same distribution. In addition, we aim to construct quantile regions that are as small as possible, reliably estimating the conditional distribution of $Y \mid X$. When $d = 1$, the quantile region reduces to a one-dimensional prediction interval, determined by lower and upper bounds, within which the response is expected to lie with probability at least $1 - \alpha$.

One of the methods that addresses this problem is *directional quantile regression* (DQR) (Kong and Mizera, 2012; Paindaveine and Šiman, 2011; Boček and Šiman, 2017). The main idea is to estimate conditional quantiles in different directions, where each defines a half-space, and the quantile region is defined as the intersection of all half-spaces. This method is simple and fast, compared to competitive methods that require approximating the entire distribution P_{XY} (Carlier et al., 2016, 2017, 2020). However, being an intersection of half-spaces, the quantile region is convex, so it might be unnecessarily large, as demonstrated in Section 2.1. Additionally, the empirical coverage of such a quantile region is lower than the nominal one, which forces the user to estimate extremal conditional quantiles, as explained in Section 4.1. Furthermore, since the conditional distribution of $Y \mid X$ is unknown, it is difficult estimating what empirical coverage DQR will achieve given a certain nominal level. This raises the problem of choosing the correct nominal level for which DQR achieves the desired coverage rate. In sum, the DQR method can work effectively only in specific cases, i.e., when the distribution of $Y \mid X$ has level sets of the density that are convex. However, even in

those cases, the user is required to estimate extremal quantiles, a process that is known to be impractical, as shown by Diebolt et al. (2000). Furthermore, as demonstrated in (Liu et al., 2019, Section 4.1), DQR is not guaranteed to capture the zone with the highest density even for convex density level sets.

In this work we develop a novel scheme to construct statistically efficient quantile regions, relying on the DQR approach. The core idea is to learn a representation of the response variable for which the DQR method is effective. Once obtaining a quantile region for that representation, we transform it to the original representation of the response variable Y . In this work, we use a *conditional variational auto-encoder* (CVAE) (Sohn et al., 2015) model to learn a mapping between these two representations, described in detail in Appendix F.4; see also (Thickstun et al., 2017; Xu and Tewari, 2021; Bengio et al., 2013; Oord et al., 2017; Kobyzev et al., 2020) for other representation learning techniques. In striking contrast to the DQR, this scheme is non-parametric and can produce non-convex regions, which are therefore smaller and more informative. In Section 2.1, and in the experiments, provided in Section 5, we show that among the four methods we examine in this work, our method consistently tracks the conditional distribution better, tested on six real data sets, and various synthetic examples. This phenomenon is supported by a theoretical guarantee of the quantile region covering only possible responses, i.e., responses in the support of $Y | X$.

Secondly, we extend ideas from conformal prediction to the multidimensional case and propose a calibration procedure that guarantees the coverage requirement (1) in the finite-sample case for any distribution. Conformal inference (Vovk et al., 2005) is a framework commonly used in the one-dimensional case ($d = 1$) (Romano et al., 2019; Izbicki et al., 2020; Sesia and Candès, 2020; Kivaranovic et al., 2020; Chernozhukov et al., 2021; Gupta et al., 2021; Izbicki et al., 2021; Guan, 2019) that provides a generic methodology for building prediction intervals that provably attain valid marginal coverage (1). See (Angelopoulos and Bates, 2021) for a recent overview of this subject. The procedure we propose is generic and can be applied to any multiple-output quantile regression method, including those we discuss in this work. In Section 4.1 we show that this procedure is vital to the DQR method, since its empirical coverage is significantly lower than the nominal level.

2. Background and Related Work

In this section, we describe existing and related methods, but first, begin with a small synthetic experiment that demonstrates the challenges of constructing informative conditional quantile regions.

2.1 A Synthetic Example

The example provided hereafter illustrates how the existing methods perform on this synthetic data, revealing their strengths and weaknesses, which will be further discussed in this section. We generate a v-shaped 2-dimensional response whose structure varies with the feature vector $X \in \mathbb{R}$. This data is visualized in Figure 1, presenting the marginal and conditional distribution of the data. Observe that as X increases, the response is shifted downwards and the slope of the valley becomes steeper. For illustration purposes, we choose to work with a one-dimensional feature vector here, however, in section 5 we also examine data sets

with responses and features of higher dimensions. The full description of the distribution of $Y | X$ is given in Appendix B.1.

Figure 2 presents the conditional distribution of two test points: $x = 1.5$ and $x = 2.5$, and the quantile regions constructed for each of them. This figure shows that the regions constructed by our method reflect the true conditional distribution. This stands in striking contrast with the competitive techniques, which are described hereafter.

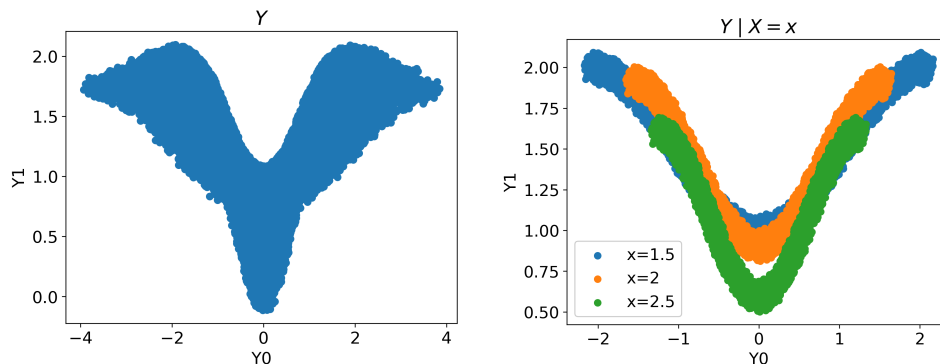


Figure 1: Synthetic data visualization. Left: scatter plot of the marginal distribution of Y . Right: scatter plot of the conditional distribution of $Y | X = x$, for $x \in \{1.5, 2, 2.5\}$.

2.2 One-dimensional Quantile Regression

Conditional quantile regression is a commonly used method to estimate a certain quantile, such as the median, of a Y conditional on X , given a sample $\{(X_i, Y_i)\}_{i=1}^n$ drawn from a distribution P_{XY} . The α -th quantile function of Y is defined as:

$$q_\alpha(x) := \inf\{y \in \mathbb{R} : F(y | X = x) \geq \alpha\},$$

where F is the CDF of $Y | X = x$. One application of conditional quantiles is to obtain a prediction interval for a one-dimensional response, as presented next. Denote by $\alpha_{lo} = \alpha/2$, $\alpha_{hi} = 1 - \alpha/2$ the lower and upper quantile levels, respectively. Given the lower and upper quantiles $q_{\alpha_{lo}}(x), q_{\alpha_{hi}}(x)$, the prediction interval for Y given $X = x$ is defined as:

$$C(x) = [q_{\alpha_{lo}}(x), q_{\alpha_{hi}}(x)].$$

By construction, the interval satisfies the requirement in (1). While the true conditional quantiles are unknown, they can be estimated empirically by solving an optimization problem, e.g., by minimizing the pinball loss (Koenker and Bassett, 1978; Koenker and Hallock, 2001; Steinwart and Christmann, 2011). This process is known to yield estimations that are asymptotically consistent under some regularity conditions (Meinshausen, 2006; Takeuchi et al., 2006; Steinwart and Christmann, 2011). Even though the estimated quantiles are not perfectly accurate, they have been shown to be adaptive to local variability (Hunter and Lange, 2000; Taylor, 2000; Koenker and Hallock, 2001; Meinshausen, 2006; Takeuchi et al., 2006; Steinwart and Christmann, 2011).

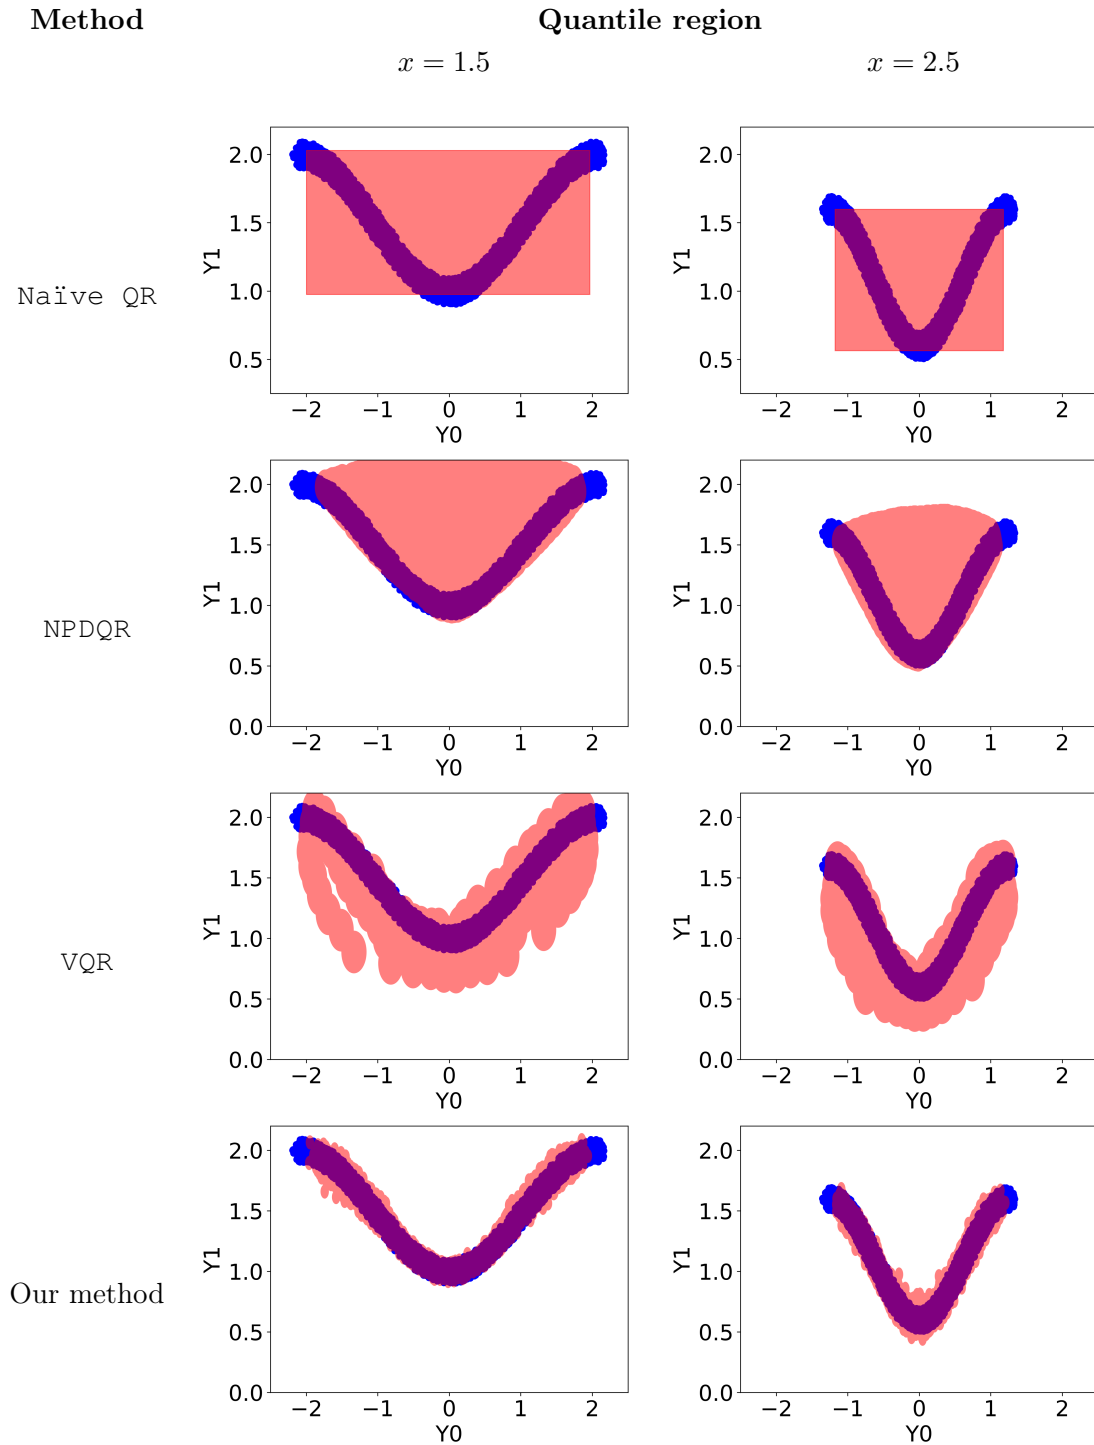


Figure 2: Quantile region obtained by each of the methods: Naïve QR, NPDQR, VQR, and our method. See more details about the synthetic data in Appendix B.1.

2.3 Naïve Multivariate Quantile Regression

The idea presented in the previous section can be extended to build a quantile region for a multivariate response. The naïve approach regresses to the upper and lower quantile for each dimension separately. The nominal coverage level for each dimension is set to be $1 - \beta$, where $\beta = \alpha/d$. This process results in a prediction interval C^j for each feature in the response vector that attains the right coverage rate in the population level:

$$\mathbb{P}[Y_{n+1} \in C^j(x)] = 1 - \beta.$$

The prediction intervals are used to construct the quantile region in the following way:

$$R(x) = C^1(x) \times C^2(x) \times \dots \times C^d(x).$$

Notice that the resulted quantile region is a rectangle, for any distribution of $Y \mid X$. Furthermore, in the ideal infinite-samples case, the produced quantile region satisfies the coverage requirement (1), as proved in Appendix A.1. Even though this method converges to the desired coverage level with infinite data, the quantile regions it produces are not flexible, and too conservative. This problem is illustrated in Figure 2, where we see that the regions produced by this naïve approach do not reflect the true distribution of the response, as opposed to the other methods: the true distribution is v-shaped, whereas the estimated regions are rectangles. Moreover, further experiments (Section 5) reveal that this method produces the largest quantile regions among all existing methods, indicating poor statistical efficiency.

2.4 Directional Quantile Regression

The next approach, which we refer to as *directional quantile regression* (DQR) (Kong and Mizera, 2012; Paindaveine and Šiman, 2011; Boček and Šiman, 2017), is not restricted to produce rectangle-shaped quantile regions, in contrast to the naïve approach, and thus can improve the statistical efficiency. However, DQR is also limited as it can only produce *convex* quantile regions, as it coincides with Tukey depth (Tukey, 1975; Kong and Mizera, 2012). Observe that the naïve method from Section 2.3 estimates the boundaries of the quantile region in four directions $u \in \{(0, 1), (1, 0), (0, -1), (-1, 0)\}$. The DQR method extends this procedure, and estimates the boundaries in all directions, as displayed in Figure 3. Formally, DQR first projects $Y \in \mathbb{R}^d$ using a direction $u \in \mathbb{S}^{d-1}$, where $\mathbb{S}^{d-1} := \{u \in \mathbb{R}^d : \|u\|_2 = 1\}$ is the unit sphere of \mathbb{R}^d . The quantile region boundaries are defined as the following hyperplanes. The order- α quantile of Y given $X = x$ in a direction $u \in \mathbb{S}^{d-1}$ is any element of the collection of hyperplanes:

$$\pi_{\alpha u, x} := \{(x, y) \in \mathbb{R}^p \times \mathbb{R}^d : u^T y = f_\theta(x, u)\},$$

with

$$\theta = \operatorname{argmin}_{\theta'} \frac{1}{n} \sum_{i=1}^n \rho_\alpha(u^T Y_i, f_{\theta'}(X_i, u)),$$

where θ are the parameters of the regression model, $f_\theta(x, u) : \mathbb{R}^{p+d} \rightarrow \mathbb{R}$ is the regression function and ρ_α is the pinball loss, expressed as:

$$\rho_\alpha(y, \hat{y}) = \begin{cases} \alpha(y - \hat{y}) & y - \hat{y} > 0, \\ (1 - \alpha)(\hat{y} - y) & \text{otherwise.} \end{cases}$$

Paindaveine and Šiman (2011) defined $f_\theta(x, u)$ as a linear function, i.e., $f_\theta(x, u) = \theta(u)^T x$, where the coefficients $\theta(u) \in \mathbb{R}^p$ are a function of the direction u . A solution for each direction defines the following half-space:

$$H_u^+(x) = \{y \in \mathbb{R}^d : u^T y \geq f_\theta(x, u)\}.$$

Figure 3 illustrates the half-spaces obtained from different directions. As the figure implies, the quantile region is defined as the intersection of all half-spaces obtained from all directions:

$$R(x) = \bigcap_{u \in \mathbb{S}^{d-1}} H_u^+(x). \tag{2}$$

As shown by Boček and Šiman (2017), the conditional quantile regions are closed, convex and nested, i.e., for any $x \in \mathcal{X}$, $R_{\alpha_1}(x) \subseteq R_{\alpha_2}(x)$ for $\alpha_1 \geq \alpha_2$. The convexity of the constructed regions is also illustrated in Figure 2. However, the quantile region achieves coverage lower than the nominal level in the population level. See Section 4.1 for more details. We note that there are more recent methods to compute the same quantile regions described in this section (Hallin et al., 2010, 2015; Charlier et al., 2020). These methods require estimating only a finite set of half-spaces, but they are all limited to construct convex regions. In fact, the method of Hallin et al. (2010) is the common way to compute these directional regions, although in this work we focus on the version proposed by Paindaveine and Šiman (2011), for simplicity.




Figure 3: Half-spaces obtained from DQR on unconditional data. The left and right panels contain 8 and 512 different half-spaces, respectively. Figure credit: Hallin et al. (2013).

2.5 Additional Related Work

The problem of estimating the quantile region for $Y | X = x$ was also tackled by Hallin et al. (2015); Charlier et al. (2020). However, their proposed methods can only construct convex regions and are based on kernel functions or on a quantization grid, which are infeasible for high-dimensional regressors. A similar technique is the one by Liu et al. (2019), which is a fast algorithm to construct half-space regions. A different approach to multivariate quantiles, called *vector quantile regression* (VQR) (Carlier et al., 2016, 2017, 2020), addresses

the estimation of the conditional distribution of $Y \mid X = x$, and can produce non-convex quantile regions. Nevertheless, it assumes that Y depends linearly on X , and Figure 2 shows that this method fails on the synthetic data, in which this assumption is not satisfied. This approach relates to the one proposed by Chernozhukov et al. (2017) that is based on statistical depth. More recent approaches to construct distribution-free quantile regions, based on geometric tools, are proposed by Hallin (2017); Hallin et al. (2021). However, these methods can only handle data without covariates.

2.6 Our Contribution

We state three features of our method, where each addresses a different limitation of the vanilla DQR method.

Flexible quantile regions. The DQR method can only construct convex regions, whereas the true distribution of the response might not be convex. As illustrated in Figure 2, the quantile regions produced by this method are too conservative and uninformative. In contrast, as indicated by this figure, the regions produced by our algorithm are non-convex, reflecting the true distribution of the response.

Feasible for high-dimensional responses. Due to the curse of dimensionality, the quantile regression problem becomes more difficult when increasing the dimension of the response. In addition, the discrepancy between the nominal coverage level and the empirical coverage rate achieved by DQR worsens as the dimension of the response increases. See Section 4.1 for more details. Moreover, the time complexity of the methods proposed by Carlier et al. (2016, 2017, 2020) grows exponentially as the dimension of the response increases. Our method overcomes these limitations by computing the directional quantiles in space \mathcal{Z} , whose dimension can be determined regardless of the dimension of the response. As a result, our method is feasible for higher-dimensional response settings.

Guaranteed coverage rate. The quantile regions constructed by DQR and VQR are not guaranteed to achieve the desired coverage rate. This problem is more severe with the DQR, whose coverage rate is significantly lower than the nominal level, even with infinite data. See Section 4.1 for additional details. To overcome this limitation, we develop a calibration scheme that guarantees the coverage requirement (1). This process is generic and can also be applied to our proposed method, DQR, VQR, and other methods. Our numerical experiments in Section 5 show that, after calibration, all methods achieve the right marginal coverage.

3. Proposed Method

In this section, we introduce the proposed algorithm, but first, extend DQR beyond linear settings.

3.1 Non-parametric Directional Quantile Regression

Before describing our contribution, we pause to extend the formulation of conditional directional quantiles as defined by Kong and Mizera (2012); Paindaveine and Šiman (2011); Boček and Šiman (2017) beyond linear models. This extension of DQR will be used as a subroutine in our main algorithm. To this end, we follow the original DQR from Section 2.4,

however, use a *non-parametric function class* for f_θ , formulated as neural networks in this work. In such a case, the quantile region for Y conditional on $X = x$ is given by

$$R(x) = \{y \in \mathbb{R}^d : u^T y \geq f_\theta(x, u), \forall u \in \mathbb{S}^{d-1}\}.$$

The latter stands in contrast with the methods proposed by Paindaveine and Šiman (2011); Boček and Šiman (2017) that allow $R(x)$ in (2) to depend only linearly on x . We refer to this new method as non-parametric DQR (NPDQR) throughout this work.

3.2 Our Method: Going Beyond Convex Quantile Regions

In this section, we present a general approach to construct quantile regions of an arbitrary shape, overcoming the convexity restriction of NPDQR. Our method relies on the following observation. When the distribution of $Y | X$ has level sets of the density that are convex, NPDQR (which must create convex regions) is still appropriate. This motivates us to transform an arbitrary response into a space where it has level sets of the density that are convex. Then, we will apply NPDQR and construct a convex quantile region in that space. Lastly, we will transform it back to the original space of the response, using the inverse of the mapping. By applying a non-linear mapping, this process will result in a quantile region that is not restricted to have a convex shape, having an arbitrary structure.

We now describe this procedure in detail. We start by learning a mapping that transforms a general distribution $Y | X = x$ into a latent distribution Z_x whose level sets are convex. In this work, we focus on mapping $Y | X$ to a r -dimensional standard normal distribution, which is not only spherical, but also has convex level sets. To learn such a mapping, we fit a *conditional variational auto-encoder* (CVAE) (Sohn et al., 2015) on the training set $\{(X_i, Y_i)\}_{i \in \mathcal{I}_1}$, and obtain the non-linear transformation between space \mathcal{Y} to space \mathcal{Z} . For technical details regarding CVAE, see Appendix F.4. For our purposes, an ideal CVAE ($\mathcal{E}(y; x), \mathcal{D}(z; x)$) should satisfy the following:

$$Z_x = \mathcal{E}(Y; X = x) \sim \mathcal{N}(0, 1)^r, \quad \mathcal{D}(Z_x; X = x) = Y.$$

Since Z_x is spherically distributed, the conditional distribution $Z_x | X_{n+1} = x$ for a new test point X_{n+1} has a convex level sets. Figure 4 illustrates this process. The top panel visualizes the non-linear mapping $Y | X \rightarrow Z_x \sim \mathcal{N}(0, 1)^3$ obtained by the CVAE model. Observe how the distribution $Z_x | X = x$ has approximately a spherical shape. Observe also that the inverse transformation is fairly accurate, so it can map samples from space \mathcal{Z} back to space \mathcal{Y} .

Since the distribution of $Z | X$ is approximately spherical, NPDQR can estimate effectively its quantile region. We therefore fit NPDQR in space \mathcal{Z} . First, we map the response vectors of the training set to space \mathcal{Z} , and obtain the transformed training set $\{(X_i, \mathcal{E}(Y_i; X_i))\}_{i \in \mathcal{I}_1}$. Next, we fit a NPDQR model on the transformed training samples, as described in Section 2.4. This process results in a model that can construct a quantile region $R_{\mathcal{Z}}(x) \subseteq \mathcal{Z}$, for any given feature vector x . Notice that even though the constructed regions are convex, they are appropriate, since the distribution of $Z | X$ is approximately spherical. That is, NPDQR is applied in a space for which it is well-suited. The procedure of fitting the NPDQR model is summarized in the bottom panel of Figure 4, displaying (in red) the quantile region constructed in space \mathcal{Z} during training, for a specific feature vector x .

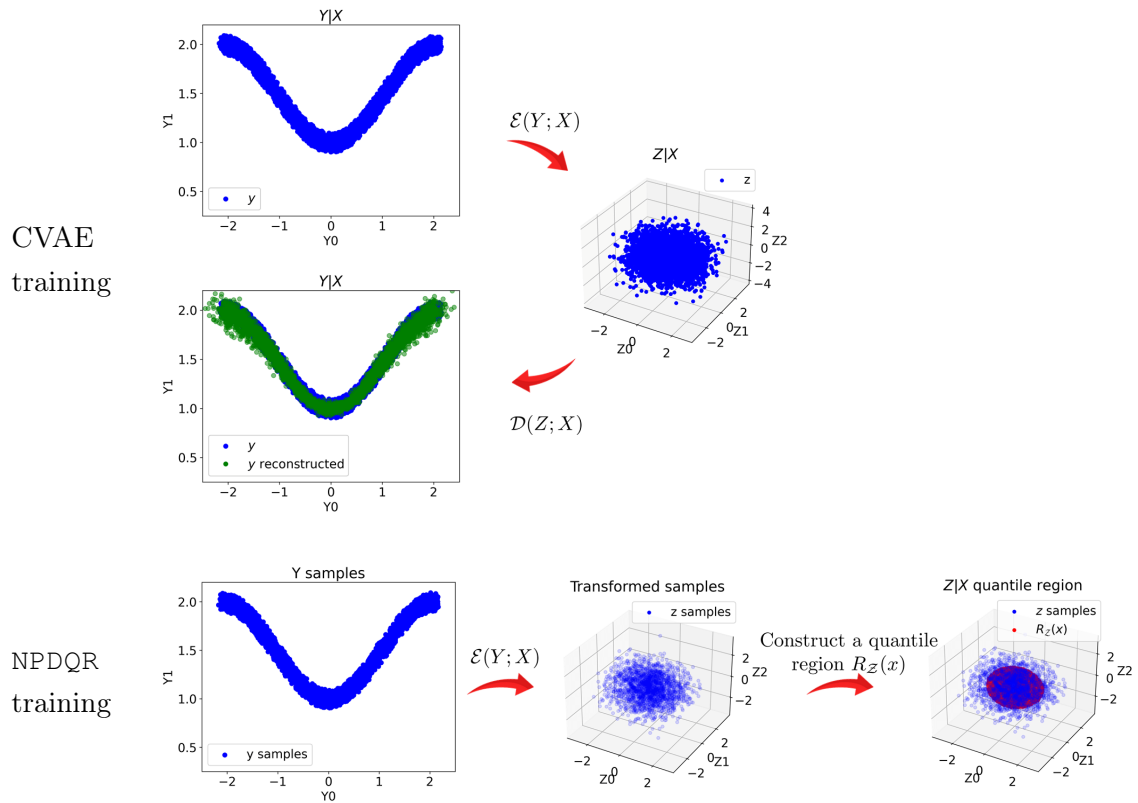


Figure 4: CVAE and NPDQR training schemes on the synthetic data. For further details regarding the synthetic data, see Appendix B.1.

At test time, given a test point X_{n+1} , we (i) construct a quantile region $R_{\mathcal{Z}}(X_{n+1}) \subseteq \mathcal{Z}$ by applying the fitted NPDQR model; and (ii) transform the estimated region to the original space \mathcal{Y} , forming the desired quantile region:

$$R_{\mathcal{Y}}(X_{n+1}) := \mathcal{D}(R_{\mathcal{Z}}(X_{n+1}); X_{n+1}) \subseteq \mathcal{Y}. \quad (3)$$

Observe that $R_{\mathcal{Y}}(X_{n+1})$ is the quantile region of X_{n+1} in \mathcal{Y} . From a practical point of view, the function \mathcal{D} can only map a discrete set of points from $R_{\mathcal{Z}}(X_{n+1})$, and therefore the resulting set $R_{\mathcal{Y}}(X_{n+1})$ is a discretization of the quantile region. We address this important issue in Section 4.2 and show how to construct a continuous region from the discretized $R_{\mathcal{Y}}(X_{n+1})$. The test procedure is illustrated in Figure 5, in which, we can see that while the quantile region in space \mathcal{Z} has a convex shape, the transformed one (in space \mathcal{Y}) has the desired non-convex structure. The whole procedure is summarized in Algorithm 1, which we refer to as *Spherically Transformed* DQR (ST-DQR).

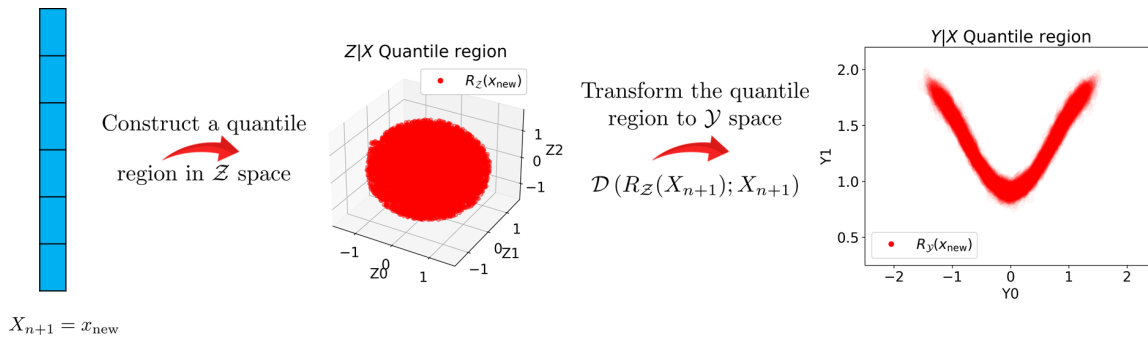


Figure 5: The test procedure on the synthetic data, given a new test point $X_{n+1} = x_{\text{new}}$.

We pause here to highlight several features of the proposed algorithm. First, we use a CVAE model which is nonlinear and non-convex, so we can obtain arbitrary quantile regions, unlike previous approaches. Second, since we apply NPDQR in a latent r -dimensional space, where r is a hyper-parameter, our method can effectively treat high-dimensional response variables by choosing $r < d$. We demonstrate this in Section 5.1, in which we display the results obtained by different methods on synthetic data sets with higher-dimensional responses.

3.3 Theoretical Results

We explain a formal property satisfied by our proposed algorithm that supports the behavior we observed in Figure 2. We would like our quantile region to reflect the true distribution of the response variable. For example, in the synthetic data from Figure 2, the quantile region should only cover blue points, i.e., areas where the response can be present. Formally, we ask that the quantile region will be contained in the support of $Y | X = x$. We now show that a quantile region constructed by our method satisfies this property.

Theorem 1 *Suppose $Y | X = x$ has a continuous distribution for all x . Suppose $(\mathcal{E}(y; x), \mathcal{D}(z; x))$ is a CVAE model that satisfies:*

$$\forall x \in \text{supp}(X) : Z_x = \mathcal{E}(Y; X = x) \in \mathbb{R}^r, \quad \mathcal{D}(Z_x; X = x) \stackrel{d}{=} Y | X = x,$$

Algorithm 1: Spherically transformed DQR (ST-DQR)

Input:

Data $(X_i, Y_i) \in \mathbb{R}^p \times \mathbb{R}^d, i \in \mathcal{I}_1$.
 Miscoverage level $\alpha \in (0, 1)$.
 Directional quantile regression algorithm, e.g., NPDQR from Section 3.1.
 Conditional variational auto-encoder algorithm $(\mathcal{E}(y; x), \mathcal{D}(z; x))$; see Section F.4.
 A test point $X_{n+1} = x$.

Training time:

Fit a CVAE model on the data $\{(X_i, Y_i)\}_{i \in \mathcal{I}_1}$. See (Sohn et al., 2015).
 Transform the response values Y_i to the space \mathcal{Z} : $Z_i = \mathcal{E}(Y_i; X_i), i \in \mathcal{I}_1$.
 Fit a directional quantile regression model on the training set in space \mathcal{Z}
 $\{(X_i, Z_i) : i \in \mathcal{I}_1\}$ to obtain a method to construct quantile regions in \mathcal{Z} , denoted by $R_{\mathcal{Z}}(x)$.

Test time:

Construct the quantile region in \mathcal{Z} $R_{\mathcal{Z}}(X_{n+1} = x)$.
 Transform the quantile region to \mathcal{Y} : $R_{\mathcal{Y}}(X_{n+1} = x) = \mathcal{D}(R_{\mathcal{Z}}(X_{n+1} = x); X_{n+1} = x)$.

Output:

A quantile region $R_{\mathcal{Y}}(X_{n+1} = x)$ for the unseen input $X_{n+1} = x$.

where \mathcal{E} and \mathcal{D} are continuous functions. Suppose $R_{\mathcal{Z}}(x)$ is a quantile region in space \mathcal{Z} . Define the quantile region in space \mathcal{Y} as: $R_{\mathcal{Y}}(x) = \mathcal{D}(R_{\mathcal{Z}}(x); x)$. Then the quantile region $R_{\mathcal{Y}}(x)$ satisfies:

$$R_{\mathcal{Y}}(x) \subseteq \text{supp}(Y | X = x). \tag{4}$$

All proofs are given in Appendix A. Even though the requirement in (4) is a modest bar, we see in Figure 2 that, unlike the proposed method, the other methods do not satisfy this property. In conclusion, we have shown that a quantile region in \mathcal{Y} constructed by our method does not contain spurious portions, since it does not cover areas outside the support of $Y | X = x$. As a complementary result, we also give a lower bound for the coverage rate of a quantile region constructed by our method in Appendix A.4. To achieve the exact nominal coverage level, we propose a calibration procedure, described in Section 4.

4. Calibration

In this section, we introduce a procedure to calibrate quantile regions to exactly achieve $1 - \alpha$ coverage. The procedure is modular and can be used with any quantile region algorithm, such as DQR, VQR, or our proposed method from the previous section. At a technical level, the calibration scheme instantiates split conformal prediction (Vovk et al., 2005) in a way that is compatible with multi-dimensional quantile regions.

4.1 DQR Requires Estimating Extreme Quantiles

To motivate our calibration scheme, we first point out that the parameter α in DQR does not correspond to the coverage level. This phenomenon is known in the literature (Zuo and

Serfling, 2000; Tukey, 1975), and in this section, we provide an intuitive explanation and an example illustrating this problem. As a result, the DQR regions have a coverage level unknown to the user without further calibration, such as the one described in this section. This problem arises from the definition of the DQR quantile region as an intersection of infinite half-spaces, where each covers $1 - \alpha$ of the distribution; see Figure 3. As a result, their intersection, i.e., the quantile region output by DQR, covers strictly less than $1 - \alpha$ of the distribution. To make this precise, we now analyze the coverage rate of a quantile region constructed with the DQR method, in the setting in which $Y | X = x \sim \mathcal{N}(0, 1)^r$ (see Appendix F.5 for the full calculation). The left panel of Figure 6 displays the coverage rate of a quantile region constructed by DQR as a function of the dimension r , when the nominal coverage level is set to 90%. The right panel in that figure presents the coverage of a DQR quantile region as a function of the directional quantile level $1 - \alpha$ for $r = 3$. We see that the achieved coverage is far below the nominal rate. For example, to construct regions that truly have coverage 90% in a three-dimensional response setting, one would need the 99.38% directional quantiles. Unfortunately, such extreme quantiles are impractical to estimate, as shown by Diebolt et al. (2000). In summary, the DQR regions do not achieve the nominal coverage rate, even for reasonable quantile levels. The coverage level is the scaling of interest to the user, so we turn to formulate a calibration scheme that guarantees the desired coverage level.




Figure 6: Coverage rate of a quantile region constructed by DQR. Left panel: marginal coverage as a function of r , when the desired level is 90%. Right panel: marginal coverage as a function of $1 - \alpha$, for $r = 3$.

4.2 Calibration Preliminaries

Recall that ST-DQR produces a discretization of the quantile region for a given test point X_{n+1} , denoted by $R_{\mathcal{Y}}(X_{n+1}) \subseteq \mathcal{Y}$; see (3). We now show how to extend this discrete set to a continuous quantile region that contains infinitely many points. In more detail, we introduce a family of continuous quantile regions, parameterized by a single number. Then, in Section 4.3, we explain how to choose this parameter to achieve the desired coverage level. The method we develop can also generate valid predictive regions for NPDQR (or any quantile method), by discretizing its quantile regions.

We begin by defining a base region $S^\gamma(x)$, which we will later expand or contract. We define that a point $y \in \mathcal{Y}$ is inside the base region of $X_{n+1} = x$ if it is close to a point in $R_{\mathcal{Y}}(x)$. Formally, the base region is given by

$$S^\gamma(x) = \left\{ y \in \mathbb{R}^d : \min_{a \in R_{\mathcal{Y}}(x)} d(a, y) \leq \gamma \right\}, \quad (5)$$

where d denotes L_2 distance, and γ is a distance threshold. We initialize γ to be γ_{init} , which is the 90-th quantile of the distance between two neighbor points in $R_{\mathcal{Y}}(X_{n+1})$; see more details in Appendix F.1. We find that this initialization performs well in the sense that it tends to transform the discrete set into a continuous region, although other options are possible.

Notice that since γ is not tuned, the coverage achieved by this method might be far from the nominal level. To tune this parameter, we first split the data into a training set, indexed by \mathcal{I}_1 , and a calibration set, indexed by \mathcal{I}_2 . Denote the coverage rate of the base regions by

$$c_{\text{init}} = \frac{1}{|\mathcal{I}_2|} |\{Y_i : Y_i \in S^{\gamma_{\text{init}}}(X_i), i \in \mathcal{I}_2\}|,$$

where $|\cdot|$ is the set size. Depending on c_{init} , we grow or shrink the base region $S^{\gamma_{\text{init}}}$ to the extent required to achieve the desired $1 - \alpha$ coverage. We describe these two cases (grow/shrink) separately next. See Appendix F.6 for an explanation of why it is important to handle the two cases separately.

Case 1: Too low coverage. In this setting, $c_{\text{init}} \leq 1 - \alpha$ and therefore we need to enlarge the base region by increasing γ . Figure 7 shows the effect of γ on the quantile region and its coverage rate. By inflating γ , we enlarge the quantile region and, as a result, increase the coverage rate. In Section 4.3 we show how to exploit the calibration set to compute γ_{cal} that rigorously achieves this nominal rate. Given γ_{cal} , the calibrated quantile region in this case is formulated as

$$S^{\gamma_{\text{cal}}}(x) = \left\{ y \in \mathbb{R}^d : \min_{a \in R_{\mathcal{Y}}(x)} d(a, y) \leq \gamma_{\text{cal}} \right\}. \quad (6)$$

In practice, DQR tends to generate regions with coverage rate below the nominal level (recall Figure 6), therefore, this regime, where $c_{\text{init}} \leq 1 - \alpha$, is most likely to happen in practice.




Figure 7: Demonstration of the quantile region under Case 1 (i.e., γ_{init} yields regions of a low coverage rate) for different values of γ .

Case 2: Too high coverage. This scenario treats the case where $c_{\text{init}} > 1 - \alpha$, which is less likely to occur in practice. In this setting, analogously to Case 1, one could decrease γ to reduce the coverage rate. This strategy, however, may result in a new region that is composed of many disjoint sub-regions, as explained in Appendix F.6. This construction is undesired and hard to interpret for continuous distributions, such as the one presented in Figure 2. We therefore alter the scheme in Case 1, and shrink the base region in a different manner. We begin by taking a set of points outside the quantile region, denoted by $R_{\mathcal{Y}}^c(x)$:

$$R_{\mathcal{Y}}^c(x) = \left\{ y \in \mathbb{R}^d : \min_{a \in R_{\mathcal{Y}}(x)} d(a, y) > \gamma_{\text{init}} \right\}.$$

Next, we say that a point y is inside the quantile region if it is far from its boundaries. Formally, the calibrated quantile region is given by

$$S^{\gamma_{\text{cal}}}(x) = \left\{ y \in \mathbb{R}^d : \min_{a \in R_{\mathcal{Y}}^c(x)} d(a, y) \geq \gamma_{\text{cal}} \right\}, \quad (7)$$

where the calibrated threshold parameter, γ_{cal} , is defined hereafter.

4.3 The Calibration Scheme

We now turn to describe how to choose the distance threshold γ in a way that guarantees the coverage requirement (1), by borrowing ideas from conformal prediction. Following the discussion from the previous subsection, we divide the calibration scheme into two cases, depending on the value of c_{init} which is evaluated on the calibration set. For $c_{\text{init}} \leq 1 - \alpha$ (Case 1), we grow the base quantile region by computing $\gamma_{\text{cal}} > \gamma_{\text{init}}$ as follows:

$$E_i^+ = \min_{a \in R_{\mathcal{Y}}(X_i)} d(a, Y_i), \forall i \in \mathcal{I}_2, \quad (8)$$

$$\gamma_{\text{cal}} := [(n_2 + 1)(1 - \alpha)]\text{-th smallest value of } \{E_i^+ : i \in \mathcal{I}_2\},$$

where $n_2 = |\mathcal{I}_2|$. The effect of γ on the coverage rate is visualized in Figure 8. The figure shows that γ_{cal} is the value for which the empirical marginal coverage rate is equal to the desired level (up to a small correction). In Case 2, where $c_{\text{init}} > 1 - \alpha$, we instead grow the




Figure 8: Quantile region coverage rate under Case 1 for different values of γ . The value for which the 90% marginal coverage rate is attained is $\gamma = \gamma_{\text{cal}}$.

Algorithm 2: Calibrating Multivariate Quantile Regression

Input:

Data $(X_i, Y_i) \in \mathbb{R}^p \times \mathbb{R}^d, 1 \leq i \leq n$.

Miscoverage level $\alpha \in (0, 1)$.

Multivariate quantile regression algorithm, e.g., NPDQR from Section 3.1.

An unseen input $X_{n+1} = x$.

Training time:

Randomly split $\{1, \dots, n\}$ into two disjoint sets $\mathcal{I}_1, \mathcal{I}_2$ of sizes n_1 and $n_2 = n - n_1$, respectively.

Fit the multivariate quantile regression algorithm on the training set $\{(X_i, Y_i) : i \in \mathcal{I}_1\}$.

Compute the coverage rate of the uncalibrated quantile regions:

$$c_{\text{init}} \leftarrow \frac{1}{n_2} |\{Y_i : Y_i \in S^{\gamma_{\text{init}}}(X_i), i \in \mathcal{I}_2\}|.$$

if $c_{\text{init}} \leq 1 - \alpha$ **then**

 Compute E_i^+ for each $i \in \mathcal{I}_2$, according to Equation (8).

 Compute γ_{cal} the $[(n_2 + 1)(1 - \alpha)]$ -th smallest value of $\{E_i^+\}_{i \in \mathcal{I}_2}$.

else

 Compute E_i^- for each $i \in \mathcal{I}_2$, according to Equation (9).

 Compute γ_{cal} the $[(n_2 + 1)\alpha]$ -th smallest value of $\{E_i^-\}_{i \in \mathcal{I}_2}$.

Test time:

Obtain a base quantile region $R_{\mathcal{Y}}(X_{n+1} = x)$ using the multivariate quantile regression algorithm. **if** $c_{\text{init}} \leq 1 - \alpha$ **then**

 Construct the calibrated quantile region $S^{\gamma_{\text{cal}}}$ according to Equation (6).

else

 Construct the calibrated quantile region $S^{\gamma_{\text{cal}}}$ according to Equation (7).

Output:

A quantile region $S^{\gamma_{\text{cal}}}(x)$ for the unseen test point $X_{n+1} = x$.

complement of the base quantile region by computing γ_{cal} as follows:

$$E_i^- = \min_{a \in R_{\mathcal{Y}}^c(X_i)} d(a, Y_i), \forall i \in \mathcal{I}_2, \tag{9}$$

$$\gamma_{\text{cal}} := [(n_2 + 1)\alpha]\text{-th smallest value of } \{E_i^-\}_{i \in \mathcal{I}_2}.$$

In words, under Case 1 (Case 2), the quantity E_i^+ (E_i^-) is the distance of Y_i from its closest point *inside* (*outside*) the base quantile region. From a computational perspective, Case 1 is more efficient since usually $|R_{\mathcal{Y}}(x)| < |R_{\mathcal{Y}}^c(x)|$. As a result, computing $E_i^+ = \min_{a \in R_{\mathcal{Y}}(x)} d(a, Y_i)$ requires less operations compared to $E_i^- = \min_{a \in R_{\mathcal{Y}}^c(x)} d(a, Y_i)$. We now state that a quantile region constructed by the above procedure, summarized in Algorithm 2, satisfies the marginal, distribution-free coverage guarantee (1). The proof is given in Section A.3.

Theorem 2 *If $(X_i, Y_i), i = 1, \dots, n+1$ are exchangeable, then the quantile region $S^{\gamma_{\text{cal}}}(X_{n+1})$ constructed by Algorithm 2 satisfies:*

$$\mathbb{P}(Y_{n+1} \in S^{\gamma_{\text{cal}}}(X_{n+1})) \geq 1 - \alpha.$$

Moreover, if the distances E_i^+, E_i^- are almost surely distinct, then the quantile region is almost perfectly calibrated:

$$\mathbb{P}(Y_{n+1} \in S^{\gamma_{cal}}(X_{n+1})) \leq 1 - \alpha + \frac{1}{n_2 + 1}.$$

We pause here to explain the significance of Theorem 2. First, the coverage guarantee of the calibration procedure applies for any sample size, and dimension of (X, Y) . In addition, once applying this procedure, we guarantee that the calibrated version of any base multivariate quantile regression method (DQR/VQR/ST-DQR) would attain the desired $1 - \alpha$ coverage. Therefore, the calibrated methods would differ only in their statistical efficiency, i.e., the area of the constructed quantile region.

5. Experiments

Herein, we systematically quantify the effectiveness of our proposed method (ST-DQR) and compare its performance to existing techniques (Naïve QR, NPDQR, and VQR). Turning to the details of our setup, for all methods except for VQR, we apply a deep neural network as a base model for constructing quantile regions with $1 - \alpha = 0.9$ coverage level. The VQR method cannot incorporate neural networks in the same way, so we applied the procedure exactly as proposed by Carrier et al. (2016); see Appendix F.7 for details. We split the data sets (both real and synthetic) into a training set (38.4%), calibration (25.6%), validation set (16%) used for early stopping, and a test set (20%) to evaluate performance. Then, we normalize the feature vectors and response variables to have a zero mean and unit variance. Appendix C gives the details about the network architecture, training strategy, and more information about this experimental protocol. The performance metrics (coverage and area, as described below) are averaged over 20 random splits of the data. For our method, we set the dimension of the latent space to $r = 3$; see F.4.2 for other choices of this hyperparameter. In all experiments, we report only the performance of the calibrated quantile regions, since this puts all methods on the same scale. Specifically, ST-DQR, NPDQR, and VQR are calibrated according to Algorithm 2 and Naïve QR is calibrated as described in Appendix C.1. Software implementing the proposed method and reproducing our experiments can be found at <https://github.com/Shai128/mqr>

We report the following two metrics, evaluated on test data:

- **Coverage:** The percentage of samples that are covered by the quantile region. The coverage of a point is determined as described in Section 4.2.
- **Area:** The area of the generated quantile region. To evaluate this metric, we take a grid in space \mathcal{Y} , and define the area to be the number of cells that fall inside the quantile region. See more details in Appendix F.2.

5.1 Synthetic Data Results

We return to the synthetic v-shaped data from Section 2.1 and extend it to higher-dimensional settings. In Appendix B.1 we describe how we generate such data for increased dimensions of X and Y . Furthermore, we explore two settings: in the first, the relationship between

the response variables and the covariates is **linear**, whereas in the second this relationship follows a **non-linear** model. In both cases the relationship between the elements in the response vector is non-linear, however, the relationship between Y and X can be either linear or non-linear; see Figure 1. We evaluate the four methods described in this paper (ST-DQR, Naïve QR, NPDQR, and VQR) using the synthetic data sets, and examine their robustness to non-linearity, a high regressor dimension, and a high dimensional response vector. We find that the VQR method is feasible only for data sets of small dimensions, so we report the results only for those data sets; see more details in Table 18 in the Appendix that summarizes VQR’s runtime and memory footprint. For NPDQR and ST-DQR, we estimate a (pre-calibration) directional quantiles of a level higher than the nominal 90% rate (see Table 7), due to the under-coverage problem presented in Section 4.1. For Naïve QR and VQR, we set the quantile level to be equal to the target $1 - \alpha = 90\%$ rate.

Table 1 displays the coverage rates and areas of the constructed quantile regions. Observe that all methods attain the nominal coverage level, a consequence of applying our proposed calibration procedure from Algorithm 2. However, the regions constructed by different methods are different in size, as presented in the same table. Here, our proposed method ST-DQR constructs quantile regions that are substantially smaller compared to all other techniques. These results are anticipated, since Naïve QR and NPDQR are restricted to produce convex quantile regions, forcing the two to cover irrelevant areas, whereas our method does not have this limitation. In addition, since the linearity assumption of VQR is not satisfied in the non-linear setting, the quantile regions it produces are unnecessarily large. Finally, following Figure 2, we can see the advantages of our non-parametric method: it produces a quantile region of an arbitrary shape, estimating well the conditional distribution of $Y \mid X$ in contrast to the competitive techniques. Table 1 also reports the performance metrics that correspond to data sets with high-dimensional features. These results indicate that VQR is infeasible when the dimension of the feature vector is not small enough, while other methods (including ours) are robust to high-dimensional regressors; see Table 18 in the Appendix for more details.

In the case where the response is a four-dimensional vector, the differences between our method and Naïve QR/NPDQR become more significant. (Once again, VQR is infeasible in this setting.) Here, Naïve QR and NPDQR produce quantile regions with an area larger by a factor of 180-440 than regions constructed by our method. This limitation of the standard methods to handle a high dimensional response is also visualized in Figure 9. One explanation for the substantial improvement that our method achieves is this: while the standard methods work in a four-dimensional space (the dimension of Y), our method works in a lower-dimensional space (in this case, the dimension of Z is three), so it can achieve a higher coverage rate for the same directional quantile level; recall Figure 6. Therefore, the calibration applied to our method is milder and does not affect much the base quantile region (5). By contrast, the uncalibrated NPDQR has an extremely low coverage rate and therefore requires an aggressive calibration. That is, it must smooth out the original fit, making it more like a round ball and less adaptive to the test point.

Coverage rate						Relative area of quantile regions				
Setting	d	p	ST-DQR	Naïve QR	NPDQR	VQR	ST-DQR	Naïve QR	NPDQR	VQR
linear	2	1	89.943	90.059	90.041	89.755	1	3.88	3.25	1.476
linear	2	10	89.926	89.789	90.131	90.065	1	4.372	4.222	1.264
linear	2	50	89.91	89.99	89.96	-	1	4.573	3.926	-
linear	2	100	89.963	89.993	90.003	-	1	3.922	3.406	-
nonlinear	2	1	90.126	90.078	90.165	90.13	1	3.369	2.934	2.73
nonlinear	3	1	90.165	90.114	90.021	90.156	1	24.611	21.396	8.161
nonlinear	3	10	89.991	89.881	90.051	-	1	35.21	27.897	-
nonlinear	4	1	90.031	90.175	89.955	-	1	72.037	217.172	-
nonlinear	4	10	89.792	89.841	89.956	-	1	183.672	440.817	-

Table 1: Simulated data experiments. The standard errors are given in Appendix E.3.1. See Appendix B.1 for more details about the synthetic data sets.

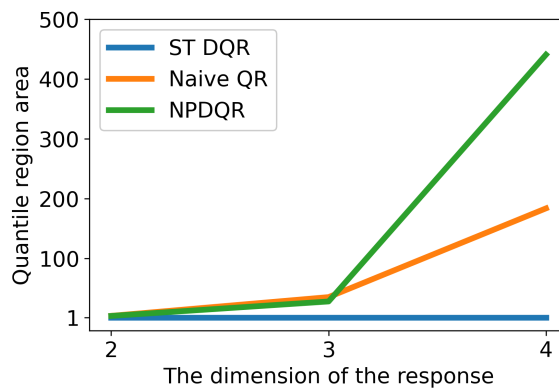


Figure 9: Quantile region area vs. the dimension of the response. The area is scaled by the area of the quantile region constructed by our method, averaged on 20 random splits of the data. The data set used is the non-linear synthetic data with $p = 10$.

5.2 Real Data Results

Next, we compare the performance of the proposed ST-DQR method to NPDQR, and Naïve QR on six benchmarks data sets as in (Romano et al., 2019; Sesia and Candès, 2020): blog feedback (blog_data), physicochemical properties of protein tertiary structure (bio), House Sales in King County, USA (house), and medical expenditure panel survey number 19-21 (meps_19, meps_20, and meps_21). We modify each data set to have a 2-dimensional response as described in Appendix B.2, which also provides additional information about each data set. We follow the experimental protocol and training strategy described in Section 5.1. Specifically, we randomly split each data set into disjoint training (38.4%), calibration (25.6%), validation (16%), and testing sets (20%), and further normalized the feature vector and response variables to have a zero mean and a unit variance each. Due

to the under-coverage problem of DQR presented in Section 4.1, we estimate a directional quantile of a level higher than the nominal 90% rate for NPDQR and ST-DQR; see Table 8. For Naïve QR and VQR, we set the quantile level to be equal to the target $1 - \alpha = 90\%$ rate.

Table 2 summarizes the performance metrics, showing that all calibrated methods consistently attain the nominal coverage rate, as guaranteed by Theorem 2. In addition, the same table indicates that the regions constructed by our method are significantly smaller than the ones produced by the competitive methods. Similar to the synthetic case, VQR is infeasible to deploy and thus omitted from that table. Instead, in Table 3 we report the results on a smaller version of the data sets, in which each feature vector is reduced to dimension 10 using PCA so that VQR is feasible. The table shows that even for these modified versions of the data sets, our method outperforms all the others.

We also test the quality of the constructed regions on sub-populations of the data as follows. We split the test set into three disjoint clusters c_0, c_1, c_2 , where each contains at least 20% of the data. The split is done using the K-means algorithm (Hartigan and Wong, 1979). For illustration purposes, we define the quantile region of a cluster c as

$$S^{\gamma_{\text{cal}}}(c) = \bigcup_{x \in c} S^{\gamma_{\text{cal}}}(x).$$

Figure 10 displays the quantile region constructed by each method for each of the three clusters for the bio data set. The figure shows that the regions constructed by our method reflect the distribution of $Y | X \in c$ better than other methods. Appendix D presents the quantile regions constructed on the house data set, leading to a similar conclusion.

Data Set name	Coverage rate			Relative area of quantile regions		
	ST-DQR	Naïve QR	NPDQR	ST-DQR	Naïve QR	NPDQR
bio	90.0	90.002	89.892	1	1.223	1.222
house	90.157	89.978	89.876	1	1.168	1.149
blog_data	90.145	90.064	90.016	1	1.551	1.821
meps_19	90.144	89.892	89.865	1	2.215	2.169
meps_20	89.997	90.0	89.913	1	2.27	2.209
meps_21	89.899	89.879	89.676	1	2.191	2.212

Table 2: Real data experiments. The standard errors are given in Appendix E.3.2. See Appendix B.2 for more details about the real data sets.

6. Conclusion

In this work, we introduced the ST-DQR method to construct non-parametric and flexible quantile regions of an arbitrary shape. We also proposed a modular extension of conformal prediction to the multivariate response setting that guarantees any pre-specified coverage level. Experiments showed that our method generates informative quantile regions for data with response vectors and features of high dimensions.

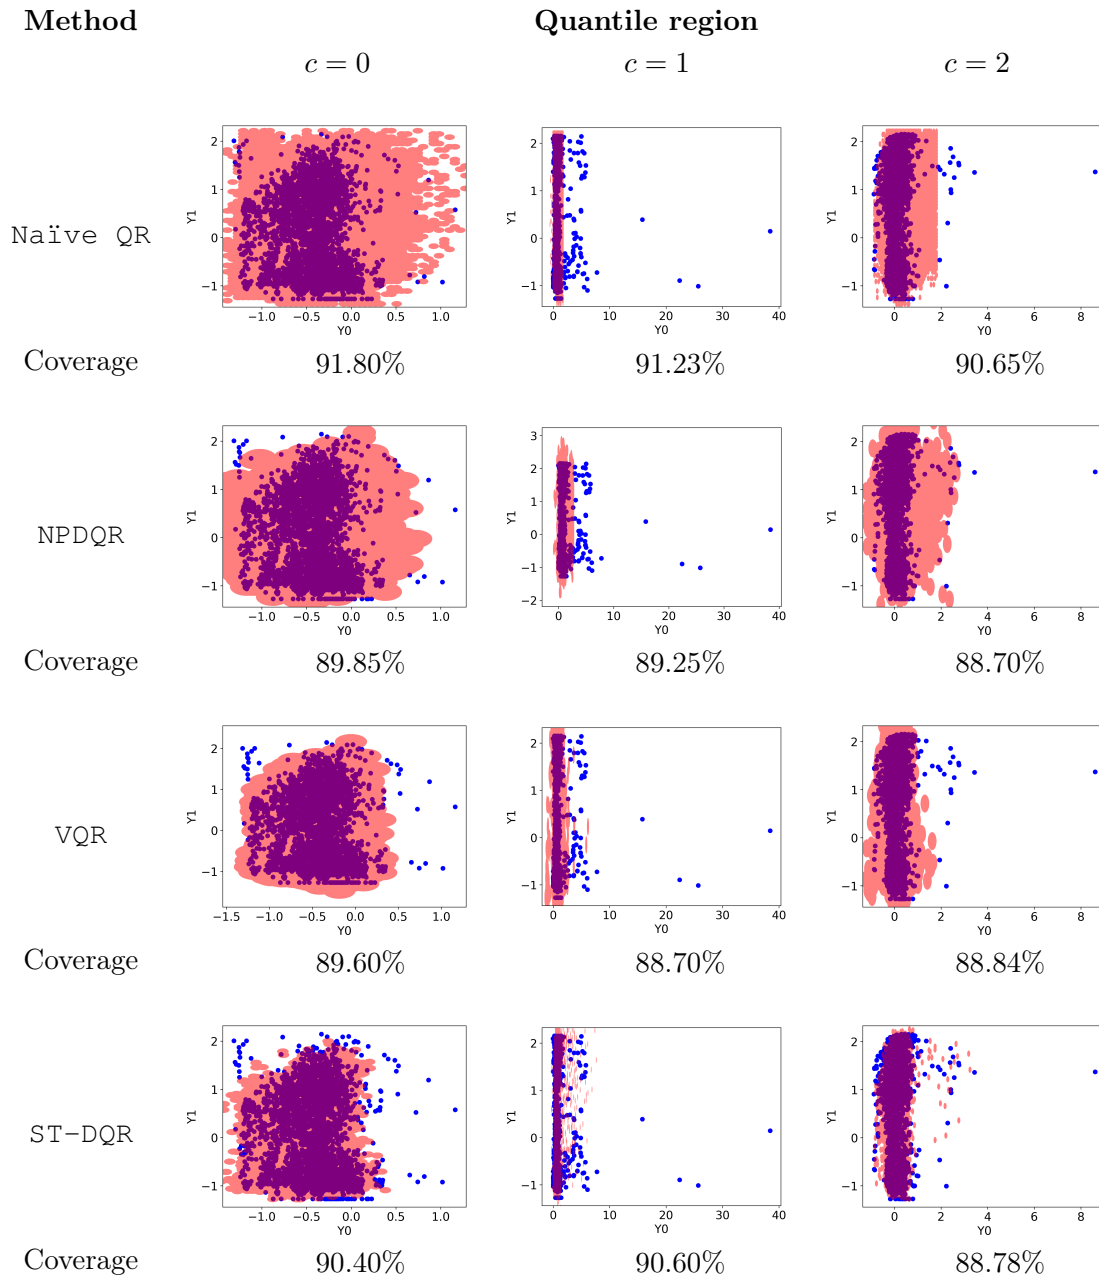


Figure 10: Quantile regions constructed for Bio data set. The regions were obtained by each of the methods: Naïve QR, NPDQR, VQR, and our ST-DQR. In this data set, Y_0 and Y_1 are two protein structural features.

Data Set name	Coverage rate				Relative area of quantile regions			
	ST-DQR	Naïve QR	NPDQR	VQR	ST-DQR	Naïve QR	NPDQR	VQR
bio	90.0	90.002	89.892	89.87	1	1.223	1.222	1.591
house	89.923	90.094	90.067	90.119	1	1.501	1.425	1.359
blog_data	90.078	90.035	89.823	90.034	1	1.744	1.95	1.658
meps_19	90.179	90.16	89.919	90.149	1	2.82	1.527	1.078
meps_20	90.087	89.925	90.036	90.134	1	2.761	1.488	1.1
meps_21	90.061	89.957	89.965	89.887	1	2.848	1.538	1.061

Table 3: Real data experiments. All feature vectors were reduced to dimension 10 using PCA. The standard errors are given in Appendix E.3.2.

A promising future direction could be to exploit the property that the response is approximately normally distributed in the latent space Z_x , and construct a quantile region for normally distributed data instead of using DQR. Another direction could be to replace the CVAE model (used to transform points from space \mathcal{Y} to space \mathcal{Z} and vice versa) with more recent techniques, such as normalizing flows (Kobyzev et al., 2020; Rezende and Mohamed, 2015). Turning the calibration procedure, for very high-dimensional responses, other notions of statistical error beyond marginal coverage—such as the false-negative rate across coordinates—may be more appropriate. Extensions of our procedure to control other error rates would be possible in combination with generalizations of conformal prediction (Bates et al., 2021; Angelopoulos et al., 2021), and we view this as an important next step. Lastly, it would be exciting to explore the conditional coverage of multivariate quantile regression methods, and offer techniques that further improve it, e.g., by generalizing the one proposed by Feldman et al. (2021) to this setting.

Acknowledgments

Y.R. and S.F. were supported by the ISRAEL SCIENCE FOUNDATION (grant No. 729/21). Y.R. thanks the Career Advancement Fellowship, Technion, for providing research support. Y.R. and S.F. thank Alex Bronstein, Sai-Sanketh Vedula, and Aviv Rosenberg for insightful discussions.

Appendix A. Theoretical Results

We now present the proofs of the theoretical results presented in the main manuscript.

A.1 Coverage Guarantee of Naïve Multivariate Quantile Regression

The naïve multivariate quantile regression introduced in Section 2.3 achieves the desired coverage level, as proved next.

$$\begin{aligned}
 \mathbb{P}[Y_{n+1} \in R(x)] &= \mathbb{P} \left[\bigwedge_{j=1}^d Y_{n+1} \in C^j(x) \right] \\
 &= 1 - \mathbb{P} \left[\bigvee_{j=1}^d Y_{n+1} \notin C^j(x) \right] \\
 &\geq 1 - \sum_{j=1}^d \mathbb{P} [Y_{n+1} \notin C^j(x)] \\
 &= 1 - \sum_{j=1}^d \alpha/d \\
 &= 1 - \alpha.
 \end{aligned}$$

A.2 Proof of Theorem 1

Proof [Proof of Theorem 1] We begin by proving that for a fixed $X = x$,

$$\mathcal{D} : \mathbb{R}^r \rightarrow \text{supp}(Y \mid X = x). \quad (10)$$

Assume for the sake of contradiction that Equation (10) does not hold. That is, there exists $y \in \text{Im}(\mathcal{D})$ such that $y \notin \text{supp}(Y \mid X = x)$. Therefore, there exists $\varepsilon > 0$ such that the ball $B := \{a \in \mathbb{R}^d : \|a - y\|_2 \leq \varepsilon\}$ satisfies:

$$B \cap \text{supp}(Y \mid X = x) = \emptyset, \quad \mathbb{P}(\mathcal{D}(Z_x; X = x) \in B \mid X = x) > 0.$$

However, under the assumption, $\mathcal{D}(Z_x; X = x) \stackrel{d}{=} Y \mid X = x$, so it follows that:

$$\mathbb{P}(Y \in B \mid X = x) = \mathbb{P}(\mathcal{D}(Z_x; X = x) \in B \mid X = x) > 0,$$

which contradicts $B \cap \text{supp}(Y \mid X = x) = \emptyset$. We conclude $\text{Im}(\mathcal{D}) \subseteq \text{supp}(Y \mid X = x)$. Finally, since a quantile region in space \mathcal{Y} satisfies $R_{\mathcal{Y}}(x) = \mathcal{D}(R_{\mathcal{Z}}(x); x)$ by construction, we have

$$R_{\mathcal{Y}}(x) \subseteq \text{Im}(\mathcal{D}) \subseteq \text{supp}(Y \mid X = x). \quad \blacksquare$$

A.3 Proof of Theorem 2

Proof [Proof of Theorem 2] We provide the proof for Case 1, in which $c_{\text{init}} \leq 1 - \alpha$. The proof for the complementary case is similar. Recall that the quantile region is defined as:

$$S^{\gamma_{\text{cal}}}(x) = \left\{ y \in \mathbb{R}^d : \min_{a \in R_{\mathcal{Y}}(x)} d(a, y) \leq \gamma_{\text{cal}} \right\},$$

where

$$\gamma_{\text{cal}} = \hat{Q}_{1-\alpha}(\{E_i^+\}_{i \in \mathcal{I}_2}), \quad E_i^+ = \min_{a \in R_{\mathcal{Y}}(X_i)} d(a, Y_i),$$

and $\hat{Q}_{1-\alpha}(\{E_i^+\}_{i \in \mathcal{I}_2})$ is the $\lceil (1-\alpha)(1+|\mathcal{I}_2|) \rceil$ -th smallest value in $\{E_i^+\}_{i \in \mathcal{I}_2}$. This implies that:

$$Y_{n+1} \in S^{\gamma_{\text{cal}}}(X_{n+1}) \iff E_{+n+1}^+ \leq \hat{Q}_{1-\alpha}(\{E_i^+\}_{i \in \mathcal{I}_2}). \quad (11)$$

Since the conformity scores $\{E_i^+\}_{i \in \mathcal{I}_2}$ and E_{+n+1}^+ are exchangeable, the probability of the event in (11) is at least $1-\alpha$. The remaining technical details for proving this statement follow from Romano et al. (2019). The upper bound guarantee of the coverage follows from (11) as well, by applying (Romano et al., 2019, Lemma 2). \blacksquare

A.4 ST-DQR Coverage Rate Lower Bound

While the property guaranteed in Theorem 1 is highly desired, it does not suffice, as it can be trivially satisfied by an empty region $R_{\mathcal{Y}}(x) = \emptyset$, which is a subset of $\text{supp}(Y | X = x)$, hence satisfies the property of Theorem 1. We therefore require the quantile region $R_{\mathcal{Y}}(x)$ to achieve a high coverage rate as well. For complex distributions of $Y | X = x$, the task of constructing a quantile region with a good coverage rate is difficult to achieve via NPDQR. By contrast, the distribution of the response in space \mathcal{Z} is spherical (recall that $Z \sim \mathcal{N}(0, 1)^r$) and thus much simpler to handle in the sense that the region constructed by NPDQR in space \mathcal{Z} is likely to achieve a better coverage rate. Therefore, we would like this coverage property to be preserved when transforming the quantile region back to space \mathcal{Y} . We now show that our method satisfies this property, i.e., the coverage attained in space \mathcal{Y} is at least as good as the one in space \mathcal{Z} .

Proposition 3 *Suppose $(\mathcal{E}(y; x), \mathcal{D}(z; x))$ is a CVAE model as in Theorem 1. Suppose $R_{\mathcal{Z}}(x)$ is a quantile region in space \mathcal{Z} , and $R_{\mathcal{Y}}(x)$ is a quantile region in space \mathcal{Y} , as defined in Theorem 1. Assuming the coverage rate of $R_{\mathcal{Z}}(x)$ is $1-\beta$, then the coverage rate of $R_{\mathcal{Y}}(x)$ is at least $1-\beta$.*

Proof By the assumption of the coverage rate:

$$\mathbb{P}(Z_x \in R_{\mathcal{Z}}(x) | X = x) = 1 - \beta.$$

Since \mathcal{D} is a function, it follows that:

$$Z_x \in R_{\mathcal{Z}}(x) \implies \mathcal{D}(Z_x; X = x) \in \mathcal{D}(R_{\mathcal{Z}}(x); X = x)$$

Therefore:

$$\mathbb{P}(\mathcal{D}(Z_x; x) \in \mathcal{D}(R_{\mathcal{Z}}(x); x) | X = x) \geq \mathbb{P}(Z_x \in R_{\mathcal{Z}}(x) | X = x) = 1 - \beta.$$

Finally, since $Y | X = x \stackrel{d}{=} \mathcal{D}(Z_x; X = x)$, and, $R_{\mathcal{Y}}(x) = \mathcal{D}(R_{\mathcal{Z}}(x); x)$, we conclude that:

$$\mathbb{P}(Y \in R_{\mathcal{Y}}(x) | X = x) = \mathbb{P}(\mathcal{D}(Z_x; x) \in \mathcal{D}(R_{\mathcal{Z}}(x); x) | X = x) \geq 1 - \beta. \quad \blacksquare$$

Appendix B. Data Sets Details

This section provides details about the generation of the synthetic data, and about the real data sets used in the experiments.

B.1 Synthetic Data Details

The generation of the feature vector and the response variable of the *linear* version of the synthetic data is done in the following way:

$$\begin{aligned}\hat{\beta} &\sim \text{Uniform}(0, 1)^p, \\ \beta &= \frac{\hat{\beta}}{\|\hat{\beta}\|_1}, \\ Z &\sim \text{Uniform}(-\pi, \pi), \\ \phi &\sim \text{Uniform}(0, 2\pi), \\ R &\sim \text{Uniform}(-0.1, 0.1), \\ X &\sim \text{Uniform}(0.8, 3.2)^p, \\ Y_0 &= \frac{Z}{\beta^T X} + R \cos(\phi), \\ Y_1 &= \frac{1}{2} (-\cos(Z) + 1) + R \sin(\phi),\end{aligned}$$

where $\text{Uniform}(a, b)$ is a uniform distribution on the interval (a, b) . The *non-linear* version of the synthetic data is generated in the same way, except for an additional non-linear dependence between X and Y_1 :

$$Y_1 = \frac{1}{2} (-\cos(Z) + 1) + R \sin(\phi) + \sin\left(\frac{1}{n} \sum_{i=1}^p X_i\right).$$

In the three-dimensional response case, we define the third response variable as:

$$Y_2 = \sin\left(\frac{Z}{\beta^T X}\right),$$

and in the four-dimensional response case, the fourth response Y_3 is defined as:

$$Y_3 = \cos\left(\sin\left(\frac{Z}{\beta^T X}\right)\right) + R \cos(\phi) \sin(\phi)$$

We report the number of samples of each data set in Table 4.

B.2 Real Data Details

The real data sets originally contained a one-dimensional response, so we increase the target dimension by considering one of the features as a response variable instead. The feature concatenated to the response was chosen to be highly correlated to it, and to have a small correlation to the other features, so it will not be easy to predict. Table 5 displays the size of each data set, the feature dimension, the response dimension, and the index of the feature that is used as a response instead of an input variable.

Setting	p	Number of Samples
linear	1	20000
linear	10	20000
linear	50	80000
linear	100	100000
non-linear	1	20000
non-linear	10	20000

Table 4: Synthetic data sets information. The number of samples as a function of the feature dimension.

Data Set Name	Number of Samples	p	d	Additional Response
blog_data	52397	279	2	The time between the blog post publication and base-time
bio	45730	8	2	F7 - Euclidean distance
house	21613	17	2	Latitude of a house
meps_19	15785	138	2	Overall rating of feelings
meps_20	17541	138	2	Overall rating of feelings
meps_21	15656	138	2	Overall rating of feelings

Table 5: Information about the real data sets.

Appendix C. Experimental Setup

The network we used receives as an input a vector of size $p + d$ (where p is the feature dimension, and d is the response dimension). The first p variables in the input vector correspond to the elements of the feature vector, and the last variables correspond to the desired quantile level. We split the data sets (both real and synthetic) into a training set (38.4%), calibration (25.6%), validation set (16%) used for early stopping, and a test set (20%) to evaluate performance. Then, the feature vectors and the responses were preprocessed using z-score normalization. The neural network consists of 3 layers of 64 hidden units, and a leaky ReLU activation function with parameter 0.2. The learning rate used is $1e^{-3}$, the optimizer is Adam (Kingma and Ba, 2015), and the batch size is 256 for all methods. The maximum number of epochs is 10000, but the training is stopped early if the validation loss does not improve for 100 epochs, and in this case, the model with the lowest loss is chosen. The number of distinct directions used in each gradient step is 32, and they are taken from a fixed collection of 2048 directions that were sampled once, before the training process. The number of directions used to determine the quantile region belonging is 256, and they are sampled from the same collection of directions. The results are averaged over all seeds in the range between 0 and 19 (inclusive). We reduce the dimension of the feature vector to 50 in meps data sets, and to 100 in blog data set, using PCA. The code we use is based on the implementation of Chung et al. (2020).

C.1 Naïve Quantile Regression Setup

We trained a quantile regression model using the pinball loss to predict two quantile levels for each dimension d : $1 - \alpha/d$ and α/d . The quantiles are then used to construct a prediction interval, according to the explanation in 2.2. The calibration scheme used extends CQR, developed by Romano et al. (2019), to the multi-dimensional response case, which we describe in detail next. Let $\alpha_{lo} = \alpha/2$ and $\alpha_{hi} = 1 - \alpha/2$. The CQR conformity scores are defined as:

$$E_i = \max\{\max\{\hat{q}_{lo}^j(X_i) - Y_i, Y_i - \hat{q}_{hi}^j(X_i)\} : j \in \{1, \dots, d\}\},$$

where $\hat{q}_{lo}^j, \hat{q}_{hi}^j$ are the estimated lower and upper quantiles of the j -th dimension of $Y | X$, respectively. We define by Q the $(1 - \alpha)(|\mathcal{I}_2| + 1)$ -th empirical quantile of $\{E_i\}_{i \in \mathcal{I}_2}$. The calibrated quantile region is given by:

$$\hat{R}(x) = \bigtimes_{j \in \{1, \dots, d\}} [\hat{q}_{lo}^j(x) - Q, \hat{q}_{hi}^j(x) + Q]$$

C.2 Directional Quantile Regression Setup

As explained in Section 4.1, the empirical coverage rate of an uncalibrated DQR model is significantly lower than the nominal level. We therefore require the model to achieve a higher coverage level, according to Tables 7, 8. The quantile region obtained by DQR was then calibrated to achieve 90% coverage rate, according to Algorithm 2. The level of the estimated directional quantiles is chosen using an independent train/validation/test split, where the value that achieves the highest coverage level is chosen. The examined directional-quantile levels are 90%, 93%, and 95%. For the four-dimensional response data sets, we also examined a directional-quantile level of 98% for NPDQR.

C.3 Implementation Details of Our Method

Table 6 displays the hidden dimension of each layer in the encoder \mathcal{E} and decoder \mathcal{D} of the CVAE. The dimensions were chosen according to the model’s performance over the linear synthetic data set, with feature vectors of different dimensions. The networks include a dropout with parameter 0.1, and a batch-norm layer for blog data set only. The learning rate used to train the CVAE is $1e^{-4}$ for the real data sets and $1e^{-3}$ for the synthetic data sets, and for both the batch size used is 512. The activation function used is the leaky ReLU function with parameter 0.2. The maximum number of epochs is 10000, but the training is stopped early if the validation loss does not improve for 200 epochs, and in this case, the model with the lowest loss is chosen. In Tables 7, 8 we report the nominal levels of the estimated directional quantiles used to construct the uncalibrated quantile regions. The regions are then calibrated according to Algorithm 2.

p	Hidden layers dimension
$p \leq 5$	32, 64, 128, 256, 128, 64, 32
$5 < p \leq 8$	64, 128, 256, 128, 64
$8 < p \leq 10$	64, 128, 256, 512, 256, 128, 64
$10 < p \leq 25$	64, 128, 256, 256, 128, 64
$25 < p$	128, 256, 512, 512, 256, 128

Table 6: Dimension of each hidden layer for the CVAE architecture as a function of p .

Data Set setting	d	p	NPDQR nominal coverage level	ST-DQR nominal coverage level
linear	2	10	95%	95%
linear	2	50	95%	95%
linear	2	100	95%	95%
non-linear	2	1	95%	93%
non-linear	2	10	95%	93%
non-linear	3	1	95%	93%
non-linear	3	10	95%	93%
non-linear	4	1	98%	93%
non-linear	4	10	98%	95%

Table 7: The directional-quantile levels used for a NPDQR model in the synthetic data sets.

C.4 Machine’s Spec

The resources used for the experiments are:

Data Set Name	NPDQR nominal coverage level	ST-DQR nominal coverage level
blog_data	95%	93%
bio	95%	95%
house	95%	95%
meps_19	95%	93%
meps_20	95%	93%
meps_21	95%	93%

Table 8: The directional-quantile levels used for a NPDQR model in the real data sets.

- **CPU:** Intel(R) Xeon(R) E5-2650 v4.
- **GPU:** Nvidia TITAN-X, 1080TI, 2080TI.
- **OS:** Ubuntu 18.04.

Appendix D. Quantile Regions Constructed for House Data Set

Similarly to figures 2 and 10, we display in Figure 11 the constructed quantile regions for the House data set. We split the data into three clusters, as described in 5.2, and color in red the quantile region of each cluster. It is clear from the figure that our method constructs the most informative quantile regions, among all existing methods. The VQR method could not produce quantile regions for this data set since the dimension of the feature vector is too large to handle for the software we used; see Table 18 in the Appendix.

Appendix E. Additional Experiments

In this section we provide additional experiments, analyzing the conditional coverage, the effect of the calibration set, and more, using the methods discussed in this work. The experimental setup is identical to the one described in Section 5, unless explicitly stated otherwise.

E.1 Calibrating on the Training Set

In this section, we display the performance of each method calibrated on the training set instead of the calibration set. That is, here, all methods were calibrated to achieve $1 - \alpha = 90\%$ coverage rate on the training set. A similar experiment examining the effect of calibrating with a training set instead of a calibration set was previously suggested by Barber et al. (2021), showing that naively calibrated intervals do not attain the right coverage level. The results given below indicate that a model calibrated using the training data does not achieve the desired coverage level, emphasizing the necessity of a calibration set.

Table 9 presents the coverage rates and the areas of each method on the real data sets. This table shows that the coverage attained by these methods is far from the nominal level. Furthermore, this table reveals that our method achieves the best marginal coverage.

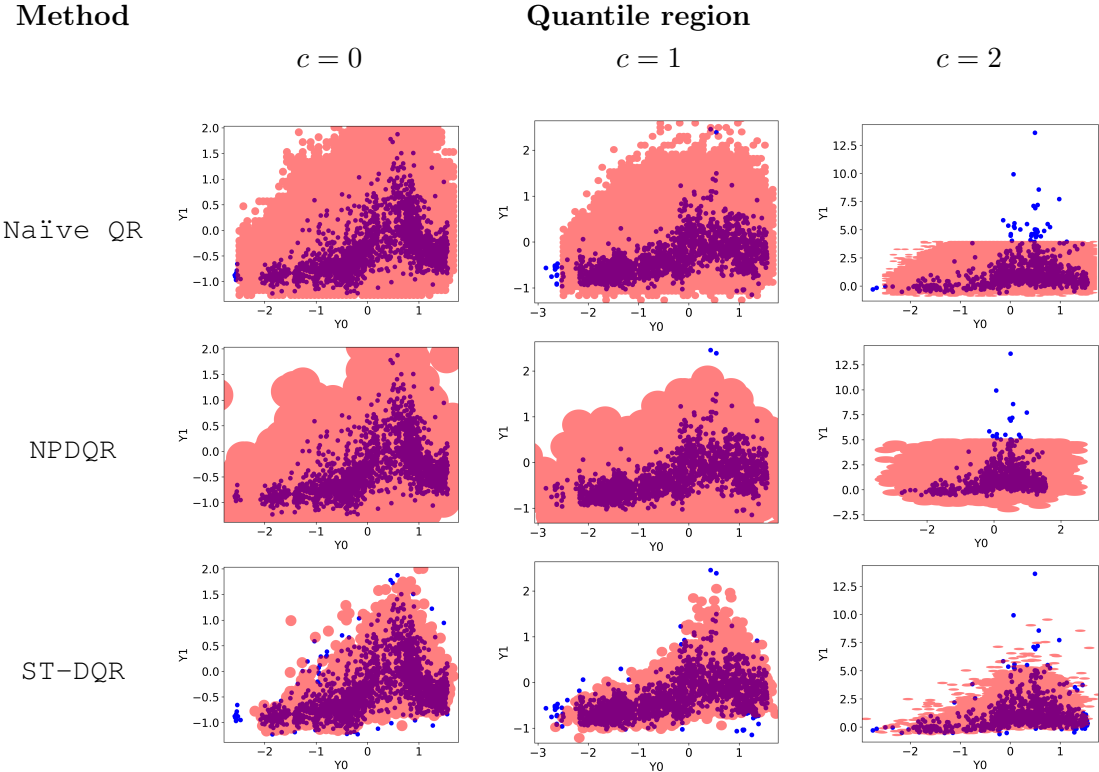


Figure 11: House data set results. Quantile region obtained by each of the methods: Naïve QR, NPDQR, and our ST-DQR. In this data set, Y_0 and Y_1 are the price and latitude of a house.

Data Set name	Coverage rate			Area of quantile regions		
	ST-DQR	Naïve QR	NPDQR	ST-DQR	Naïve QR	NPDQR
bio	88.51 (.101)	86.902 (.147)	85.907 (.173)	306.77 (3.869)	388.469 (4.699)	325.729 (5.141)
house	88.29 (.164)	82.228 (.348)	75.555 (.28)	338.502 (6.657)	347.108 (7.483)	259.874 (4.539)
blog_data	87.675 (.159)	82.802 (.186)	78.9 (.32)	95.812 (7.953)	184.51 (4.369)	146.073 (5.806)
meps_19	88.028 (.171)	83.432 (.231)	82.653 (.871)	160.33 (6.21)	399.095 (22.885)	320.96 (16.72)
meps_20	88.28 (.148)	84.595 (.179)	81.536 (.946)	162.464 (3.133)	419.837 (13.336)	311.72 (17.832)
meps_21	87.664 (.207)	83.456 (.278)	82.996 (.943)	162.164 (4.409)	409.609 (17.192)	344.641 (16.135)

Table 9: Real data experiments. Coverage, area, and their standard error of the quantile regions constructed by each method calibrated on the training set.

E.2 Evaluating Conditional Coverage

In this section, we analyze the conditional validity of the proposed method and compare it to the existing ones.

E.2.1 SYNTHETIC DATA SETS

Figure 12 presents the coverage as a function of the feature vector on the non-linear synthetic data with $d = 2$ response values and $n = 10000$ samples. We generated the feature vector to have the value of x repeated $p = 10$ times. This figure indicates that our ST-DQR achieves the best conditional coverage compared to the competitors. Table 13 shows the coverage as a function of the sample size on the synthetic data described above for different values of x . This figure reveals that as the sample size increases, the conditional coverage of all methods gets closer to the nominal level. Additionally, while ST-DQR attains poor conditional coverage for small samples, it performs better than existing methods for larger samples. Furthermore, it constructs quantile regions that represent better underlying uncertainty compared to competitors, as discussed next. Figures 14 display the quantile regions constructed by ST-DQR on the non-linear synthetic data described above with different sample sizes. This figure shows that even with a small data set, the regions reflect well the conditional uncertainty of the response. This stands in striking contrast to the regions constructed by NPDQR, visualized in Figure 15. In addition, the regions seem to converge to the true conditional density of $Y | X = x$ as the sample size increases.

The experiments conducted in this section indicate that our ST-DQR achieves good conditional coverage and performs better as the sample size increases.

E.2.2 REAL DATA SETS

In this section, we assess conditional coverage violation on the real data sets by measuring the deviation of the cluster coverage from the nominal level. Formally, we define the $\Delta\text{Coverage}$ as:

$$\Delta\text{Coverage} = \frac{1}{|C|} \sum_{c \in C} \left| \frac{1}{|c|} \sum_{x_i \in c} \mathbb{1}\{y_i \in S^{\gamma_{\text{cal}}}(x_i)\} - (1 - \alpha) \right|,$$

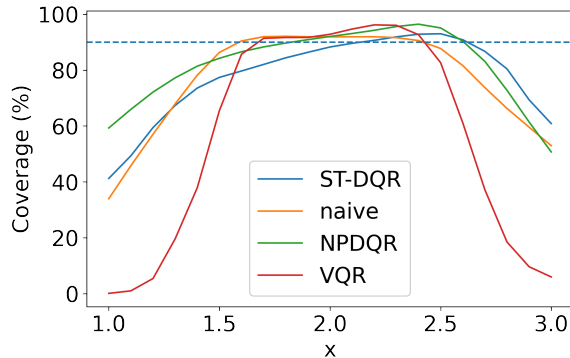


Figure 12: Conditional coverage achieved over the non-linear synthetic data with 10000 samples.

where C is a split of the test set split into clusters, as defined in Section 5.2. Figure 16 displays the $\Delta\text{Coverage}$ achieved by the techniques discussed in this paper on the real data sets introduced in Section 5.2. This figure reveals that our ST-DQR attains the best $\Delta\text{Coverage}$, indicating for a good conditional coverage compared to existing methods.

In Figure 17 we report the $\Delta\text{Coverage}$ on a different version of the real data sets, in which each feature vector is reduced to dimension 10 using PCA so that VQR is feasible. The table shows that even for these modified versions of the data sets, our method achieves better conditional coverage compared to VQR.

E.3 Standard Errors

In this section, we report the standard errors of the metrics reported in the experiments section.

E.3.1 SYNTHETIC DATA

We report the coverage rate along with its standard error, and the area along with its standard error in Tables 10, 11 respectively.

E.3.2 REAL DATA

Table 12 shows the standard error of the coverage and area of the quantile regions constructed for each of the real data sets using the methods discussed in this paper. Additionally, we report the area of quantile regions constructed by our method with different values of r in Table 13. Table 14 displays the standard error of each of the metrics for the reduced version of the real data sets.

Appendix F. Technical Details

In this section, we provide technical details regarding techniques used in this work.




Figure 13: Conditional coverage achieved over the non-linear synthetic data with increasing sample size and for different feature vectors.

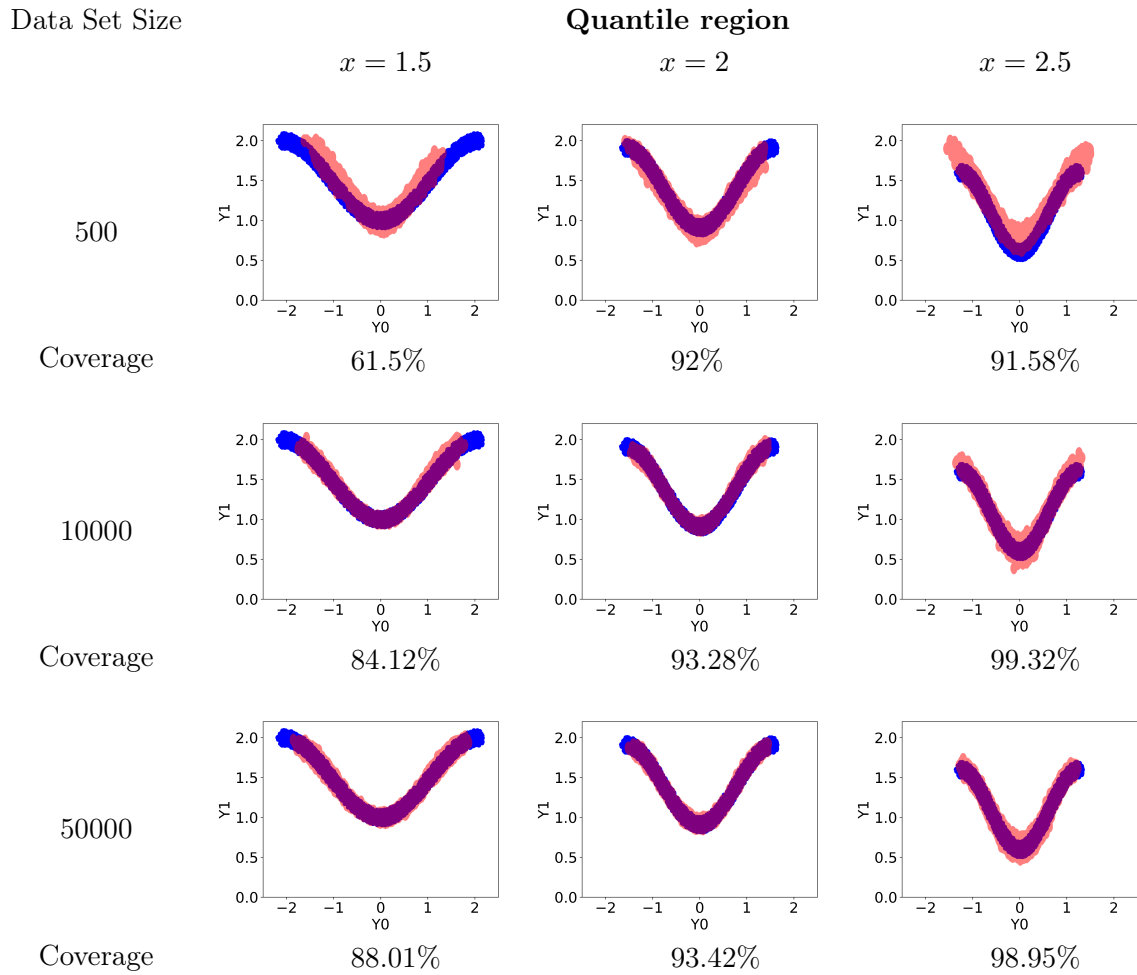


Figure 14: Calibrated quantile regions constructed by ST-DQR for the linear synthetic data set with $p = 10$ for different data set sizes.

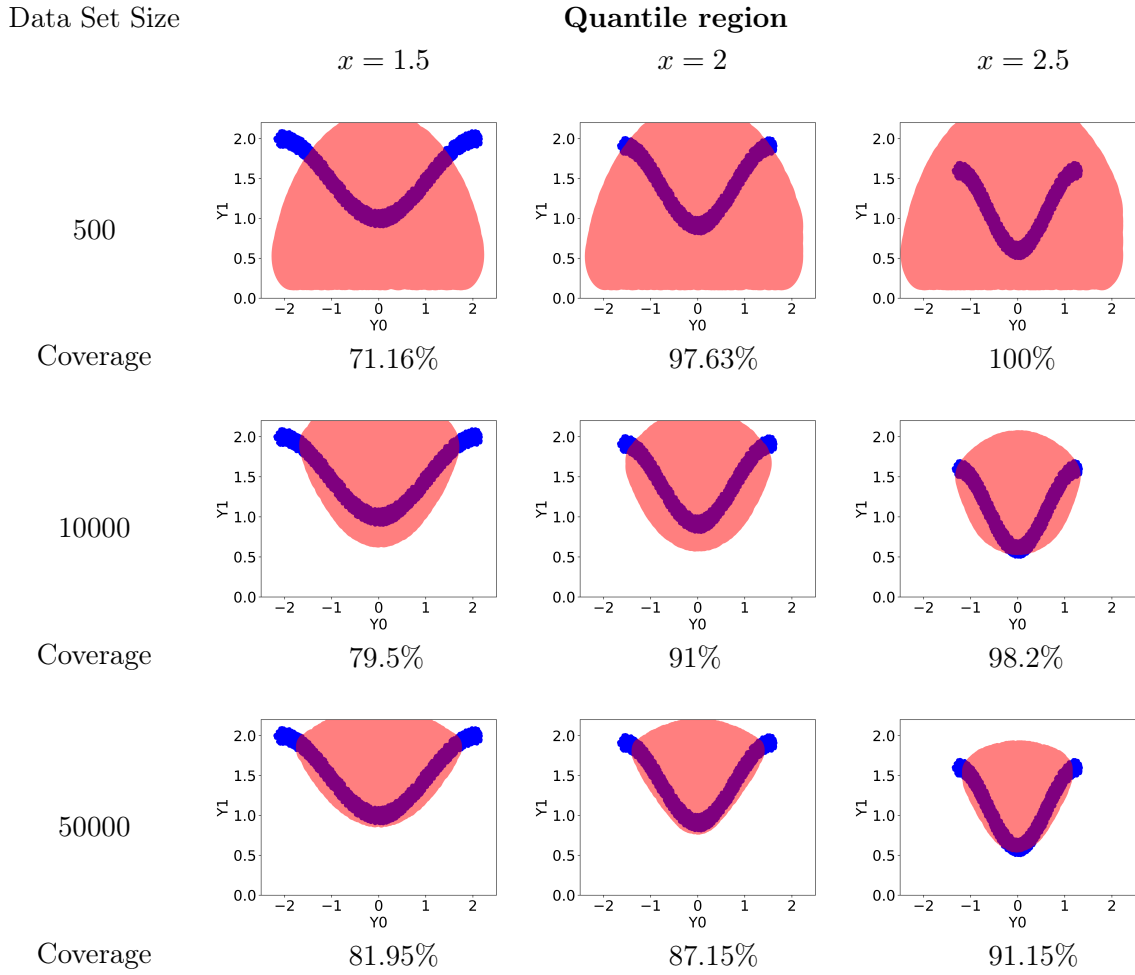


Figure 15: Calibrated quantile regions constructed by NPDQR for the linear synthetic data set with $p = 10$ for different data set sizes.

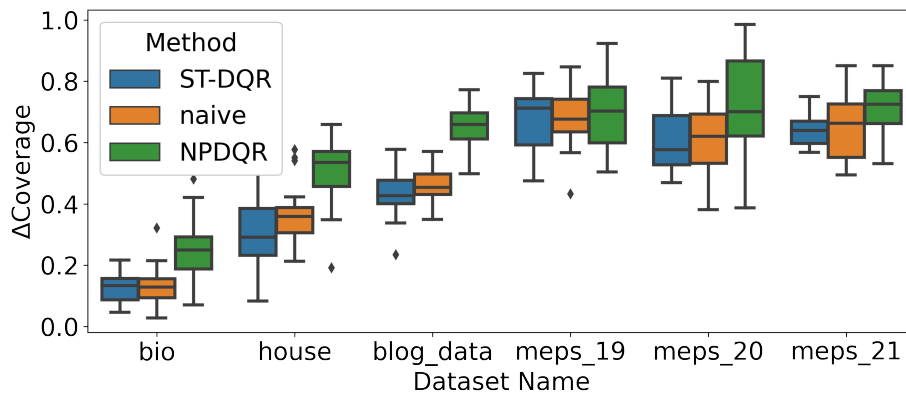


Figure 16: Δ Coverage achieved on the real data sets.

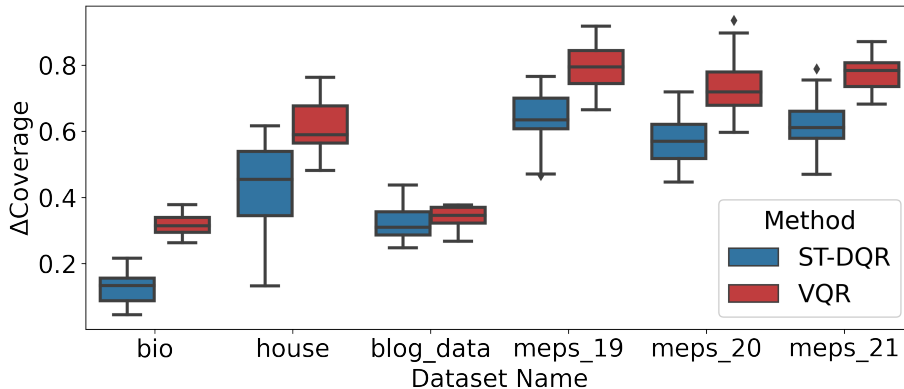


Figure 17: Δ Coverage achieved on the reduced version of the real data sets where all feature vectors were reduced to dimension 10 using PCA.

Setting	d	p	ST-DQR	Naive QR	NPDQR	VQR
linear	2	1	89.943 (.126)	90.059 (.147)	90.041 (.162)	89.755 (.101)
linear	2	10	89.926 (.137)	89.789 (.154)	90.131 (.115)	90.065 (.125)
linear	2	50	89.91 (.08)	89.99 (.058)	89.96 (.07)	-
linear	2	100	89.963 (.082)	89.993 (.059)	90.003 (.062)	-
nonlinear	2	1	90.126 (.169)	90.078 (.171)	90.165 (.163)	90.13 (.145)
nonlinear	3	1	90.165 (.139)	90.114 (.16)	90.021 (.141)	90.156 (.131)
nonlinear	3	10	89.991 (.164)	89.881 (.173)	90.051 (.133)	-
nonlinear	4	1	90.031 (.125)	90.175 (.109)	89.955 (.121)	-
nonlinear	4	10	89.792 (.141)	89.841 (.155)	89.956 (.138)	-

Table 10: Simulated data experiments. Coverage rate and standard error achieved with each method.

Setting	d	p	ST-DQR	Naïve QR	NPDQR	VQR
linear	2	1	315.313 (1.539)	1223.439 (3.374)	1024.657 (13.404)	465.427 (3.019)
linear	2	10	384.33 (1.805)	1680.308 (3.408)	1622.519 (15.904)	485.761 (2.263)
linear	2	50	388.561 (1.97)	1777.014 (1.816)	1525.568 (13.966)	-
linear	2	100	456.75 (5.01)	1791.202 (1.751)	1555.475 (23.214)	-
nonlinear	2	1	236.668 (2.11)	797.438 (2.514)	694.339 (8.421)	646.065 (2.643)
nonlinear	3	1	933.924 (12.592)	22984.579 (108.375)	19982.523 (1050.929)	7621.711 (135.156)
nonlinear	3	10	1251.648 (32.406)	44070.685 (140.94)	34917.304 (2498.402)	-
nonlinear	4	1	379.863 (7.52)	27364.108 (154.317)	82495.522 (5832.992)	-
nonlinear	4	10	293.998 (5.958)	53999.244 (210.159)	129599.146 (9656.909)	-

Table 11: Simulated data experiments. Area and standard error of the quantile regions constructed by each method.

Data Set name	Coverage rate			Area of quantile regions		
	ST-DQR	Naïve QR	NPDQR	ST-DQR	Naïve QR	NPDQR
bio	90.0 (.093)	90.002 (.106)	89.892 (.102)	333.057 (3.8)	407.304 (5.333)	406.852 (6.955)
house	90.157 (.133)	89.978 (.133)	89.876 (.143)	360.616 (6.789)	421.087 (4.062)	414.327 (5.349)
blog_data	90.145 (.089)	90.064 (.076)	90.016 (.075)	38.527 (11.535)	214.8 (4.857)	252.249 (7.093)
meps_19	90.144 (.121)	89.892 (.108)	89.865 (.163)	188.503 (7.585)	417.525 (19.792)	408.791 (11.748)
meps_20	89.997 (.145)	90.0 (.135)	89.913 (.126)	190.762 (3.866)	433.061 (12.122)	421.309 (11.149)
meps_21	89.899 (.14)	89.879 (.092)	89.676 (.15)	196.867 (7.593)	431.402 (15.341)	435.418 (10.448)

Table 12: Real data experiments. Coverage, area, and standard error of the quantile regions constructed by each method.

F.1 The Initial Distance Threshold

In this section, we formally define γ_{init} . Let $R_{\mathcal{Y}} = \{a_1, a_2, \dots, a_m\}$, and denote the distance of a_i from its neighbor by $D_i = \min_{a \in R_{\mathcal{Y}}(x)} d(a_i, a)$. The initial distance threshold γ_{init} is defined as the 90-th smallest quantile of $\{D_i : i = 1, \dots, m\}$.

Data Set	$r = 1$	$r = 2$	$r = 3$	$r = 4$
bio	317.316 (3.004)	321.482 (3.447)	333.057 (3.8)	342.314 (3.941)
house	268.405 (2.949)	326.061 (6.588)	360.616 (6.789)	359.062 (5.866)
meps_20	192.848 (4.43)	198.335 (4.181)	190.762 (3.866)	190.021 (4.225)

Table 13: Area and its standard error of quantile regions constructed on real data sets using ST-DQR with different values of r .

Data Set	Coverage rate				Area of quantile regions			
	ST-DQR	Naive QR	NPDQR	VQR	ST-DQR	Naive QR	NPDQR	VQR
bio	90.0 (.093)	90.002 (.106)	89.892 (.102)	89.87 (.104)	333.057 (3.8)	407.304 (5.333)	406.852 (6.955)	530.053 (5.189)
house	89.923 (.159)	90.094 (.111)	90.067 (.149)	90.119 (.172)	384.985 (6.94)	577.808 (5.167)	548.684 (6.955)	523.282 (4.645)
blog_data	90.078 (.075)	90.035 (.072)	89.823 (.078)	90.034 (.076)	208.05 (3.63)	362.808 (6.332)	405.793 (8.749)	344.921 (4.684)
meps_19	90.179 (.121)	90.16 (.105)	89.919 (.093)	90.149 (.128)	249.926 (9.727)	704.739 (27.762)	381.729 (11.357)	269.517 (11.503)
meps_20	90.087 (.158)	89.925 (.139)	90.036 (.158)	90.134 (.179)	240.692 (5.099)	664.476 (19.175)	358.091 (7.766)	264.666 (6.205)
meps_21	90.061 (.152)	89.957 (.149)	89.965 (.166)	89.887 (.111)	249.156 (8.992)	709.56 (24.683)	383.216 (13.327)	264.41 (10.889)

Table 14: Reduced real data experiments. Coverage, area, and standard error of the quantile regions constructed by each method. All feature vectors were reduced to dimension 10 using PCA.

F.2 Grid Size

Recall that our method maps a discretized quantile region $R_{\mathcal{Z}}(x)$ to space \mathcal{Y} . We also discretize the space \mathcal{Y} to compute the area of a quantile region. In this section, we describe how we make these discretizations in practice. For both tasks, we define a grid in space \mathcal{Y} or space \mathcal{Z} , and set the grid’s boundaries in each dimension to be the 99% and 1% empirical quantiles of the training set’s response variables, respectively. We also increase the upper and lower boundaries by 1 for a grid used to discretize a quantile region, i.e., $R_{\mathcal{Y}}(x)$ or $R_{\mathcal{Z}}(x)$, and by 0.2 for a grid used to measure the area of a region in \mathcal{Y} . Table 15 shows the number of points in the grid, depending on its purpose. In all cases, the grid is equally spaced, having the same number of points in each dimension.

Space dimension	Area calculation	$R_{\mathcal{Y}}(x)$ or $R_{\mathcal{Z}}(x)$
2	3025	$1e^4$
3	103823	42875
4	234256	104976

Table 15: Number of cells in a grid used for area calculation, or for discretizing the quantile regions $R_{\mathcal{Y}}(x)$ and $R_{\mathcal{Z}}(x)$.

F.3 Dealing With High-Dimensional Responses

In many cases, it is required to quantify the uncertainty of a high-dimensional response that has more than three dimensions. In this section, we present several ways to interpret the quantile regions produced by our ST-DQR in the high-dimensional setting.

Dimension-wise prediction interval. The simplest way to interpret a high-dimensional quantile region is to construct an interval for each dimension. For instance, given a quantile region $R(X) \subseteq \mathcal{Y}$, the prediction interval for the j -th variable of the response can be defined as: $C^j(X) = [\min(R(X))^j, \max(R(X))^j]$. The Cartesian product of these intervals forms

the complement of $R(X)$ to a rectangle. Furthermore, this rectangle is more conservative as it covers more area than the region $R(X)$.

Querying a point coverage. The quantile region can be used to query whether certain points fall inside it or not. By doing so, the user can check if a scenario is likely and take a decision according to the response.

Visualizing 3d slices. Assuming there are d responses, one can visualize the quantile region of each triplet $(i_1, i_2, i_3) \in \{1, \dots, d\}^3$ while the other coordinates are held fixed at some value. This way, the user can learn the relation between each response triplet.

F.4 Conditional Variational Auto Encoder

A *conditional variational auto encoder* (CVAE), proposed by Sohn et al. (2015), is a model that aims to capture the implicit conditional distribution of $Y | X$ by reconstructing the input data from its encoded representation. Such a model is composed of an encoder \mathcal{E} that encodes $Y | X$ into a latent variable Z_x , and a decoder \mathcal{D} that reconstructs $(Z_x; x)$ back to the original variable $Y | X$. Formally, CVAE is a pair $(\mathcal{E}(\cdot; x), \mathcal{D}(\cdot; x))$ that ideally satisfies:

$$\begin{aligned} Z_x &= \mathcal{E}(Y; X = x) \sim \mathcal{N}(0, 1)^r, \\ \mathcal{D}(\mathcal{E}(Y; X = x); X = x) &\stackrel{d}{=} Y | X = x, \end{aligned}$$

where $\mathcal{N}(0, 1)$ is a standard normal distribution. In practice, we estimate $(\mathcal{E}, \mathcal{D})$ with a neural network. The loss used for training the CVAE model includes two components: a reconstruction loss (measuring the mean squared error) that penalizes for inaccurate reconstruction, and an additional term that encourages the encoded vector to be normally distributed. In this work, we used a loss that measures the KL-divergence (Kullback, 1997; Kullback and Leibler, 1951), between the encoded representations and the multivariate normal distribution, multiplied by a coefficient λ . Below, we examine the effects of the KL-divergence penalty coefficient and the dimension of the encoded space on the performance of the CVAE model. To maximize performance, we suggest choosing these hyper-parameters via cross-validation.

F.4.1 KL-DIVERGENCE LOSS COEFFICIENT

In this section, we investigate the effects of the KL-divergence loss penalty coefficient, denoted by λ , on the performance of the resulting CVAE model. Figure 18 presents the encoded and reconstructed samples, computed by CVAE models with different values of λ on the non-linear synthetic data set, with $p = 1$; see more details about the synthetic data in Section B.1. We report the results only for the conditional distribution of $Y | X = 2$. The figure shows that as λ increases, the distribution of Z_x is closer to standard normal, but the reconstruction becomes worse. On the other hand, as λ decreases, the CVAE recovers better the input data, but the distribution of Z_x is far from being standard normal. Relying on that figure, we set λ to be equal to 0.01 for all of the experiments presented in this paper.

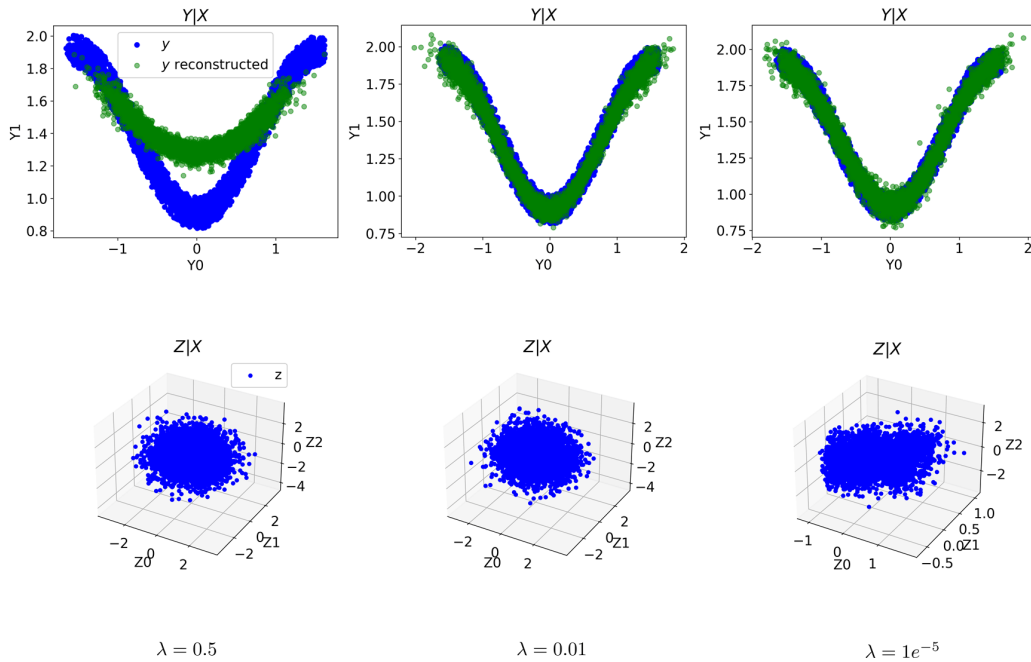


Figure 18: CVAE results for different λ s.

F.4.2 THE DIMENSION OF THE LATENT SPACE

In this section, we examine the second hyper-parameter of the CVAE, which is the dimension of the latent vector Z , denoted by r . Specifically, we study the effects of r on the CVAE performance as well as the efficiency of the constructed quantile regions. The reconstruction errors for different values of r are summarized in Table 16. The table shows that the CVAE reconstructs better the data when using a larger latent space dimension. This trend is also visualized in Figure 19. This is anticipated, as the expressiveness of the CVAE depends on r , in the sense that as r increases, the CVAE is more flexible. In practice, we recommend setting this hyper-parameter to be the dimension of the response, or choosing it using cross-validation.

Data Set	$r = 1$	$r = 2$	$r = 3$	$r = 4$
bio	0.034 (6e-3)	0.021 (7e-3)	0.021 (7e-3)	0.021 (.001)
house	0.036 (6e-3)	0.031 (.002)	0.026 (.002)	0.025 (.002)
meps_20	0.104 (.003)	0.025 (6e-3)	0.024 (7e-3)	0.025 (4e-3)

Table 16: CVAE reconstruction loss and its standard error on real data sets achieved with different latent space dimensions.

We now turn to test how r affects the area of a quantile region constructed to cover $1 - \alpha = 90\%$ of the data. The levels of the estimated directional quantiles used for constructing the quantile region are summarized in Table 8. Table 17 displays the area of quantile regions

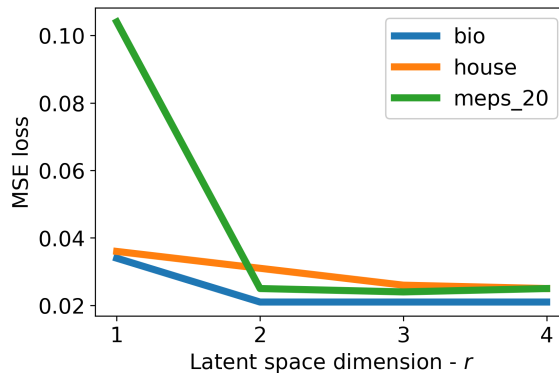


Figure 19: CVAE reconstruction (MSE) loss as a function of the latent space dimension for bio, house, and meps_20 data sets.

constructed by our method with different values of r , scaled by the area achieved by the best method. Following that table, the best latent space dimension for bio and house data sets is $r = 1$, whereas for meps_20 data set it is $r = 4$. This phenomenon can be explained in the following way. The coverage of a quantile region constructed in space \mathcal{Z} depends on the dimension r ; see Figure 6. As the dimension increases, the coverage rate of the DQR region decreases, and this coverage is preserved when transforming the region to space \mathcal{Y} ; see Proposition 3. On the other hand, as r increases, the CVAE reconstructs better the data, which results in more accurate regions. Therefore, there is a trade-off between the reconstruction loss and the coverage attained in space \mathcal{Z} , where both are determined by r . Since the CVAE’s reconstruction on bio and house data sets does not significantly improve for larger values of r , the limitation of the DQR has a stronger impact on the produced region, and as a result, ST-DQR with r set to 1 constructs smaller quantile regions. However, meps_20 data set requires a large value of r to reconstruct well the data, as presented in Figure 19, and therefore using $r = 4$ results in the best quantile regions. To conclude, the best value of r varies between the data sets, so we recommend finding it via cross-validation to maximize the performance of the method.

Data Set	$r = 1$	$r = 2$	$r = 3$	$r = 4$
bio	1	1.013	1.05	1.079
house	1	1.215	1.344	1.338
meps_20	1.015	1.044	1.004	1

Table 17: Area of quantile regions constructed on real data sets using our method with different values of r , scaled by the area achieved by the best method. The standard errors are reported in Table 13

F.5 Coverage Rate of a DQR Quantile Region

In this section, we compute the coverage rate of a quantile region constructed with NPDQR when the nominal coverage level is set to $1 - \alpha \geq 0.5$. In this setting, the training samples given to NPDQR are $\{(X_i, Z_i)\}_{i=1}^n$, so the model is fitted in space \mathcal{Z} . We assume that the CVAE is ideal, i.e., $Z_i \sim \mathcal{N}(0, 1)^r$, where $r \in \mathbb{N}^+$.

For any direction $u \in \mathbb{S}^{r-1}$, the half-space defined by NPDQR satisfies:

$$\mathbb{P}(Z^r \in H_u^+(x)) = 1 - \alpha.$$

Let $Z^r \sim \mathcal{N}(0, 1)^r$, and $Z \sim \mathcal{N}(0, 1)$. Note that

$$1 - \alpha = \mathbb{P}(Z^r \in H_u^+(x)) = \mathbb{P}(u^T Z^r \geq C),$$

for $C = \Phi^{-1}(\alpha) \leq 0$. Denote by χ_r^2 a chi-squared random variable with r degrees of freedom. The probability to lie inside the quantile region is:

$$\begin{aligned} \mathbb{P}(Z^r \in R(x)) &= \mathbb{P}(\forall u \in \mathbb{S}^{r-1} : Z^r \in H_u^+(x)) \\ &= \mathbb{P}(\forall u \in \mathbb{S}^{r-1} : u^T Z^r \geq C) \\ &= \mathbb{P}\left(\max_{u \in \mathbb{S}^{r-1}} u^T Z^r \leq -C\right) \\ &= \mathbb{P}\left(\frac{Z^{rT}}{\|Z^r\|_2} Z^r \leq -C\right) \\ &= \mathbb{P}(\|Z^r\|_2 \leq -C) \\ &= \mathbb{P}(\chi_r^2 \leq C^2) \\ &= \mathbb{P}(\chi_r^2 \leq \Phi^{-1}(\alpha)^2) \\ &= F_{\chi_r^2}(\Phi^{-1}(\alpha)^2), \end{aligned}$$

where $F_{\chi_r^2}$ is the CDF of chi-squared distribution with r degrees of freedom. In Section 4.1 we use this computation to report the exact coverage rate of NPDQR as a function of the nominal coverage level, and the dimension of the response r .

F.6 Calibration - Shrinking The Quantile Region

In this section, we demonstrate the necessity of handling the two calibration cases (Case 1 and Case 2) separately. Recall that in Case 1 $c_{\text{init}} \leq 1 - \alpha$, therefore we grow the base region (5), and in the complementary case, Case 2, we shrink it. If these two cases would not be treated separately, one would shrink a large region by following the technique presented in Case 1, which is mainly suited for growing small regions; see Section 4.2. Figure 20 illustrates the outcome of shrinking a large region using the procedure described in Case 1. The figure shows that the calibrated region is not continuous, and its boundaries are the same as those of the base region. That is, shrinking the base region according to Case 1 procedure does not yield the desired region. However, shrinking the base region according to the method described in Case 2 (see Section 4.2) does result in a smaller, continuous quantile region, as displayed in Figure 21. Therefore, we suggest treating the two cases separately.

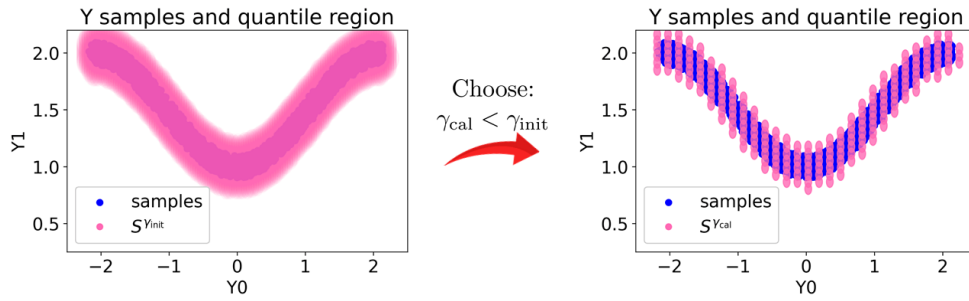


Figure 20: Visualization of a large quantile region (pink) calibrated according to Case 1. The resulted calibrated region is not continuous and is not a smaller version of the base region.




Figure 21: Visualization of the calibration of the large base region (pink) from Figure 20, where the calibration is done according to Case 2. The calibrated region (red) is both continuous, and smaller than the base one.

F.7 Vector Quantile Regression

Given a test point $X_{n+1} = x$, we estimate the vector quantile function $Q(\cdot; x)$ as described by Carlier et al. (2016), using Gurobi solver (Gurobi Optimization, LLC, 2021).¹ This process results in a set of m quantiles $\{Q(U_i; x)\}_{i=1}^m$, for $U_i \sim \text{Uniform}[0, 1]^d$, where $\text{Uniform}(a, b)$ is a uniform distribution on the interval (a, b) . To construct the quantile region, we follow the procedure developed by Chernozhukov et al. (2017). We sample m directions, denoted by $\{S_j\}_{j=1}^m$, from a spherical uniform distribution of dimension d . Next, we match every element in $\{S_j\}_{j=1}^m$ to an element in $\{U_i\}_{i=1}^m$ by solving the following assignment problem:

$$\begin{aligned} & \text{minimize } \sum_{i \in [m]} \|U_i - f(U_i)\|_2, \\ & \text{subject to } \forall i \in [m] : f(U_i) \in \{S_j\}_{j=1}^m \wedge f \text{ is a permutation.} \end{aligned}$$

The discrete set of points inside the quantile region is given by:

$$R_{\mathcal{Y}}(x) = \{Q(U_i; x) : i \in [m] \wedge \|f(U_i)\|_2 \leq 1 - \alpha\}.$$

Setting	d	p	Run time (hours)	Memory consumption (GB)
linear	2	1	< 1	< 13
linear	2	10	< 1	< 13
linear	2	50	> 1	< 13
non-linear	2	1	< 1	< 13
non-linear	3	1	< 1	< 13
non-linear	3	10	-	> 13
non-linear	4	1	-	> 13
non-linear	4	10	-	> 13

Table 18: VQR run time and memory usage on the synthetic data.

References

Anastasios N. Angelopoulos, Stephen Bates, Emmanuel J. Candès, Michael I. Jordan, and Lihua Lei. Learn then test: Calibrating predictive algorithms to achieve risk control. *arXiv preprint*, 2021. arXiv:2110.01052.

Anastasios Nikolas Angelopoulos and Stephen Bates. A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv preprint*, 2021. arXiv:2107.07511.

1. We used the implementation from <https://github.com/alfredgalichon/VQR>

- Rina Foygel Barber, Emmanuel J Candes, Aaditya Ramdas, and Ryan J Tibshirani. Predictive inference with the jackknife+. *The Annals of Statistics*, 49(1):486–507, 2021.
- Stephen Bates, Anastasios Angelopoulos, Lihua Lei, Jitendra Malik, and Michael I. Jordan. Distribution-free, risk-controlling prediction sets. *Journal of the ACM*, 68(6), September 2021. ISSN 0004-5411. doi: 10.1145/3478535. URL <https://doi.org/10.1145/3478535>.
- Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- bio. Physicochemical properties of protein tertiary structure data set. <https://archive.ics.uci.edu/ml/datasets/Physicochemical+Properties+of+Protein+Tertiary+Structure>. Accessed: January, 2019.
- blog_data. Blogfeedback data set. <https://archive.ics.uci.edu/ml/datasets/BlogFeedback>. Accessed: January, 2019.
- Pavel Boček and Miroslav Šiman. Directional quantile regression in r. *Kybernetika*, 53(3):480–492, 2017.
- Guillaume Carlier, Victor Chernozhukov, and Alfred Galichon. Vector quantile regression: an optimal transport approach. *The Annals of Statistics*, 44(3):1165–1192, 2016.
- Guillaume Carlier, Victor Chernozhukov, and Alfred Galichon. Vector quantile regression beyond the specified case. *Journal of Multivariate Analysis*, 161:96–102, 2017.
- Guillaume Carlier, Victor Chernozhukov, Gwendoline De Bie, and Alfred Galichon. Vector quantile regression and optimal transport, from theory to numerics. *Empirical Economics*, pages 1–28, 2020.
- Isabelle Charlier, Davy Paindaveine, and Jérôme Saracco. Multiple-output quantile regression through optimal quantization. *Scandinavian Journal of Statistics*, 47(1):250–278, 2020.
- Victor Chernozhukov, Alfred Galichon, Marc Hallin, and Marc Henry. Monge–kantorovich depth, quantiles, ranks and signs. *The Annals of Statistics*, 45(1):223–256, 2017.
- Victor Chernozhukov, Kaspar Wüthrich, and Yinchu Zhu. Distributional conformal prediction. *arXiv preprint*, 2021. arXiv:1909.07889.
- Youngseog Chung, Willie Neiswanger, Ian Char, and Jeff Schneider. Beyond pinball loss: Quantile methods for calibrated uncertainty quantification. *arXiv preprint*, 2020. arXiv:2011.09588.
- Jean Diebolt, Mhamed-Ali El-Aroui, Valérie Durbec, and Benjamain Villain. Estimation of extreme quantiles: Empirical tools for methods assessment and comparison. *International Journal of Reliability Quality and Safety Engineering*, 7:75–94, 03 2000.
- Shai Feldman, Stephen Bates, and Yaniv Romano. Improving conditional coverage via orthogonal quantile regression. *arXiv preprint*, 2021. arXiv:2106.00394.

- Leying Guan. Conformal prediction with localization. *arXiv preprint*, 2019. arXiv:1908.08558.
- Chirag Gupta, Arun K. Kuchibhotla, and Aaditya K. Ramdas. Nested conformal prediction and quantile out-of-bag ensemble methods. *arXiv preprint arXiv:1910.10562*, 2021.
- Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2021.
- Marc Hallin. On distribution and quantile functions, ranks and signs in \mathbb{R}^d . *ECARES WP*, 2017.
- Marc Hallin, Davy Paindaveine, and Miroslav Šiman. Multivariate quantiles and multiple-output regression quantiles: From L1 optimization to halfspace depth. *The Annals of Statistics*, 38(2):635 – 669, 2010. doi: 10.1214/09-AOS723. URL <https://doi.org/10.1214/09-AOS723>.
- Marc Hallin, Zudi Lu, Davy Paindaveine, and Miroslav Šiman. DQR half-spaces figure. <https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.800.2596>, 2013. Accessed: September, 2021.
- Marc Hallin, Zudi Lu, Davy Paindaveine, and Miroslav Šiman. Local bilinear multiple-output quantile/depth regression. *Bernoulli*, 21(3):1435–1466, 2015.
- Marc Hallin, Eustasio Del Barrio, Juan Cuesta-Albertos, and Carlos Matrán. Distribution and quantile functions, ranks and signs in dimension d: A measure transportation approach. *The Annals of Statistics*, 49(2):1139–1165, 2021.
- J. A. Hartigan and M. A. Wong. Algorithm AS 136: A K-Means clustering algorithm. *Applied Statistics*, 28(1):100–108, 1979.
- house. House sales in king county, USA. <https://www.kaggle.com/harlfoxem/housesalesprediction/metadata>. Accessed: July, 2021.
- David R. Hunter and Kenneth Lange. Quantile regression via an MM algorithm. *Journal of Computational and Graphical Statistics*, 9(1):60–77, 2000. ISSN 10618600.
- Rafael Izbicki, Gilson Shimizu, and Rafael Stern. Flexible distribution-free conditional predictive bands using density estimators. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108, pages 3068–3077. PMLR, 26–28 Aug 2020.
- Rafael Izbicki, Gilson Shimizu, and Rafael B. Stern. CD-split and HPD-split: efficient conformal regions in high dimensions. *arXiv preprint arXiv:2007.12778*, 2021.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations*, 2015.
- Danijel Kivaranovic, Kory D Johnson, and Hannes Leeb. Adaptive, distribution-free prediction intervals for deep networks. In *International Conference on Artificial Intelligence and Statistics*, pages 4346–4356. PMLR, 2020.

- Ivan Kobyzev, Simon Prince, and Marcus Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- Roger Koenker and Gilbert Bassett. Regression quantiles. *Econometrica*, 46(1):33–50, 1978. ISSN 00129682, 14680262.
- Roger Koenker and Kevin F. Hallock. Quantile regression. *Journal of Economic Perspectives*, 15(4):143–156, December 2001.
- Linglong Kong and Ivan Mizera. Quantile tomography: using quantiles with multivariate data. *Statistica Sinica*, pages 1589–1610, 2012.
- S. Kullback and R. A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79 – 86, 1951.
- Solomon Kullback. *Information Theory and Statistics*. Courier Corporation, 1997.
- Xiaohui Liu, Karl Mosler, and Pavlo Mozharovskyi. Fast computation of tukey trimmed regions and median in dimension $p > 2$. *Journal of Computational and Graphical Statistics*, 28(3):682–697, 2019.
- Nicolai Meinshausen. Quantile regression forests. *J. Mach. Learn. Res.*, 7:983–999, December 2006. ISSN 1532-4435.
- meps_19. Medical expenditure panel survey, panel 19. https://meps.ahrq.gov/mepsweb/data_stats/download_data_files_detail.jsp?cboPufNumber=HC-181. Accessed: January, 2019.
- meps_20. Medical expenditure panel survey, panel 20. https://meps.ahrq.gov/mepsweb/data_stats/download_data_files_detail.jsp?cboPufNumber=HC-181. Accessed: January, 2019.
- meps_21. Medical expenditure panel survey, panel 21. https://meps.ahrq.gov/mepsweb/data_stats/download_data_files_detail.jsp?cboPufNumber=HC-192. Accessed: January, 2019.
- Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. *arXiv preprint arXiv:1711.00937*, 2017. arXiv: arXiv:1711.00937.
- Davy Paindaveine and Miroslav Šíman. On directional multiple-output quantile regression. *Journal of Multivariate Analysis*, 102(2):193–212, 2011.
- Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pages 1530–1538. PMLR, 2015.
- Yaniv Romano, Evan Patterson, and Emmanuel Candes. Conformalized quantile regression. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

- B Schuell, T Gruenberger, GV Kornek, N Dworan, D Depisch, F Lang, B Schneeweiss, and W Scheithauer. Side effects during chemotherapy predict tumour response in advanced colorectal cancer. *British Journal of Cancer*, 93(7):744–748, 2005.
- Matteo Sesia and Emmanuel J Candès. A comparison of some conformal quantile regression methods. *Stat*, 9(1):e261, 2020.
- Kihyuk Sohn, Honglak Lee, and Xinchun Yan. Learning structured output representation using deep conditional generative models. *Advances in Neural Information Processing Systems*, 28:3483–3491, 2015.
- Ingo Steinwart and Andreas Christmann. Estimating conditional quantiles with the help of the pinball loss. *Bernoulli*, 17(1), Feb 2011. ISSN 1350-7265.
- Ichiro Takeuchi, Quoc V. Le, Timothy D. Sears, and Alexander J. Smola. Nonparametric quantile estimation. *Journal of Machine Learning Research*, 7(45):1231–1264, 2006.
- James W. Taylor. A quantile regression neural network approach to estimating the conditional density of multiperiod returns. *Journal of Forecasting*, 19(4):299–311, 2000.
- John Thickstun, Zaid Harchaoui, and Sham M. Kakade. Learning features of music from scratch. In *International Conference on Learning Representations (ICLR)*, 2017.
- John W Tukey. Mathematics and the picturing of data. In *Proceedings of the International Congress of Mathematicians, Vancouver, 1975*, volume 2, pages 523–531, 1975.
- Vladimir Vovk, Alex Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*. Springer, 2005. doi: 10.1007/b106715.
- Ziping Xu and Ambuj Tewari. Representation learning beyond linear prediction functions. *arXiv preprint*, 2021. arXiv:2105.14989.
- Yijun Zuo and Robert Serfling. General notions of statistical depth function. *Annals of statistics*, pages 461–482, 2000.