

Conformal Thresholded Intervals for Efficient Regression

Rui Luo^{1*} and Zhixin Zhou²

¹Department of Systems Engineering, City University of Hong Kong, Hong Kong SAR, China

²Alpha Benito Research, Los Angeles, USA
 rui Luo@cityu.edu.hk, zzhou@alphabenito.com

Abstract

This paper introduces Conformal Thresholded Intervals (CTI), a novel conformal regression method that aims to produce the smallest possible prediction set with guaranteed coverage. Unlike existing methods that rely on nested conformal frameworks and full conditional distribution estimation, CTI estimates the conditional probability density for a new response to fall into each interquantile interval using off-the-shelf multi-output quantile regression. By leveraging the inverse relationship between interval length and probability density, CTI constructs prediction sets by thresholding the estimated conditional interquantile intervals based on their length. The optimal threshold is determined using a calibration set to ensure marginal coverage, effectively balancing the trade-off between prediction set size and coverage. CTI’s approach is computationally efficient and avoids the complexity of estimating the full conditional distribution. The method is theoretically grounded, with provable guarantees for marginal coverage and achieving the smallest prediction size given by Neyman-Pearson. Extensive experimental results demonstrate that CTI achieves superior performance compared to state-of-the-art conformal regression methods across various datasets, consistently producing smaller prediction sets while maintaining the desired coverage level. The proposed method offers a simple yet effective solution for reliable uncertainty quantification in regression tasks, making it an attractive choice for practitioners seeking accurate and efficient conformal prediction.

Code — <https://github.com/luo-lorry/CTI>

Extended version — <https://arxiv.org/abs/2407.14495>

1 Introduction

Conformal prediction is a powerful framework for constructing prediction sets with finite-sample coverage guarantees. By leveraging exchangeability of the data, conformal methods can convert the output of any machine learning algorithm into a set-valued prediction satisfying the required coverage level, without assumptions on the data distribution. This paper develops a novel conformal prediction method for regression that aims to produce the smallest possible prediction set with guaranteed coverage.

*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Most existing conformal methods for regression either directly predict the lower and upper endpoints of the interval using quantile regression models (Romano, Patterson, and Candès 2019; Kivaranovic, Johnson, and Leeb 2020; Sesia and Candès 2020; Gupta, Kuchibhotla, and Ramdas 2022) or first estimate the full conditional distribution of the response and then invert it to obtain prediction sets (Izbicki, Shimizu, and Stern 2020a; Chernozhukov, Wüthrich, and Zhu 2021). While these approaches perform well in many situations, they may produce sub-optimal prediction sets if the conditional distribution is skewed. Conformal quantile regression typically yields equal-tailed intervals, but the shortest valid interval may be unbalanced. On the other hand, density-based methods can adapt to skewness but typically involve many tuning parameters and more difficult interpretation, which can be complex for practitioners.

To address these limitations, we propose conformal thresholded intervals (CTI), a conformal inference method that seeks the smallest possible prediction set. Instead of relying on an estimate of the full conditional distribution, we use off-the-shelf multi-output quantile regression and construct prediction set by thresholding the estimated conditional interquantile intervals. Compared with conformal histogram regression (CHR) (Sesia and Romano 2021), which first partitions the response space into bins, CTI directly trains a multi-output quantile regression model that uses equiprobable quantiles. This allows us to estimate the conditional probability density for a new response to fall into each interquantile interval, without the need for explicitly binning the response space.

For each sample in the calibration set, we obtain the interquantile interval that its response falls into and find the corresponding probability density estimate. We compute the non-conformity scores based on these estimates. Intuitively, the non-conformity score is higher for a sample that falls into a long interquantile interval and lower for a sample that falls into a short interquantile interval. By adopting a similar thresholding idea as in conformal classification (Sadinle, Lei, and Wasserman 2019; Luo and Zhou 2024), we threshold the intervals according to their length, the inverse of which corresponds to the probability density estimate. At test time, the threshold, i.e., the quantile for non-conformity scores, is used in constructing prediction sets for test samples. Specifically, the interquantile intervals are sorted in as-

ending order of length, and the first ones shorter than or equal to the threshold are kept. We show that the prediction sets generated from thresholding interquantile intervals guarantee marginal coverage and can achieve desired conditional coverage as well as the smallest expected prediction interval length if the multi-output quantile regression model produces true conditional probability density estimates.

The main contributions of this paper are as follows:

1. We propose a novel conformal prediction method for regression, CTI, that aims to produce the smallest possible prediction set with guaranteed coverage. CTI leverages multi-output quantile regression to estimate conditional interquantile intervals and constructs prediction sets by thresholding these intervals.
2. We provide a theoretical analysis showing that CTI guarantees marginal coverage and can achieve desired conditional coverage as well as the smallest expected prediction interval length under certain conditions.
3. We conduct extensive numerical experiments on both simulated and real datasets, demonstrating that CTI compares favorably to existing conformal regression methods in terms of prediction set size while maintaining valid coverage.
4. The CTI method is simple to implement and interpret, making it an attractive choice for practitioners seeking reliable uncertainty quantification in regression tasks.

The rest of this paper is organized as follows. We discuss related work in Section 2. Section 3 introduces the problem setup and competitive methods. Section 4 describes the proposed CTI method in detail. Section 5 presents a theoretical analysis of CTI. Section 6 provides numerical experiments comparing CTI to existing conformal regression methods on both simulated and real data. Finally, Section 7 concludes with a discussion of the main results and future directions.

2 Related Work

Quantile regression (Koenker 2005) estimates the τ -th conditional quantile function by minimizing the check function loss:

$$\min_{f_\tau} \sum_{i=1}^n \rho_\tau(y_i - f_\tau(x_i)),$$

where

$$\rho_\tau(r) = \begin{cases} \tau r, & \text{if } r > 0 \\ -(1 - \tau)r, & \text{otherwise} \end{cases}$$

is the check function representing the absolute loss.

Quantile regression has been widely used to construct prediction intervals by estimating conditional quantile functions at specific levels, such as the 5% and 95% levels for 90% nominal coverage (Hunter and Lange 2000; Taylor 2000; Meinshausen 2006; Takeuchi et al. 2006; Steinwart and Christmann 2011). This approach adapts to local variability, even for highly heteroscedastic data.

Simultaneous estimation of multiple quantiles is asymptotically more efficient than separate estimation of individual regression quantiles or ignoring within-subject dependency (Cho, Kim, and Kim 2017). However, this approach

does not guarantee non-crossing quantiles, which can affect the validity of the predictions and introduce critical issues in certain scenarios. To address this limitation, research on non-crossing multiple quantile regression has gained attention in recent years, with several methods proposed to ensure non-crossing quantile estimates, including stepwise approaches (Liu and Wu 2009), non-parametric techniques (Cannon 2018), and deep learning-based models (Moon et al. 2021; Brando et al. 2022).

However, the validity of the produced intervals is only guaranteed for specific models under certain regularity and asymptotic conditions (Steinwart and Christmann 2011; Takeuchi et al. 2006; Meinshausen 2006). Many related methods for constructing valid prediction intervals can be encompassed within the nested conformal prediction framework, where a nested sequence of prediction sets is generated by thresholding nonconformity scores derived from various approaches, such as residual-based methods (Papadopoulos et al. 2002; Balasubramanian, Ho, and Vovk 2014; Lei et al. 2018), quantile regression (Romano, Patterson, and Candès 2019; Kivaranovic, Johnson, and Leeb 2020; Sesia and Candès 2020; Chernozhukov, Wüthrich, and Zhu 2021), density estimation (Izbicki, Shimizu, and Stern 2020b; Sesia and Romano 2021; Izbicki, Shimizu, and Stern 2022), and their combinations with ensemble methods (Gupta, Kuchibhotla, and Ramdas 2022) and localized methods (Papadopoulos, Gammerman, and Vovk 2008; Colombo 2023; Luo and Colombo 2024). However, as noted by (Lei, Robins, and Wasserman 2013), the optimal conditionally-valid prediction regions are level sets of conditional densities, which need not be intervals, suggesting that constructing possibly non-convex prediction sets might lead to more efficient conformal predictors.

Our proposed method for constructing non-convex prediction sets is related to the work of (Izbicki, Shimizu, and Stern 2022), who introduce a profile distance to measure the similarity between features and construct prediction sets based on neighboring samples. In contrast, our method directly estimates the conditional probability density for a new response to fall into each interquantile interval based on a multi-output quantile regression model. By thresholding the interquantile intervals based on their length, which is inversely proportional to the estimated probability density, we can construct efficient prediction sets that adapt to the local density of the data. This approach allows us to generate prediction sets that are not restricted to intervals and can potentially achieve better coverage and efficiency compared to interval-based methods.

Another related approach (Guha et al. 2024) converts regression to a classification problem and employs a conditional distribution with a smoothness-enforcing penalty. This method is orthogonal to our approach and can be potentially combined with our multi-output quantile regression framework to further improve the efficiency of the constructed prediction sets.

3 Preliminary and Problem Setup

In this section, we first provide an overview of conformal prediction methods for regression problems in the literature.

We then introduce a simple setting for studying conformal interval arithmetic, where the calibration and test sets are at a group level. In this setting, we propose using the residual of sums to perform conformal prediction for the test group.

Conformal Prediction for Regression

Conformal prediction (Vovk, Gammerman, and Shafer 2005; Shafer and Vovk 2008) is a framework for constructing prediction intervals with guaranteed marginal coverage in regression problems. It leverages *conformity scores* to measure how well each sample fits the model and uses the empirical distribution of these scores from a calibration set to determine prediction interval sizes for new test points.

Given a dataset $\{(x_i, y_i)\}_{i=1}^n$, split into training, calibration, and test sets with index sets $\mathcal{I}_{\text{train}}, \mathcal{I}_{\text{cal}}, \mathcal{I}_{\text{test}}$, a model \hat{f} is first trained on $\mathcal{I}_{\text{train}}$ to produce predictions $\hat{y}_i = \hat{f}(x_i)$. Conformity scores S_i are then computed for the calibration data:

$$S_i = s(x_i, y_i), \quad \forall i \in \mathcal{I}_{\text{cal}}, \quad (1)$$

where s is a score function. In the *split conformal method* (Papadopoulos et al. 2002; Vovk, Gammerman, and Shafer 2005), the score is typically the absolute residual:

$$S_i^{\text{Split}} = |y_i - \hat{f}(x_i)|. \quad (2)$$

For a new test point x_{n+1} , the $1 - \alpha$ prediction interval is:

$$\mathcal{C}_{n,\alpha}(x_{n+1}) = \{y \in \mathbb{R} : s(x_{n+1}, y) \leq t_{1-\alpha}^{\text{Split}}\}, \quad (3)$$

where $t_{1-\alpha}^{\text{Split}}$ is the $(1 - \alpha)(1 + 1/|\mathcal{I}_{\text{cal}}|)$ -th empirical quantile of $\{S_i^{\text{Split}}\}_{i \in \mathcal{I}_{\text{cal}}} \cup \{\infty\}$. Under the exchangeability assumption, this guarantees marginal coverage:

$$\mathbb{P}(Y_{n+1} \in \mathcal{C}(x_{n+1})) \geq 1 - \alpha. \quad (4)$$

Two advanced methods, *Conformal Quantile Regression (CQR)* (Romano, Patterson, and Candès 2019) and *Conformal Histogram Regression (CHR)* (Sesia and Romano 2021), extend this framework:

- **CQR** constructs intervals based on quantile regression estimates:

$$\mathcal{C}^{\text{CQR}}(x_{n+1}) = \left[\hat{q}_{\frac{\alpha}{2}}(x_{n+1}) - t_{1-\alpha}^{\text{CQR}}, \hat{q}_{1-\frac{\alpha}{2}}(x_{n+1}) + t_{1-\alpha}^{\text{CQR}} \right], \quad (5)$$

where $\hat{q}_{\frac{\alpha}{2}}$ and $\hat{q}_{1-\frac{\alpha}{2}}$ are conditional quantile estimates, and $t_{1-\alpha}^{\text{CQR}}$ is the $(1 - \alpha)(1 + 1/|\mathcal{I}_{\text{cal}}|)$ -th empirical quantile of $\{S_i^{\text{CQR}}\}_{i \in \mathcal{I}_{\text{cal}}} \cup \{\infty\}$, with:

$$S_i^{\text{CQR}} = \min(\hat{q}_{\alpha/2}(x_i) - y_i, y_i - \hat{q}_{1-\alpha/2}(x_i)). \quad (6)$$

- **CHR** constructs prediction intervals by estimating the full conditional density $f_{Y|X}$ using histograms and finding the shortest interval (a, b) such that:

$$\mathcal{C}^{\text{CHR}}(x_{n+1}) = \arg \min_{a < b} (b - a), \quad (7)$$

$$\text{s.t. } \int_a^b \hat{f}_{Y|X}(y|x_{n+1}) dy \geq 1 - \alpha. \quad (8)$$

Both methods improve efficiency (tighter intervals) compared to the split conformal method, with CQR adapting to local variability and CHR accommodating non-standard distributions.

Problem Setup

We consider a regression problem with a dataset $\{(x_i, y_i)\}_{i=1}^n$, where $x_i \in \mathcal{X} \subseteq \mathbb{R}^d$ and $y_i \in \mathcal{Y} \subseteq \mathbb{R}$. The dataset is split into training, calibration, and test sets with index sets $\mathcal{I}_{\text{train}}, \mathcal{I}_{\text{cal}}, \mathcal{I}_{\text{test}}$. We assume the calibration and test samples are exchangeable.

The goal is to construct a *conformal predictor* that outputs a prediction set $\mathcal{C}(x) \subseteq \mathcal{Y}$ for each test input x , such that the true response y satisfies:

$$\mathbb{P}(Y \in \mathcal{C}(X)) \geq 1 - \alpha, \quad (9)$$

where $\alpha \in (0, 1)$ is a user-specified significance level. While ensuring valid marginal coverage, we aim to minimize the expected size of the prediction sets:

$$\mathbb{E}[\mu(\mathcal{C}(X))] = \int_{\mathcal{X}} \mu(\mathcal{C}(x)) dP(x), \quad (10)$$

where μ is the Lebesgue measure and $P(x)$ is the marginal distribution of X .

Existing methods like CQR and CHR produce interval-based prediction sets, which may be suboptimal for non-unimodal or non-symmetric conditional distributions. Our objective is to develop a more flexible approach that generates general, possibly non-convex prediction sets, improving efficiency while maintaining valid coverage.

4 Conformal Thresholded Interquantile Intervals

First, we apply quantile regression on the training set $\mathcal{D}_{\text{train}}$ to predict the τ -th quantile of the conditional distribution $Y|X = x$ for every $x \in \mathcal{X}$, where τ takes values from 0 to 1 in increments of $1/K$. The estimated quantile for $\tau = k/K$ is denoted by

$$\hat{q}_k(x) \quad \text{for } k = 0, 1, \dots, K. \quad (11)$$

We then define the interquantile intervals as

$$I_k(x) = (\hat{q}_{k-1}(x), \hat{q}_k(x)] \quad \text{for } k = 1, \dots, K. \quad (12)$$

Assuming the quantile regression provides sufficiently accurate estimations, each interval should have approximately the same probability, $1/K$, of covering the true label Y . To minimize the size of the prediction set, it is more efficient to include intervals with smaller sizes. This strategy leads us to define the confidence set as:

$$\mathcal{C}(x) = \bigcup \{I_k(x) : \mu(I_k(x)) \leq t, k = 1, \dots, K\}, \quad (13)$$

where t is a threshold determined in a marginal sense, meaning it is independent of x . To determine t , we utilize the calibration set. We want t to satisfy the condition that $y_i \in \mathcal{C}(x_i)$ for at least $\lceil (1 + |\mathcal{I}_{\text{cal}}|)(1 - \alpha) \rceil$ instances in the calibration set, where $i \in \mathcal{I}_{\text{cal}}$. We define the threshold t as follows:

$$t = (1 - \alpha)\text{-th quantile of the empirical distribution} \\ \frac{1}{(1 + |\mathcal{I}_{\text{cal}}|)} \sum_{i \in \mathcal{I}_{\text{cal}}} \delta_{\mu(I_{k(y_i)}(x_i))} + \delta_{\infty} \quad (14)$$

where $k(y)$ is the index that of the interval that y belongs, i.e., $y \in I_{k(y)}(x)$. By plugging t back into (13), we obtain the prediction set for every $x \in \mathcal{X}$. The above procedure is summarized the following algorithm.

Algorithm 1: Conformalized Thresholded Intervals

- 1: **Input:** labeled data $\{(x_i, y_i)\}_{i \in \mathcal{I}}$, unlabeled test data $\{x_i\}_{i \in \mathcal{I}_{\text{test}}}$, a data split ratio, black-box learning algorithm \mathcal{B} , level $\alpha \in (0, 1)$, number of interquantile intervals K
 - 2: Randomly split the indices \mathcal{I} into $\mathcal{I}_{\text{train}}$ and \mathcal{I}_{cal} .
 - 3: Train \mathcal{B} on samples in $\mathcal{I}_{\text{train}}$, and obtain quantile estimation functions \hat{q}_k for $k = 0, 1, \dots, K$.
 - 4: For every $i \in \mathcal{I}_{\text{cal}} \cup \mathcal{I}_{\text{test}}$, evaluate $\hat{q}_k(x_i)$ for $k = 0, 1, \dots, K$.
 - 5: For every $i \in \mathcal{I}_{\text{cal}} \cup \mathcal{I}_{\text{test}}$, define the interquantile intervals $I_k(x_i) = (\hat{q}_{k-1}(x_i), \hat{q}_k(x_i)]$ for $k = 1, \dots, K$.
 - 6: $t \leftarrow (1 - \alpha)$ -th quantile of the empirical distribution $\frac{1}{(1 + |\mathcal{I}_{\text{cal}}|)} \sum_{i \in \mathcal{I}_{\text{cal}}} \delta_{\mu(I_{k(y_i)}(x_i))} + \delta_{\infty}$.
 - 7: For $i \in \mathcal{I}_{\text{test}}$, $\mathcal{C}(x_i) = \bigcup \{I_k(x_i) : \mu(I_k(x_i)) \leq t, k = 1, \dots, K\}$.
 - 8: **Output:** $\mathcal{C}(x_i)$ for $i \in \mathcal{I}_{\text{test}}$.
-

Remark. Our approach can also be considered in terms of conformity scores. Using the definition of the prediction set in equation (13), the value of label y is contained within a small interval. More formally, let $k(y)$ be the index such that $y \in I_{k(y)}(x)$. We can then define the score function (1) for our proposed method as:

$$s(x, y) = \mu(I_{k(y)}(x)).$$

This score function assigns a score to each label y based on the size of the interval $I_{k(y)}(x)$ in which it falls. A smaller score indicates that the label y is more likely to be the true label for the input x . In the context of conformal prediction, labels with smaller scores are given priority for inclusion in the prediction set. We present a detailed discussion of the connection of Algorithm 1 and Theorem 1 in the Appendix.

5 Theoretical Analysis

In the context of the entire population, CTI shares a very similar formulation with the Least Ambiguous Set method used for classification, as described in (Sadinle, Lei, and Wasserman 2019). If we assume that our quantile regression model is sufficiently accurate, CTI has the potential to achieve the optimal size for prediction sets when considering the marginal distribution. To understand this better, let's first take a look at the Neyman-Pearson Lemma:

Lemma 1 (Neyman-Pearson). *Let f and g be two nonnegative measurable functions. Then the optimizer of the problem*

$$\min_C \int_C g \quad \text{s.t.} \quad \int_C f \geq 1 - \alpha,$$

is given by $C = \{x : f(x)/g(x) \geq t'\}$ if there exists t such that $\int_{f/g \geq t'} f = 1 - \alpha$.

To formalize the problem of minimizing the expected length of the prediction set subject to $1 - \alpha$ coverage, we can write the problem as:

$$\begin{aligned} \min_{\mathcal{C}(x)} \int_{\mathcal{X}} \int_{\mathcal{C}(x)} 1 d\mu(y) dP(x) \\ \text{s.t.} \quad \int_{\mathcal{X}} \int_{\mathcal{C}(x)} f(y|x) d\mu(y) dP(x) \geq 1 - \alpha. \end{aligned}$$

The Neyman-Pearson Lemma implies that the optimal solution for $\mathcal{C}(x)$ has the form:

$$\mathcal{C}(x) = \{y : f(y|x) \geq t'\} \quad (15)$$

for some suitable threshold t' . Indeed, this threshold can be defined as

$$t' = \inf\{t \in \mathbb{R} : \mathbb{P}(f(Y|X) \geq t) \geq 1 - \alpha\}. \quad (16)$$

which is discussed in the Appendix in detail. Our algorithm is an empirical construction of such an interval. Suppose the quantile regression approximates \hat{q}_τ well. In that case, we have:

$$\int_{y \in I_k(x)} f(y|x) d\mu(y) = \mathbb{P}(Y \in I_k(X)) \approx 1/K.$$

As K approaches infinity, $\mu(I_k(x))$ tends to 0. If $f(y|x)$ is sufficiently smooth, then $f(y|x) \approx 1/(K\mu(I_k(x)))$. The threshold on the length of intervals $\mu(I_k(x)) \leq t$ in equation (13) approximately implies $f(y|x) \geq 1/(Kt)$, which is optimal in the sense of the Neyman-Pearson Lemma. This means that our algorithm, which constructs prediction sets based on the threshold on interval lengths, is an empirical approximation of the optimal solution prescribed by the Neyman-Pearson Lemma.

The demonstration of the coverage probability for CTI follows the same reasoning as the traditional proof used in the general conformal prediction framework.

Theorem 1 (Coverage Probability). *Suppose the samples in $\{(X_i, Y_i)\}_{i \in \mathcal{I}_{\text{cal}} \cup \mathcal{I}_{\text{test}}}$ are exchangeable, then for (X, Y) in the test set, the coverage probability*

$$\mathbb{P}(Y \in \mathcal{C}(X)) \geq 1 - \alpha.$$

The upcoming proposition will demonstrate that the threshold t , which is defined in equation (14) for the length of interquantile intervals, results in a suitable threshold for the distribution of the conditional density $f(Y|X)$, as shown in equation (16).

Proposition 1 (Threshold Consistency). *Suppose the interquantile intervals satisfy*

$$\sup_{t'} \left| F_{f(Y|X)}(t') - \frac{1}{|\mathcal{I}_{\text{cal}}|} \sum_{i \in \mathcal{I}_{\text{cal}}} \mathbf{1}\{\mu(I_{k(y_i)}(x_i)) > \frac{1}{Kt'}\} \right| \leq \epsilon$$

Let t be as defined in (14) and let $1 - \alpha' = \lceil (1 + |\mathcal{I}_{\text{cal}}|)(1 - \alpha) \rceil / |\mathcal{I}_{\text{cal}}|$, then $F_{f(Y|X)}(1/(Kt)) \geq \alpha' - \epsilon$.

This proposition establishes that if the value of ϵ is sufficiently small and the size of the calibration set is large enough, then the value $t' = 1/(Kt)$, where t is the threshold defined in equation (14), will be close to the α -th quantile of the distribution $f(Y|X)$.

Lastly, the theorem presented below demonstrates that the size of the prediction set obtained using CTI will not significantly exceed the size of the theoretically optimal prediction set, which is defined in equation (15).

Theorem 2 (Prediction Set). *Suppose for $x \in \mathcal{X}$,*

$$P(Y \in I_k(x) | X = x) = \int_{y \in I_k(x)} f(y|x) dy \geq \frac{1 - \delta(x)}{K},$$

and suppose $f(y|x)$ has a Lipschitz constant $L(x)$, then

$$\mathcal{C}_{1-\alpha}(x) \subseteq \left\{ y : f(y|x) \geq \frac{1 - \delta_k(x)}{Kt} - \frac{L(x)t}{2} \right\}.$$

The previous proposition demonstrates that the value $1/(Kt)$, where t is the threshold defined in equation (14), is slightly smaller than the α -th quantile of $f(Y|X)$, which represents the theoretically optimal threshold as shown in equation (16). This theorem further illustrates that the prediction set obtained using CTI will include values that are even more conservative. However, we understand that $1/(Kt)$, where t is the threshold defined in equation (14), serves as a relatively stable estimate of the quantile, which is asymptotically equivalent to a constant value. As the number of interquantile intervals K approaches infinity, the threshold t converges to 0. If we additionally assume that our quantile regression model is accurate, meaning that the error term $\delta_k(x)$ is small, then the prediction set $\mathcal{C}_{1-\alpha}(x)$ obtained using CTI will be close to the theoretically optimal prediction set defined in equation (13).

Remark (Comparison with existing methods). *Both CTI and CHR use multi-output quantile regression but differ in their approach to achieving coverage. CHR constructs prediction sets with an expected $1 - \alpha$ coverage for each input x , focusing on conditional coverage. In contrast, CTI adopts a global perspective by thresholding interquantile intervals across all x values. This global approach improves efficiency in achieving marginal coverage by deprioritizing x values where interquantile ranges are large and uncertain. By focusing on the overall distribution, CTI allocates coverage more effectively, resulting in efficient prediction sets.*

Multi-output quantile regression often requires training separate models for each quantile, which can be computationally expensive. CTI addresses this by using Quantile Regression Forests (RF) (Meinshausen 2006) and Quantile Regression Neural Networks (NN) to estimate multiple quantiles simultaneously. RF aggregates quantile estimates from an ensemble of decision trees, avoiding separate optimization for each quantile. Similarly, NN shares parameters across a single neural network to output multiple quantiles as a vector. This efficiency allows CTI to generate prediction sets tailored to user-specified tolerance levels without redundant model training.

To balance expressiveness and computational efficiency, we fix the number of quantiles at $K = 100$ for all datasets. This choice ensures affordable runtime for CTI while maintaining flexibility for various tasks. We present a comprehensive discussion of the choice of K in the Appendix.

As shown in Lemma 1 and Theorem 2, the optimality of CTI is guaranteed by the Neyman-Pearson Lemma. Specifically, as the error term $\delta_k(x) \rightarrow 0$, $K \rightarrow \infty$, and $t \rightarrow 0$, the prediction set S converges to

$$S \rightarrow \{x \mid p(x) \geq \frac{1}{Kt}\},$$

ensuring both validity and efficiency in prediction set construction.

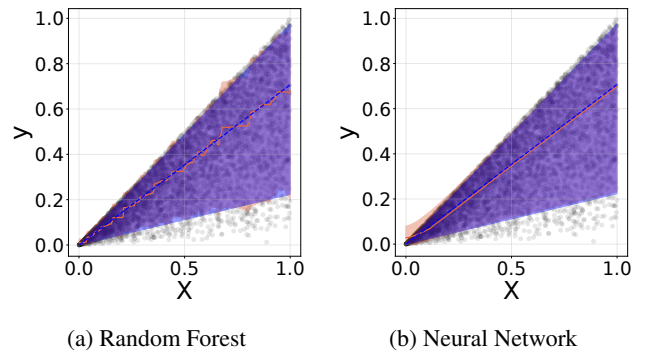


Figure 1: Quantile regression results for the synthetic data. (a) Quantile regression forest and (b) quantile regression neural network show the estimated and theoretical 5%, 50%, and 95% quantiles as well as the 90% intervals. Theoretical quantiles are in blue, empirical in red.

6 Experiment

Simulation Study

We presented an experiment on simulated data to illustrate the importance of adaptivity in conformal prediction. We also derive the analytical quantile function at different levels and the expected prediction set sizes of various conformal methods and demonstrate the superiority of CTI.

To generate the training data, we draw $n = 10000$ independent, univariate predictor samples X_i from the uniform distribution on the interval $[0, 1]$. The response variable is then sampled i.i.d. according to:

$$y \sim \text{Triangular}(0, x, x),$$

where $\text{Triangular}(0, x, x)$ is the Triangular distribution with lower limit 0, upper limit x , and mode x . The conditional density is:

$$f(y|x) = \frac{2y}{x^2} \mathbb{1}\{y \in (0, x)\}.$$

Figure 1 presents the estimated quantiles alongside the theoretical quantile functions. In what follows, we derive the analytical forms of the prediction sets for CQR, CHR, and CTI based on the theoretical quantile functions.

Figure 2 illustrates the theoretical and estimated prediction set sizes as a function of x at $\alpha = 0.1$. The theoretical expected set sizes for CQR, CHR, and CTI are 0.376, 0.342, and 0.317, respectively. The results demonstrate that CTI achieves the smallest set size while maintaining guaranteed coverage.

Real Data

Following the methodology outlined in (Sesia and Candès 2020), we rescale the response Y by the mean absolute value. We randomly allocate 20% of the samples for testing, and from the remaining data, we utilize 70% for training the quantile regression model and 30% for calibration. This split has been validated in (Sesia and Candès 2020). We repeat

Dataset	Metric	CTI(RF)	CTI(NN)	CHR(RF)	CHR(NN)	CQR	Split
synthetic	Coverage	0.900 (0.006)	0.900 (0.007)	0.903 (0.007)	0.901 (0.006)	0.903 (0.008)	0.903 (0.007)
	Size	0.345 (0.005)	0.369 (0.015)	0.375 (0.006)	0.370 (0.014)	0.440 (0.022)	0.482 (0.016)
bike	Coverage	0.898 (0.007)	0.899 (0.007)	0.898 (0.010)	0.900 (0.006)	0.906 (0.009)	0.899 (0.008)
	Size	1.032 (0.029)	0.720 (0.028)	1.124 (0.028)	0.758 (0.047)	1.599 (0.054)	1.345 (0.053)
bio	Coverage	0.900 (0.004)	0.902 (0.004)	0.899 (0.005)	0.900 (0.004)	0.900 (0.003)	0.901 (0.004)
	Size	1.295 (0.018)	1.474 (0.030)	1.450 (0.023)	1.576 (0.012)	2.005 (0.016)	1.961 (0.039)
blog	Coverage	0.910 (0.002)	0.900 (0.004)	0.902 (0.004)	0.902 (0.003)	0.940 (0.009)	0.910 (0.006)
	Size	0.709 (0.031)	1.003 (0.024)	1.567 (0.074)	1.737 (0.154)	3.259 (0.327)	1.453 (0.113)
community	Coverage	0.909 (0.018)	0.908 (0.021)	0.903 (0.015)	0.905 (0.021)	0.889 (0.024)	0.902 (0.024)
	Size	1.611 (0.088)	1.275 (0.095)	1.637 (0.096)	1.588 (0.100)	1.680 (0.078)	2.132 (0.188)
concrete	Coverage	0.908 (0.024)	0.900 (0.031)	0.899 (0.022)	0.900 (0.023)	0.901 (0.024)	0.896 (0.021)
	Size	0.967 (0.035)	0.473 (0.050)	0.933 (0.041)	0.505 (0.144)	0.692 (0.051)	0.619 (0.029)
facebook1	Coverage	0.909 (0.003)	0.899 (0.003)	0.901 (0.004)	0.900 (0.004)	0.945 (0.009)	0.903 (0.002)
	Size	0.766 (0.033)	0.780 (0.023)	1.595 (0.088)	1.379 (0.086)	2.627 (0.329)	2.252 (0.208)
facebook2	Coverage	0.911 (0.002)	0.900 (0.001)	0.899 (0.002)	0.899 (0.002)	0.943 (0.006)	0.904 (0.002)
	Size	0.735 (0.017)	0.773 (0.023)	1.533 (0.053)	1.382 (0.057)	2.661 (0.272)	2.100 (0.108)
homes	Coverage	0.900 (0.005)	0.900 (0.006)	0.899 (0.005)	0.895 (0.007)	0.898 (0.006)	0.897 (0.005)
	Size	0.640 (0.011)	0.515 (0.008)	0.682 (0.012)	0.535 (0.010)	0.851 (0.052)	0.825 (0.072)
meps19	Coverage	0.907 (0.008)	0.902 (0.007)	0.901 (0.007)	0.902 (0.004)	0.932 (0.007)	0.902 (0.010)
	Size	1.760 (0.087)	1.795 (0.061)	2.388 (0.195)	2.602 (0.128)	2.923 (0.170)	3.092 (0.377)
meps20	Coverage	0.904 (0.004)	0.901 (0.007)	0.901 (0.007)	0.901 (0.006)	0.927 (0.009)	0.902 (0.005)
	Size	1.883 (0.067)	1.921 (0.091)	2.376 (0.105)	2.594 (0.140)	2.925 (0.193)	3.154 (0.217)
meps21	Coverage	0.906 (0.005)	0.900 (0.008)	0.900 (0.006)	0.898 (0.004)	0.928 (0.007)	0.905 (0.004)
	Size	1.832 (0.089)	1.866 (0.076)	2.510 (0.167)	2.609 (0.145)	2.971 (0.179)	3.046 (0.199)
star	Coverage	0.903 (0.018)	0.910 (0.017)	0.907 (0.021)	0.897 (0.018)	0.901 (0.016)	0.910 (0.024)
	Size	0.186 (0.006)	0.197 (0.009)	0.182 (0.005)	0.204 (0.009)	0.181 (0.005)	0.181 (0.008)

Table 1: The coverage and size results for various methods are presented in the table. Our proposed Conformalized Thresholded Interval (CTI) method, which utilizes quantile regression based on either random forest (RF) or neural network (NN) models, demonstrates superior performance compared to other methods on most datasets.

all experiments 10 times, starting from the initial data splitting and the training procedure of quantile regression using both random forest (RF) and neural network (NN) models.

Both the CTI and CHR incorporate the quantile regression results from both RF and NN. For CQR, we use the NN results, as it is the choice in the original paper (Romano, Patterson, and Candès 2019). We assess the performance of the generated prediction intervals in terms of coverage and efficiency. Since CTI and CHR utilize the same quantile regression results, the comparison between these two methods is fair with respect to the quality of the quantile regression.

Table 1 presents the coverage probabilities and prediction set sizes for the various methods. CTI, which incorporates quantile regression based on either RF or NN models, outperforms other methods on most datasets.


In a small dataset “star”, which has a relatively small sample size ($n = 2161, d = 39$), CQR outperforms CTI and CHR. The limited number of samples in this dataset may hinder the performance of the multi-output quantile regression model, as it requires sufficient data to accurately capture the underlying relationships between the features and the response variable. We also notice a similar trend in the relative performance comparison of CTI based on random forest and

CTI based on neural network, as well as CHR based on random forest and CHR based on neural network. This suggests that the efficiency of the conformal prediction sets depends on the quality of the multi-output quantile regression. The choice of the underlying model plays a crucial role in the performance of the conformal prediction methods.

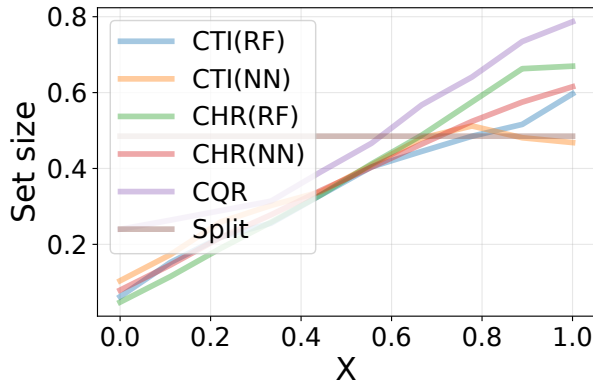
To evaluate the performance of CTI, we compare the lengths of response intervals (intervals containing the actual responses) with the lengths of all intervals generated by the multi-output quantile regression model across all datasets. The distributions of interval lengths for response intervals (blue histogram) and all intervals (red histogram) on the test set are shown in Figure 3. Although, theoretically, the difference in means between the two distributions should be zero, we observe that it is not. This discrepancy highlights the dependence of the prediction sets produced by CTI on the specific dataset and the underlying quantile regression model.

7 Conclusion

Conformal Thresholded Intervals (CTI) is a groundbreaking conformal prediction method specifically designed for re-



(a) Theoretical prediction set at $\alpha = 0.1$



(b) Prediction set sizes at $\alpha = 0.1$

Figure 2: Prediction set sizes for the synthetic data at $\alpha = 0.1$. (a) Theoretical prediction set for different conformal methods. The expected set sizes for CQR, CHR, and CTI are 0.376, 0.342, and 0.317, respectively. (b) Prediction set sizes as a function of x using the estimated quantile functions (RF for random forests and NN for neural network). CTI achieves the smallest set size while maintaining guaranteed coverage.

gression tasks. It aims to construct the smallest possible prediction set while maintaining a guaranteed coverage level. The key innovation behind CTI lies in its clever utilization of multi-output quantile regression and the strategic thresholding of the estimated conditional interquantile intervals based on their length. By doing so, CTI effectively adapts to the local density of the data, enabling it to generate general prediction sets that are not limited to intervals. The ability of CTI to create prediction sets that are not restricted to intervals sets it apart from other conformal prediction methods. This flexibility allows CTI to better capture the intricate patterns and structures present in the data, ultimately leading to more precise and informative prediction sets. However, it is important to note that the current implementation of CTI does not guarantee that the resulting prediction sets will always be intervals. This limitation opens up an exciting avenue for future research, where the goal would be to develop a variant of CTI that can produce interval-based prediction sets while still preserving its adaptivity and efficiency.

Acknowledgments

The work described in this paper was partially supported by a grant from City University of Hong Kong (Project No.9610639).

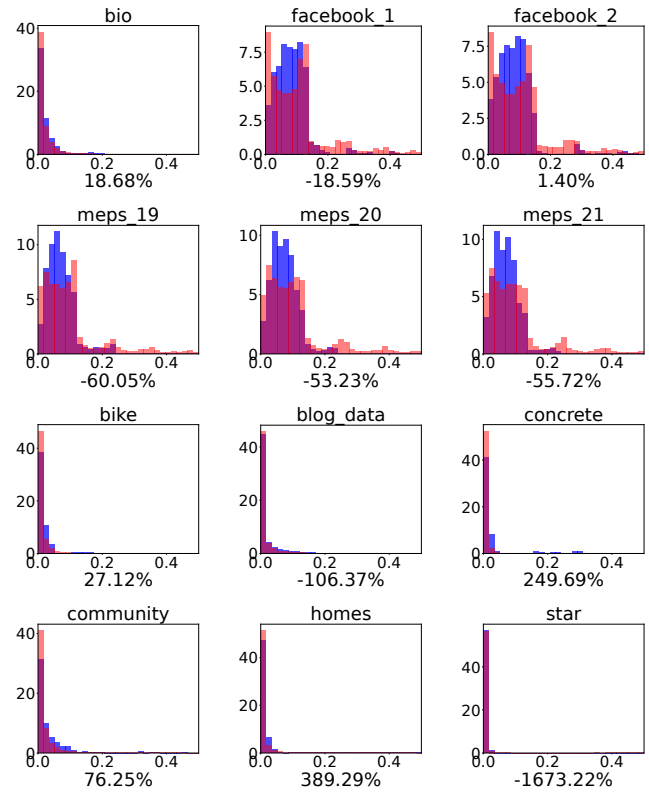


Figure 3: Comparison of interval lengths across datasets: Each subfigure shows the distribution of interval lengths for a specific dataset. The blue histogram represents the intervals containing the actual responses, while the red histogram shows all intervals from the multi-output quantile regression model on the test set. The difference in average length is shown on the x-axis.

References

Balasubramanian, V.; Ho, S.-S.; and Vovk, V. 2014. *Conformal prediction for reliable machine learning: theory, adaptations and applications*. Newnes.

Brando, A.; Center, B. S.; Rodríguez-Serrano, J.; Vitrià, J.; et al. 2022. Deep non-crossing quantiles through the partial derivative. In *International Conference on Artificial Intelligence and Statistics*, 7902–7914. PMLR.

Cannon, A. J. 2018. Non-crossing nonlinear regression quantiles by monotone composite quantile regression neural network, with application to rainfall extremes. *Stochastic environmental research and risk assessment*, 32(11): 3207–3225.

Chernozhukov, V.; Wüthrich, K.; and Zhu, Y. 2021. Distributional conformal prediction. *Proceedings of the National Academy of Sciences*, 118(48): e2107794118.

Cho, H.; Kim, S.; and Kim, M.-O. 2017. Multiple quantile regression analysis of longitudinal data: Heteroscedasticity and efficient estimation. *Journal of Multivariate Analysis*, 155: 334–343.

- Colombo, N. 2023. On training locally adaptive CP. In *Conformal and Probabilistic Prediction with Applications*, 384–398. PMLR.
- Guha, E. K.; Natarajan, S.; Möllenhoff, T.; Khan, M. E.; and Ndiaye, E. 2024. Conformal Prediction via Regression-as-Classification. In *The Twelfth International Conference on Learning Representations*.
- Gupta, C.; Kuchibhotla, A. K.; and Ramdas, A. 2022. Nested conformal prediction and quantile out-of-bag ensemble methods. *Pattern Recognition*, 127: 108496.
- Hunter, D. R.; and Lange, K. 2000. Quantile regression via an MM algorithm. *Journal of Computational and Graphical Statistics*, 9(1): 60–77.
- Izbicki, R.; Shimizu, G.; and Stern, R. 2020a. Flexible distribution-free conditional predictive bands using density estimators. In *International Conference on Artificial Intelligence and Statistics*, 3068–3077. PMLR.
- Izbicki, R.; Shimizu, G.; and Stern, R. 2020b. Flexible distribution-free conditional predictive bands using density estimators. In Chiappa, S.; and Calandra, R., eds., *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, 3068–3077. PMLR.
- Izbicki, R.; Shimizu, G.; and Stern, R. B. 2022. Cd-split and hpd-split: Efficient conformal regions in high dimensions. *Journal of Machine Learning Research*, 23(87): 1–32.
- Kivaranovic, D.; Johnson, K. D.; and Leeb, H. 2020. Adaptive, Distribution-Free Prediction Intervals for Deep Networks. In *International Conference on Artificial Intelligence and Statistics*, 4346–4356. PMLR.
- Koenker, R. 2005. *Quantile regression*, volume 38. Cambridge university press.
- Lei, J.; G’Sell, M.; Rinaldo, A.; Tibshirani, R. J.; and Wasserman, L. 2018. Distribution-Free Predictive Inference for Regression. *Journal of the American Statistical Association*, 113(523): 1094–1111.
- Lei, J.; Robins, J.; and Wasserman, L. 2013. Distribution-free prediction sets. *Journal of the American Statistical Association*, 108(501): 278–287.
- Liu, Y.; and Wu, Y. 2009. Stepwise multiple quantile regression estimation using non-crossing constraints. *Statistics and its Interface*, 2(3): 299–310.
- Luo, R.; and Colombo, N. 2024. Conformal Load Prediction with Transductive Graph Autoencoders. *arXiv preprint arXiv:2406.08281*.
- Luo, R.; and Zhou, Z. 2024. Trustworthy Classification through Rank-Based Conformal Prediction Sets. *arXiv preprint arXiv:2407.04407*.
- Meinshausen, N. 2006. Quantile regression forests. *Journal of Machine Learning Research*, 7: 983–999.
- Moon, S. J.; Jeon, J.-J.; Lee, J. S. H.; and Kim, Y. 2021. Learning multiple quantiles with neural networks. *Journal of Computational and Graphical Statistics*, 30(4): 1238–1248.
- Papadopoulos, H.; Gammelman, A.; and Vovk, V. 2008. Normalized nonconformity measures for regression conformal prediction. In *Proceedings of the IASTED International Conference on Artificial Intelligence and Applications (AIA 2008)*, 64–69.
- Papadopoulos, H.; Proedrou, K.; Vovk, V.; and Gammelman, A. 2002. Inductive confidence machines for regression. In *Machine Learning: ECML 2002: 13th European Conference on Machine Learning Helsinki, Finland, August 19–23, 2002 Proceedings 13*, 345–356. Springer.
- Romano, Y.; Patterson, E.; and Candès, E. 2019. Conformalized quantile regression. In *Advances in Neural Information Processing Systems*, 3538–3548.
- Sadinle, M.; Lei, J.; and Wasserman, L. 2019. Least ambiguous set-valued classifiers with bounded error levels. *Journal of the American Statistical Association*, 114(525): 223–234.
- Sesia, M.; and Candès, E. 2020. A comparison of some conformal quantile regression methods. *Stat*, 9(1).
- Sesia, M.; and Romano, Y. 2021. Conformal prediction using conditional histograms. *Advances in Neural Information Processing Systems*, 34: 6304–6315.
- Shafer, G.; and Vovk, V. 2008. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9: 371–421.
- Steinwart, I.; and Christmann, A. 2011. Estimating conditional quantiles with the help of the pinball loss.
- Takeuchi, I.; Le, Q. V.; Sears, T. D.; Smola, A. J.; and Williams, C. 2006. Nonparametric quantile estimation. *Journal of machine learning research*, 7(7).
- Taylor, J. W. 2000. A quantile regression neural network approach to estimating the conditional density of multiperiod returns. *Journal of Forecasting*, 19(4): 299–311.
- Vovk, V.; Gammelman, A.; and Shafer, G. 2005. *Algorithmic learning in a random world*. Springer.