

Fast Conformal Prediction Using Conditional Interquantile Intervals

Naixin Guo^{1*}, Rui Luo^{2*}, Zhixin Zhou^{3*}

¹Department of Decision Analytics and Operations, City University of Hong Kong, Hong Kong SAR, China

²Department of Systems Engineering, City University of Hong Kong, Hong Kong SAR, China

³Alpha Benito Research, Los Angeles, USA

naixinguo2-c@my.cityu.edu.hk, ruiluo@cityu.edu.hk, zzhou@alphabenito.com

Abstract

We introduce Conformal Interquantile Regression (CIR), a conformal regression method that efficiently constructs near-minimal prediction intervals with guaranteed coverage. CIR leverages black-box machine learning models to estimate outcome distributions through interquantile ranges, transforming these estimates into compact prediction intervals while achieving approximate conditional coverage. We further propose CIR+ (Conditional Interquantile Regression with More Comparison), which enhances CIR by incorporating a width-based selection rule for interquantile intervals. This refinement yields narrower prediction intervals while maintaining comparable coverage, though at the cost of slightly increased computational time. Both methods address key limitations of existing distributional conformal prediction approaches: they handle skewed distributions more effectively than Conformalized Quantile Regression, and they achieve substantially higher computational efficiency than Conformal Histogram Regression by eliminating the need for histogram construction. Extensive experiments on synthetic and real-world datasets demonstrate that our methods optimally balance predictive accuracy and computational efficiency compared to existing approaches.

Code — <https://github.com/orince/CIR>

1 Introduction

Conformal prediction is a powerful framework for constructing prediction intervals with finite-sample validity guarantees. By leveraging data exchangeability, conformal methods transform outputs from arbitrary machine learning algorithms into set-valued predictions that achieve the desired coverage level without imposing distributional assumptions on the underlying data.

Existing conformal regression methods primarily fall into two categories: directly predicting interval endpoints (Romano, Patterson, and Candès 2019; Kivaranovic, Johnson, and Leeb 2020; Sesia and Candès 2020; Gupta, Kuchibhotla, and Ramdas 2022) or inverting estimated full conditional distributions (Izbicki, Shimizu, and Stern 2020a; Chernozhukov, Wüthrich, and Zhu 2021). While effective

in many cases, these approaches may produce suboptimal intervals if the data is skewed. Conformalized quantile regression (CQR) (Romano, Patterson, and Candès 2019) may yield unbalanced intervals when the conditional distribution is skewed, whereas density-based methods can adapt to skewness but involve complex tuning and interpretation. The conformal histogram regression (CHR) method (Sesia and Romano 2021) approximates the conditional distribution of $Y|X$ using histograms and seeks the shortest intervals with the desired coverage. It extracts more information from the estimated conditional distribution compared to other methods, such as CQR (Romano, Patterson, and Candès 2019). A key advantage of CHR is its ability to automatically adapt to the potential skewness of the data distribution, which sets it apart from methods that tend to produce symmetric intervals with fixed lower and upper miscoverage rates and may be suboptimal when dealing with data of unknown skewness (Romano, Patterson, and Candès 2019; Izbicki, Shimizu, and Stern 2020b; Chernozhukov, Wüthrich, and Zhu 2021).

However, CHR faces a potential limitation in terms of computational efficiency, particularly due to the process of discretizing the response space into bins, constructing histograms, and searching through them to locate interval boundaries. To overcome these inefficiencies, we propose Conformal Interquantile Regression (CIR), a novel method that directly uses the structure of conditional quantiles to determine intervals. CIR eliminates explicit response space binning by directly estimating the conditional probability of a new response falling into each interquantile interval, using a multi-output quantile regression model with equiprobable quantiles. This design enables CIR to adapt to data skewness while providing enhanced computational efficiency.

For each calibration sample, CIR begins with the narrowest interquantile interval and iteratively incorporates the next-shortest adjacent interval until the observed response Y is covered. The number of intervals needed defines the conformity score. Intuitively, samples lying in long interquantile intervals receive larger scores, whereas those lying in shorter intervals receive smaller ones. Following the thresholding principle in (Luo and Zhou 2025b; Romano, Sesia, and Candès 2020), we select a data-dependent cutoff on these scores so that the proportion of calibration responses contained in the corresponding merged intervals attains the target coverage level. This cutoff is then applied to test inputs to form

*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

the final prediction intervals.

To optimize the prediction interval further, we also propose Conditional Interquantile Regression with More Comparison (CIR+) as an extension of CIR. CIR+ incorporates an additional decision mechanism that evaluates whether to retain or discard specific interquantile intervals based on their length. Specifically, if the length of the interquantile interval containing Y falls below a threshold determined during calibration, it is discarded. This step yields slightly narrower prediction set widths while maintaining comparable coverage performance. Our analysis shows that prediction intervals obtained by thresholding the number of interquantile intervals in CIR and CIR+ guarantee marginal coverage. Furthermore, these methods can achieve the desired conditional coverage and minimize expected prediction interval length asymptotically. Experimental results on both synthetic and real-world datasets showcase the superiority of our proposed methods. CIR and CIR+ exhibit performance comparable to the state-of-the-art CHR method while significantly reducing computational demands. This combination of statistical efficiency and computational economy positions our approaches as powerful alternatives in the field of conditional prediction intervals.

The rest of this paper is structured as follows. Section 2 discuss related work which provides context for our contributions. Section 3 introduces the proposed CIR and CIR+ methods in detail. Following this, we present a brief theoretical analysis of our methods in Section 4. Section 5 presents numerical experiments comparing our methods with existing conformal regression techniques on simulated and real datasets. Finally, we conclude the paper in Section 6.

2 Related Work

Conformal Prediction (CP) has been successfully applied to both classification (Luo and Zhou 2025f; Luo and Colombo 2024; Luo and Zhou 2025d; Wang et al. 2025; Zhang et al. 2025) and regression tasks (Luo and Zhou 2025b,e; Bao et al. 2025a), demonstrating its flexibility across diverse real-world scenarios such as segmentation (Luo and Zhou 2025a), games (Luo et al. 2024; Bao et al. 2025b), time-series forecasting (Su, Zhou, and Luo 2024), and graph-based applications (Luo, Nettasinghe, and Krishnamurthy 2023; Tang et al. 2025; Luo and Zhou 2025c; Wang, Zhou, and Luo 2025; Luo and Colombo 2025).

Our work focuses on conformal regression and aims to produce compact, interpretable prediction intervals with guaranteed coverage. While (Izbicki, Shimizu, and Stern 2022) and (Luo and Zhou 2025b) (CTI) combine profile-distance-based similarity with multi-output quantile regression to approximate conditional densities, we instead prioritize interval-based uncertainty quantification, which—as emphasized by (Sesia and Romano 2021)—offers clearer interpretability and avoids irregular, overconfident regions. CTI employs a global length threshold on interquantile segments, often yielding highly fragmented prediction sets, whereas our CIR and CIR+ procedures maintain finite-sample marginal and conditional coverage while striving to produce a single contiguous interval whenever possible. Because CTI explicitly trades coverage for narrower sets, its

behavior remains fundamentally distinct from ours even under unimodal conditional distributions and perfectly estimated quantiles. (Gao et al. 2025) construct unconditional label-space level sets that do not depend on the conditional distribution, targeting density-based highest posterior density set rather than regression intervals. However, their procedure requires repeatedly evaluating empirical coverage over all calibration-sample-induced breakpoints and solving a discrete global optimization over these candidate sets, which makes the method computational expensive.

3 Proposed Method

Problem Setup

Consider a general regression problem with a dataset $\{(x_i, y_i)\}_{i=1}^n$, where $x_i \in \mathcal{X} \subseteq \mathbb{R}^d$ is the input feature vector and $y_i \in \mathcal{Y} \subseteq \mathbb{R}$ is the corresponding continuous response variable. The dataset is split into three parts: a training set $\mathcal{D}_{\text{train}}$, a calibration set \mathcal{D}_{cal} , and a test set $\mathcal{D}_{\text{test}}$. The corresponding indices set are denoted by $\mathcal{I}_{\text{train}}$, \mathcal{I}_{cal} and $\mathcal{I}_{\text{test}}$ respectively. We assume that the examples in these sets are exchangeable. Our goal is to find a prediction interval $C(X) \subseteq \mathcal{Y}$ for each test input X such that the true response value Y is included in $C(X)$ with a probability of at least $1 - \alpha$, where $\alpha \in (0, 1)$ is a target significance level. Specifically, $C(X)$ should guarantee the finite-sample marginal coverage among all the samples in \mathcal{D}_{cal} and $\mathcal{D}_{\text{test}}$:

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

for joint distribution for X and Y .

While achieving marginal coverage, we also aim for $C(X)$ to approximately achieve conditional coverage at level $1 - \alpha$, where $\alpha \in (0, 1)$:

$$\mathbb{P}[Y \in C(x)|X = x] \geq 1 - \alpha$$

meaning that in practice the procedure should approximate this objective and, under appropriate conditions, achieve it asymptotically as the sample size tends to infinity. Lastly, the prediction intervals are expected to be as narrow as possible.

Imagine an oracle with access to $P_{Y|X}$, the distribution of Y conditional on X , which leverages such information to construct optimal prediction intervals as follows. For simplicity, suppose $P_{Y|X}$ has a continuous density $f(y | x)$ with respect to the Lebesgue measure, although this could be relaxed with more involved notation. Then, the oracle interval for $Y | X = x$ would be:

$$C_{\text{oracle}}^\alpha(x) = [\ell_{1-\alpha}(x), v_{1-\alpha}(x)], \quad (1)$$

where, for any $\tau \in (0, 1]$, $\ell_\tau(x)$ and $v_\tau(x)$ are defined as:

$$\arg \min_{(\ell, v) \in \mathbb{R}^2: v \geq \ell} \left\{ |v - \ell| : \int_\ell^v f(y | x) dy \geq \tau \right\}. \quad (2)$$

This is the shortest interval with conditional coverage. Recall that CHR (Sesia and Romano 2021) replaces f in Equation (2) with a histogram approximation, which is computation expensive. Specifically, if we partition the domain of Y into T equal parts based on quantiles and obtain the corresponding quantiles $q_0(x), \dots, q_T(x)$, where $q_t(x)$ is the

Algorithm 1: Conditional Interquantile Regression (CIR)

Input: Labeled data $\{(x_i, y_i)\}_{i \in \mathcal{I}}$, a data split ratio, unlabeled test data $\{x_i\}_{i \in \mathcal{I}_{\text{test}}}$, black-box learning algorithm \mathcal{B} , No. of interquantile intervals T , level $\alpha \in (0, 1)$.

- 1: Randomly split the indices \mathcal{I} into $\mathcal{I}_{\text{train}}$ and \mathcal{I}_{cal} .
- 2: Train \mathcal{B} on samples in $\mathcal{I}_{\text{train}}$, and obtain quantile estimation functions \hat{q}_t for $t = 0, 1, \dots, T$.
- 3: For every $i \in \mathcal{I}_{\text{cal}} \cup \mathcal{I}_{\text{test}}$, evaluate $\hat{q}_t(x_i)$ for $t = 0, 1, \dots, T$.
- 4: For any given k and $i \in \mathcal{I}_{\text{cal}}$, define l_k as the lower endpoint of the shortest interval that encompasses k consecutive interquantile intervals and the corresponding interval is defined as $C_k^\alpha(x_i) = (\hat{q}_{l_k}(x_i), \hat{q}_{l_k+k}(x_i))$.
- 5: For $i \in \mathcal{I}_{\text{cal}}$, let k_i be the smallest k such that $y_i \in C_k^\alpha(x_i)$ (equivalent to $s(x_i, y_i)$ in Eq.9).
- 6: Compute \hat{k} as the r_α smallest k_i , where $r_\alpha = \lceil (1 - \alpha)(1 + |\mathcal{I}_{\text{cal}}|) \rceil$.

Output: A prediction interval $C_{\hat{k}}^\alpha(x_i)$.

t/T -quantile of the conditional distribution of $Y|X = x$, the interval derived from following expression approximates the one from Equation(2) when T is sufficiently large:

$$\arg \min_{\substack{l=0, \dots, T-1 \\ l < u \leq T}} \left\{ |q_u(x) - q_l(x)| : \sum_{t=l}^{u-1} \mathbb{P}[Y \in (q_t(x), q_{t+1}(x))] \geq \tau \right\}. \quad (3)$$

The corresponding interval of it would be:

$$C^\alpha(x) = (q_{l_\tau}(x), q_{u_\tau}(x)). \quad (4)$$

This optimization over consecutive quantile-based bins motivates our interquantile-interval approach, which avoids explicit response-space binning.

Conditional Interquantile Regression

Our method directly selects consecutive interquantile intervals starting from the smallest ones, hence the name *conformal interquantile regression (CIR)*. Specifically, we apply quantile regression on the training set $\mathcal{D}_{\text{train}}$ to predict the t/T -th quantile of the conditional distribution $Y|X = x$ for every $x \in \mathcal{X}$, where t takes values from 0 to T in increments of 1. The estimated quantile for t/T is denoted by $\hat{q}_t(x)$ and the interquantile intervals are then defined as:

$$I_t(x) = (\hat{q}_{t-1}(x), \hat{q}_t(x)] \quad \text{for } t = 1, \dots, T. \quad (5)$$

Then Equation 3 can be approximately written as:

$$\arg \min_{\substack{l=0, \dots, T-1, \\ 1 \leq k \leq T-l-1}} \left\{ \sum_{t=l}^{l+k} |I_{t+1}(x)| : \sum_{t=l}^{l+k} \mathbb{P}[Y \in I_{t+1}(x)] \geq \tau \right\}, \quad (6)$$

where l denotes the starting index of the interval and k represents the number of consecutive interquantile intervals to

be included, both determined by τ . And the corresponding interval can be defined as

$$C_{k(\tau)}^\alpha(x) = (\hat{q}_{l(\tau)}(x), \hat{q}_{l(\tau)+k(\tau)}(x)),$$

Specifically, when $P_{Y|X}$ is unimodal, our estimated interval can be derived from

$$C_{\hat{k}}^\alpha(x) = (\hat{q}_{l_{\hat{k}}}(x), \hat{q}_{l_{\hat{k}}+\hat{k}}(x)), \quad (7)$$

where

$$\hat{k} := k(\hat{\tau}); l_{\hat{k}} := l(\hat{\tau}) = \arg \min_{l=0, \dots, T-1} \left\{ \sum_{t=l}^{l+\hat{k}} |I_{t+1}(x)| \right\}. \quad (8)$$

Notice that $l_{\hat{k}}$ represents the lower endpoint index, with the subscript \hat{k} reflecting the unimodal property that the number of included interquantile intervals uniquely determines both the narrowest interval and its corresponding lower endpoint. $\hat{\tau}$ will be determined by suitable conformity scores evaluated on the hold-out data, and it usually larger than the true target coverage τ if the model for f is not very accurate to ensure this condition is met. However, if the quantiles are accurately estimated, ensuring that each interval has approximately the same probability, $1/T$, of covering the true label Y , and if the number of interquantile intervals approaches infinity, Equation (7) will closely resemble that of the oracle in Equation (1).

The key quantity determining our prediction interval is \hat{k} , the number of interquantile intervals to include. To determine \hat{k} , we first introduce a conformity score function:

$$s(x, y) = \min \left\{ k \in 1, \dots, T : y \in C_k^\alpha(x) \right\}. \quad (9)$$

This function computes the minimum number of interquantile intervals needed to contain y . The implementation follows an iterative process. Starting with the shortest interquantile interval, for each $i \in \mathcal{I}_{\text{cal}}$, if y_i falls outside this interval, we expand our search to adjacent interquantile intervals (on either side, or one side at boundaries). We select the narrowest adjacent interval and combine it with our original interval. This process continues until y_i falls within the expanded interval. The total number of interquantile intervals in this final interval defines k_i .

For computational efficiency, we implement the conformity score as $s(x_i, y_i) = k_i$ for $i \in \mathcal{I}_{\text{cal}}$. We then define $r_\alpha = \lceil (1 - \alpha)(1 + |\mathcal{I}_{\text{cal}}|) \rceil$ and select

$$\hat{k} \leftarrow \text{the } r_\alpha\text{-th smallest } k_i, \quad (10)$$

ensuring $y_i \in C_{\hat{k}}^\alpha(x_i)$ holds for at least r_α instances in the calibration set. Finally, we use this \hat{k} in Equation (7) to obtain prediction sets for all $x_i, i \in \mathcal{I}_{\text{test}}$. The procedure is detailed in Algorithm 1.

Conditional Interquantile Regression with More Comparison (CIR+)

In CIR, the conformity score is defined as the minimum number of interquantile intervals required to contain the response y . We then select the r_α -th smallest number in the

Algorithm 2: CIR with More Comparison (CIR+)

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} , No. of interquantile intervals T .

- 1: Randomly split the indices \mathcal{I} into $\mathcal{I}_{\text{train}}$ and \mathcal{I}_{cal}
- 2: Train \mathcal{B} on samples in $\mathcal{I}_{\text{train}}$, and obtain quantile estimation functions \hat{q}_t for $t = 0, 1, \dots, T$.
- 3: For every $i \in \mathcal{I}_{\text{cal}} \cup \mathcal{I}_{\text{test}}$, evaluate $\hat{q}_t(x_i)$ for $t = 0, 1, \dots, T$.
- 4: For any given k and $i \in \mathcal{I}_{\text{cal}}$, define l_k as the lower endpoint of the shortest interval that encompasses k consecutive interquantile intervals and the corresponding interval is defined as $C_k^\alpha(x_i) = (\hat{q}_{l_k}(x_i), \hat{q}_{l_k+k}(x_i)]$.
- 5: Define $s(x_i, y_i) = k_i - 1 + e_{k_i}(x_i)$, where k_i is the smallest k such that $y_i \in C_k^\alpha(x_i)$ for $i \in \mathcal{I}_{\text{cal}}$ and $e_{k_i}(x_i) \in (0, 1)$ is a rescaled length of the k_i -th interquantile interval of x_i as defined in Equation (12).
- 6: Define $r_\alpha = \lceil (1 - \alpha)(1 + |\mathcal{I}_{\text{cal}}|) \rceil$ and then the threshold \hat{s} is defined as the r_α -th-smallest value of $s(x_i, y_i)$ for $i \in \mathcal{I}_{\text{cal}}$.
- 7: For $i \in \mathcal{I}_{\text{test}}$, $\hat{k}_i = \lfloor \hat{s} \rfloor + \mathbf{1}\{e_{\lfloor \hat{s} \rfloor}(x_i) > \hat{s} - \lfloor \hat{s} \rfloor\}$.

Output: $C_{\hat{k}_i}^\alpha(x_i)$ for y_i , where $i \in \mathcal{I}_{\text{test}}$.

calibration set as the threshold for the minimum number of interquantile intervals in the prediction interval. However, when multiple intervals in the calibration set share the same width at the r_α -th position, using only the interval count may reduce accuracy. For example, suppose 5 samples in the CIR calibration process share rank r_α , the selected \hat{k} would span ranks r_α to $r_\alpha + 4$. To obtain a more precise ranking, we incorporate additional information and modify CIR into Conditional Interquantile Regression with More Comparison (CIR+). The conformity score is defined as:

$$s(x, y) = \min \left\{ k - 1 + e_k(x) : y \in C_k^\alpha(x) \right\}, \quad (11)$$

where

$$e_k(x) = \frac{1}{c} \min \left\{ \hat{q}_{l_{k-1}}(x) - \hat{q}_{l_{k-1}-1}(x), \right. \\ \left. \hat{q}_{l_{k-1}+k}(x) - \hat{q}_{l_{k-1}+k-1}(x) \right\}, \quad (12)$$

and c can be any fixed constant which makes $e_k(x) \in (0, 1)$. In other words, when expanding our search from the shortest interquantile interval, we refine the process of incrementing k . Rather than simply increasing from $k - 1$ to k when we find an interquantile interval containing y , we add a rescaled length of the newly included interval—specifically, the adjacent interquantile interval with the smaller size. Typically, since the predicted quantiles for different samples are different, the lengths of the interquantile intervals are often different as well. This allows us to accurately select the r_α -th smallest from these conformity scores. In calibration step, we still first find k_i as in CIR. Then, we only need to focus on the lengths of the interquantile intervals containing y_i that correspond to the tied ranks at the r_α -th smallest position. We then re-rank these tied scores according to Equa-

tion (11) and define the threshold as:

$$\hat{s} \leftarrow \text{the } r_\alpha\text{-th smallest } k_i - 1 + e_{k_i}(x_i). \quad (13)$$

In the prediction phase, we identify $C_{\lfloor \hat{s} \rfloor}^\alpha(x_i)$ for each sample in the test set. Subsequently, we determine whether to include the $\lfloor \hat{s} \rfloor$ -th merged interquantile interval based on its length. Specifically, we compare the scaled length to $\hat{s} - \lfloor \hat{s} \rfloor$. If greater, we retain the segment and the prediction set becomes $C_{\lfloor \hat{s} \rfloor}^\alpha(x_i)$. If less, we discard it, leaving the set as $C_{\lfloor \hat{s} \rfloor}^\alpha(x_i)$. The procedure is shown in Algorithm 2.

4 Theoretical Results

In this section, we mainly present the theoretical results for CIR and provide additional details, including the analysis of CIR+, in the appendix. First, we show that CIR has the potential to achieve the optimal size for prediction intervals when considering the marginal distribution if we assume that our quantile regression model is sufficiently accurate. The proof of marginal coverage probability for CIR follows the same logic as the standard argument used in the general conformal prediction framework.

Theorem 4.1. (Finite-sample) *If (X_i, Y_i) , for $i \in \mathcal{I}_{\text{cal}} \cup \mathcal{I}_{\text{test}}$, are exchangeable, then for $i \in \mathcal{I}_{\text{test}}$, the output of Alg. 1 satisfies:*

$$\mathbb{P} \left[Y_i \in C_{\hat{k}_i}^\alpha(X_i) \right] \geq 1 - \alpha. \quad (14)$$

Proof. Note that the score function $s(X_i, Y_i) = \min \left\{ k \in 1, \dots, T : Y_i \in C_k^\alpha(X_i) \right\}$ for $i \in \mathcal{I}_{\text{cal}} \cup \mathcal{I}_{\text{test}}$ are also exchangeable. For any (X_i, Y_i) in the test set, the rank of $s(X_i, Y_i)$ is smaller than \hat{k} which in Alg. 1 with probability $\frac{\lfloor (1 + |\mathcal{I}_{\text{cal}}|)(1 - \alpha) \rfloor}{1 + |\mathcal{I}_{\text{cal}}|} \geq 1 - \alpha$. \square

This theoretical result also holds for CIR+ (Alg. 2) and the detailed discussion is shown in the appendix. As mentioned in Section 3, if the partition T is large enough, the output in (3)–(4) is approximate to (1)–(2). The following result shows, if the quantiles are correctly estimated and the samples are i.i.d., the \hat{k} computed by CIR (Alg. 1) are asymptotically equal to the optimal as the sample size $|\mathcal{I}| \rightarrow \infty$ (also indicating $|\mathcal{I}_{\text{cal}}| \rightarrow \infty$). Next theory relies on some assumptions:

1. (i.i.d.) The samples are i.i.d., which is stronger than exchangeability.
2. (Consistency) When the black-box model estimates $P_{Y|X}$ consistently, each interquantile interval should contain the true label Y with probability $1/T$, and the optimal number of interquantile intervals containing Y should be $\lceil T(1 - \alpha) \rceil$ to achieve $(1 - \alpha)$ coverage of true labels in the calibration set. This assumption is crucial, which similar to that in (Romano, Patterson, and Candès 2019; Sesia and Romano 2021). While difficult in practice, it has been theoretically justified for some models like in (Meinshausen 2006).
3. (Unimodality) The true $P_{Y|X}$ is unimodal.

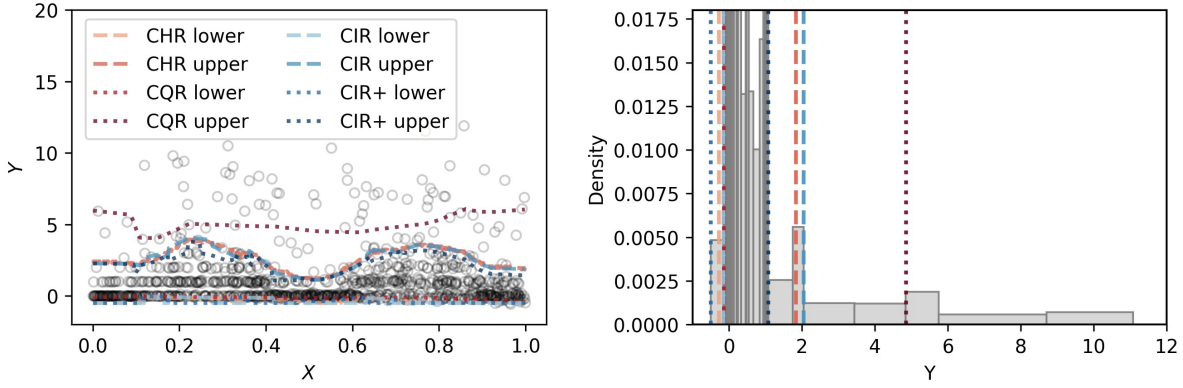


Figure 1: Comparison of CIR, CIR+, CQR (Romano, Patterson, and Candès 2019), and CHR (Sesia and Romano 2021) in a single-variable example. All methods use the same deep quantile model and guarantee 90% marginal coverage. Left: Prediction bands as a function of X . Empirical marginal and estimated conditional coverage for all methods is 0.9 (except for CQR conditional coverage: 0.7). Average lengths: CIR (3.0), CIR+ (2.7), CHR (3.0), CQR (5.2). Right: Interquantile intervals at $X \approx 0.4$, using 50 quantiles with truncated high-density regions. Narrower intervals indicate higher densities. CIR and CIR+ select narrower intervals, with CIR+ choosing one fewer due to its scaled length being less than $\hat{s} - s$.

Theorem 4.2. (Infinite-sample) Under Assumptions 1– 3, for $(X, Y) \in \mathcal{D}^{test}$, the \hat{k} obtained from Alg. 1 exhibits the following property:

$$\mathbb{P} \left[|\hat{k} - \lceil T(1 - \alpha) \rceil | > \epsilon_n \right] \leq \eta_n, \quad (15)$$

$$\mathbb{P} \left[\mathbb{P} \left[Y \in C_{\hat{k}}^\alpha(X) \mid X \right] \geq 1 - \alpha - \gamma_n \right] \geq 1 - \zeta_n \quad (16)$$

where we assume the calibration size $|\mathcal{I}_{cal}| = n \rightarrow \infty$ and $\epsilon_n \rightarrow 0, \eta_n \rightarrow 0, \gamma_n \rightarrow 0, \zeta_n \rightarrow 0$.

Intuitively, as T large enough, the additional selection step in Alg. 2 becomes negligible, causing CIR+ to converge theoretically to CIR. The details of proof are shown in the appendix.

Remark: The unimodality setting is not necessary for implementing our methods or achieving marginal coverage in Thm. 4.1 but only for properties of optimal interval length and conditional coverage in Thm. 4.2. In non-unimodal cases, the optimal prediction set may consist of multiple disjoint intervals rather than a single connected interval (guaranteed in the unimodal setting). To address this, a more effective approach involves ordering the interquantile intervals by length and making selections from this ordered set. Nevertheless, empirical analysis in later sections and the appendix shows both methods maintain robust performance even on non-unimodal distributions, demonstrating their practical versatility.

5 Experiments

Our experimental implementation is based on the publicly available code for Conformal Histogram Regression (CHR) by (Sesia and Romano 2021), which can be accessed at <https://github.com/msesia/chr>. We are deeply grateful for their work, serving as a significant source of

$ \mathcal{I} $		CHR	CQR	CIR	CIR+
500	Marginal	0.902 (0.001)	0.901 (0.001)	0.900 (0.002)	0.897 (0.002)
	Condit.	0.881 (0.003)	0.879 (0.004)	0.878 (0.004)	0.871 (0.004)
	Width	3.670 (0.039)	5.165 (0.043)	3.930 (0.057)	3.790 (0.053)
	Time	36.597 (0.396)	0.008 (0.000)	0.072 (0.001)	0.097 (0.001)
3000	Marginal	0.901 (0.001)	0.900 (0.001)	0.905 (0.001)	0.900 (0.001)
	Condit.	0.887 (0.003)	0.875 (0.003)	0.895 (0.003)	0.886 (0.003)
	Width	3.425 (0.017)	5.320 (0.019)	3.764 (0.028)	3.632 (0.028)
	Time	73.735 (0.785)	0.014 (0.000)	0.404 (0.005)	0.526 (0.007)
5000	Marginal	0.901 (0.001)	0.900 (0.001)	0.905 (0.001)	0.900 (0.001)
	Condit.	0.889 (0.003)	0.875 (0.003)	0.895 (0.003)	0.889 (0.003)
	Width	3.389 (0.013)	5.345 (0.016)	3.776 (0.022)	3.643 (0.022)
	Time	103.637 (1.105)	0.019 (0.000)	0.663 (0.008)	0.862 (0.010)

Table 1: Performance comparison of our methods (CIR and CIR+) and benchmarks on synthetic data for different sample size ($|\mathcal{I}|$), with skewness equals to 2.7. The corresponding standard deviations are shown in parentheses.

inspiration for our work. Our numerical experiments were conducted on Intel Xeon 2.10GHz CPUs in a computing cluster and each data set was analyzed using a single core.

Synthetic Data Analysis

We simulate a synthetic dataset with a one-dimensional feature X and a continuous response Y following (Romano, Patterson, and Candès 2019; Sesia and Romano 2021) and attach the details in the appendix. Figure 1 illustrates the application of our methods to this toy data, comparing CIR and CIR+ with conformalized histogram regression (CHR) (Sesia and Romano 2021) and conformalized quantile regression (CQR) (Romano, Patterson, and Candès 2019). The

Skew		CHR	CQR	CIR	CIR+
0.66	Marginal	0.901 (0.001)	0.900 (0.001)	0.906 (0.001)	0.902 (0.001)
	Condit.	0.891 (0.005)	0.883 (0.005)	0.898 (0.004)	0.889 (0.005)
	Width	5.627 (0.046)	5.950 (0.031)	5.900 (0.049)	5.715 (0.054)
	Time	106.995 (2.049)	0.018 (0.000)	0.710 (0.015)	0.911 (0.019)
1.71	Marginal	0.901 (0.001)	0.900 (0.001)	0.907 (0.001)	0.902 (0.001)
	Condit.	0.891 (0.003)	0.888 (0.003)	0.897 (0.003)	0.891 (0.003)
	Width	4.785 (0.025)	5.436 (0.018)	5.110 (0.026)	4.870 (0.030)
	Time	103.026 (1.150)	0.020 (0.000)	0.687 (0.009)	0.887 (0.011)
2.70	Marginal	0.901 (0.001)	0.900 (0.001)	0.905 (0.001)	0.900 (0.001)
	Condit.	0.889 (0.003)	0.875 (0.003)	0.895 (0.003)	0.889 (0.003)
	Width	3.389 (0.013)	5.345 (0.016)	3.776 (0.022)	3.643 (0.022)
	Time	103.637 (1.105)	0.019 (0.000)	0.663 (0.008)	0.862 (0.010)

Table 2: Performance comparison of our methods (CIR and CIR+) and benchmarks on synthetic data for different skewness of the conditional distribution of the response (Skew), with a sample size of 5000.

left-hand side figure visualizes the resulting prediction bands for independent test data, comparing the analogous outputs across all methods. Both CIR and CIR+ are based on the same deep neural network and guarantee 90% marginal coverage. CHR is capable of extracting information from all conditional quantiles estimated by the base model and automatically adapting to the estimated data distribution and produces relatively narrow intervals. Although our approaches omit the histogram step, it does not significantly impact our results. Both CIR and CIR+ yield intervals as narrow as CHR. In contrast, CQR can only leverage pre-specified lower and upper quantiles (e.g., 5% and 95%) and is therefore not adaptive to skewness. The right-hand side figure displays the interquantile intervals and their corresponding probabilities for $X \approx 0.4$. To elucidate the internal processes of the CIR and CIR+ algorithms, we have set only 50 quantiles for a tidier presentation. Additionally, we have truncated the regions with high density and focused on the areas where the information is concentrated. Note that the narrower interquantile intervals have higher densities. We can observe that both CIR and CIR+ select relatively narrower interquantile intervals at each step. In this sample, CIR+ selects one fewer interval compared to CIR because the scaled length of that interval is smaller than $\hat{s} - s$, leading to its exclusion. We simulate a synthetic dataset with a one-dimensional feature X and a continuous response Y , drawn from the distribution illustrated in Figure 1. Our methods are applied to 2,000 independent observations from this distribution, utilizing the first 1,000 for training a deep quantile regression model and the remainder for calibration. Aligning with (Sesia and Romano 2021), we set 100 bins for CHR.

Table 1 shows CIR and CIR+ performance across sample sizes ($|\mathcal{I}|$) over 1000 independent experiments, with $|\mathcal{I}|$ split evenly between training and calibration and the number of interquantile intervals fixed at $T = 100$. We evaluate marginal coverage, worst-slab conditional coverage as

in (Cauchois, Gupta, and Duchi 2021; Romano, Sesia, and Candès 2020), average interval width, and computational cost per sample for calibration and prediction, with all methods using the same base model. Both our methods and CHR achieve narrower intervals than CQR while maintaining marginal and conditional coverage. At smaller sample sizes, our method shows fluctuations due to imprecise quantile estimation from limited data, but as samples increase, quantile regression becomes more accurate and results stabilize. CIR yields higher marginal coverage than the target level due to conformity score ties. In the ideal scenario, CIR+ should closely approach the coverage level when N is large while CIR should show an additional approximately $1/(2T) = 0.005$ than coverage level $1 - \alpha = 0.9$ (As each interquantile interval carries $1/T$ probability, the threshold could occur anywhere, but on average falls at the midpoint). Notably, CHR is computationally intensive, with its runtime increasing as the sample size grows, which is primarily due to the increasingly time-consuming histogram construction process. As our confidence regions are computed directly from interquantile intervals and do not involve a binning process, the computational time required for calibrating and predicting each sample is negligible.

Table 2 shows results with 5000 samples under varying distributional skewness. Following (Sesia and Romano 2021), we use a biased coin flip to invert Y to $-Y$ for each data point, with coin bias as control. This generates distributions from skewed (Table 1) to symmetric $P_{Y|X}$. Results are shown against expected skewness $\mathbb{E}[(Y - \mu(X))^3 / \sigma^3(X)]$, where $\mu(X)$ and $\sigma(X)$ are mean and standard deviation of $Y | X$. For symmetric $P_{Y|X}$ (skewness near 0), all methods produce similar interval lengths. As skewness increases, our methods match CHR’s performance across different asymmetry levels. While CHR yields narrowest intervals throughout, its high computational cost persists. CIR and CIR+ maintain efficiency regardless of skewness, balancing performance with practicality. The appendix provides extended analyses, including comparisons with additional benchmarks such as distributional conformal prediction (DCP) (Chernozhukov et al., 2019) and DistSplit (Izbicki et al., 2020), along with evaluations across broader ranges of sample sizes, skewness levels, and results using random forest models.

Real Data Analysis

We evaluate CIR and CIR+ on seven public-domain datasets previously analyzed in (Romano, Patterson, and Candès 2019): physicochemical properties of protein tertiary structure (bio) (Rana 2019), blog feedback (blog) (Buza 2014), Facebook comment volume variants one (fb1) and two (fb2) (Singh 2016), and Medical Expenditure Panel Survey numbers 19 (meps19), 20 (meps20), and 21 (meps21) (AfHRAQ 2021). The first four datasets are from the UCI Machine Learning Repository (Dua and Graff 2017), while the MEPS datasets are from (Cohen, Cohen, and Banthin 2009). For detailed descriptions of these datasets, we refer readers to (Romano, Patterson, and Candès 2019). We compare our methods against CHR, CQR, DistSplit, and DCP-CQR. The latter, proposed in (Sesia and Romano 2021), is designed

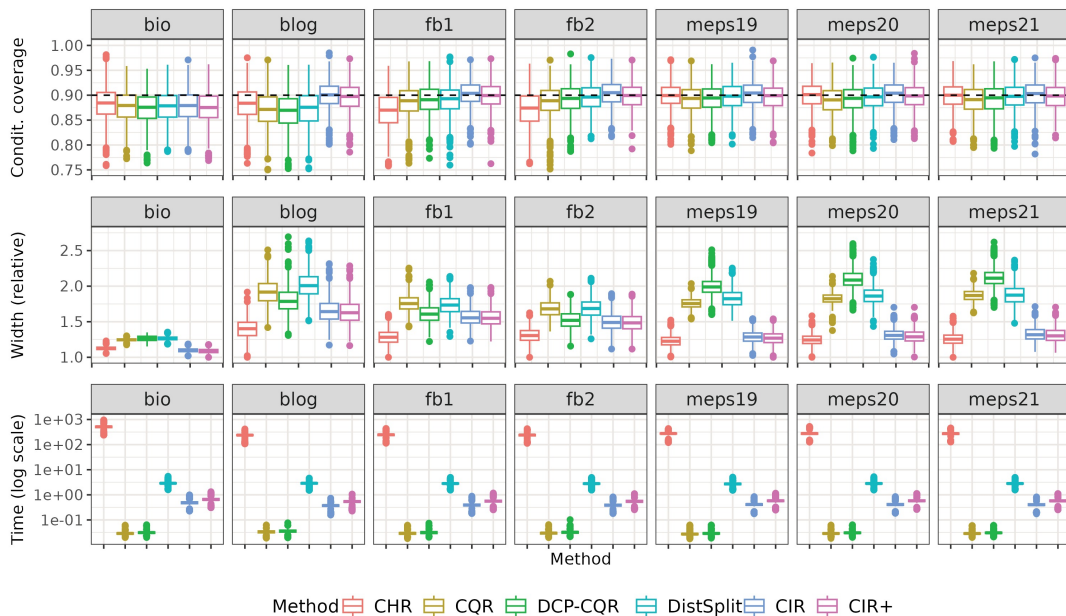


Figure 2: Performance of our method (CIR) compared to that of naive uncalibrated prediction intervals based on the same deep neural network regression model. Note that the top part of this plot shows marginal coverage.

to enhance the stability of DCP (Chernozhukov, Wüthrich, and Zhu 2019) by integrating it with CQR (Romano, Patterson, and Candès 2019). All methods employ an identical deep quantile regression model for prediction. For comparison, we also present results using Random Forest method in the appendix. The Forests method (Meinshausen 2006) aggregates quantile estimates across an ensemble of decision trees, while the Neural Network approach (Taylor 2000) outputs multiple quantiles as a vector through shared parameters in a single network. Both methods are inherently designed to avoid the problem of crossing quantiles. Aligning with (Sesia and Romano 2021), we set 1000 bins for CHR for this part. We assess their performance following the protocol described in the previous section, averaging results over 100 independent experiments per dataset. Each experiment utilizes 2000 samples for training, 2000 for calibration, and the remaining samples for testing. All features are standardized to zero mean and unit variance. The nominal coverage rate is set at 90%.

Figure 2 compares conditional coverage and interval widths across different methods and datasets. For meaningful cross-dataset comparison, prediction interval widths are normalized within each dataset, setting the smallest width to one. Computational times are presented on a logarithmic scale after similar normalization. The results show all methods achieve satisfactory conditional coverage. While CHR and CIR+ alternately produce the most efficient (narrowest) prediction intervals across datasets, CIR+ and CIR exhibit substantially reduced computational overhead compared to CHR. This positions CIR+ as the optimal choice, effectively balancing computational efficiency with predictive performance. CQR and DistSplit demonstrate comparable per-

formance levels, whereas DCP-CQR tends to produce wider prediction intervals in some cases. A comprehensive analysis, including marginal coverage (theoretically guaranteed for all methods) is provided in the appendix.

6 Conclusion

In this paper, we introduce two novel conformal regression methods: Conformal Interquantile Regression (CIR) and its enhanced variant, CIR+. These methods are designed to construct prediction intervals that offer both guaranteed coverage and computational efficiency, effectively addressing key limitations of existing approaches. Specifically, our methods overcome the challenges of suboptimal performance on skewed data while avoiding the computational burden. Our theoretical analysis demonstrates that both CIR and CIR+ not only guarantee marginal coverage but also achieve desired conditional coverage under appropriate conditions, while minimizing expected prediction interval length. Through extensive testing on synthetic and real datasets, we demonstrate that these methods match the performance of leading alternatives with substantially reduced computational cost. The two methods offer different solutions to the coverage-efficiency balance: CIR produces wider, more conservative intervals, while CIR+ generates narrower intervals by analyzing endpoint interquantile widths at a slight computational premium. The performance of CIR and CIR+ can be affected with small datasets due to challenges in training quantile regression models effectively. Our future work will focus on extending these frameworks to time series data, preserving its theoretical guarantees while expanding its applications.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant 62506315) and City University of Hong Kong (Grants 9610639, 7020161).

References

- AfHRAQ. 2021. Medical expenditure panel survey. https://meps.ahrq.gov/mepsweb/data_stats/data_overview.jsp.
- Bao, J.; Colombo, N.; Manokhin, V.; Cao, S.; and Luo, R. 2025a. A Review and Comparative Analysis of Univariate Conformal Regression Methods. In *Fourteenth Symposium on Conformal and Probabilistic Prediction with Applications (COPA 2025)*, 282–304. PMLR.
- Bao, J.; Dang, C.; Luo, R.; Zhang, H.; and Zhou, Z. 2025b. Enhancing Adversarial Robustness with Conformal Prediction: A Framework for Guaranteed Model Reliability. In *Proceedings of the Forty-second International Conference on Machine Learning (ICML)*.
- Buza, K. 2014. BlogFeedback. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C58S3F>.
- Cauchois, M.; Gupta, S.; and Duchi, J. C. 2021. Knowing what You Know: valid and validated confidence sets in multiclass and multilabel prediction. *Journal of Machine Learning Research*, 22(81): 1–42.
- Chernozhukov, V.; Wüthrich, K.; and Zhu, Y. 2019. Distributional conformal prediction. *preprint at arXiv:1909.07889*.
- Chernozhukov, V.; Wüthrich, K.; and Zhu, Y. 2021. Distributional conformal prediction. *Proceedings of the National Academy of Sciences*, 118(48): e2107794118.
- Cohen, J. W.; Cohen, S. B.; and Banthin, J. S. 2009. The medical expenditure panel survey: a national information resource to support healthcare cost research and inform policy and practice. *Medical care*, S44–S50.
- Dua, D.; and Graff, C. 2017. UCI Machine Learning Repository.
- Gao, C.; Shan, L.; Srinivas, V.; and Vijayaraghavan, A. 2025. Volume Optimality in Conformal Prediction with Structured Prediction Sets. *arXiv preprint arXiv:2502.16658*.
- Gupta, C.; Kuchibhotla, A. K.; and Ramdas, A. 2022. Nested conformal prediction and quantile out-of-bag ensemble methods. *Pattern Recognition*, 127: 108496.
- 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.
- Luo, R.; Bao, J.; Zhou, Z.; and Dang, C. 2024. Game-Theoretic Defenses for Robust Conformal Prediction Against Adversarial Attacks in Medical Imaging. *arXiv preprint arXiv:2411.04376*.
- Luo, R.; and Colombo, N. 2024. Entropy Reweighted Conformal Classification. In *The 13th Symposium on Conformal and Probabilistic Prediction with Applications*, 264–276. PMLR.
- Luo, R.; and Colombo, N. 2025. Conformal load prediction with transductive graph autoencoders. *Machine Learning*, 114(3): 1–22.
- Luo, R.; Nettasinghe, B.; and Krishnamurthy, V. 2023. Anomalous edge detection in edge exchangeable social network models. In *Conformal and Probabilistic Prediction with Applications*, 287–310. PMLR.
- Luo, R.; and Zhou, Z. 2025a. Conditional Conformal Risk Adaptation. *arXiv preprint arXiv:2504.07611*.
- Luo, R.; and Zhou, Z. 2025b. Conformal thresholded intervals for efficient regression. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(18): 19216–19223.
- Luo, R.; and Zhou, Z. 2025c. Conformalized interval arithmetic with symmetric calibration. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(18): 19207–19215.
- Luo, R.; and Zhou, Z. 2025d. Conformity score averaging for classification. In *Proceedings of the Forty-second International Conference on Machine Learning (ICML)*.
- Luo, R.; and Zhou, Z. 2025e. Density-Sorted Prediction Set: Efficient Conformal Prediction for Multi-Target Regression. *Pattern Recognition*, 112513.
- Luo, R.; and Zhou, Z. 2025f. Reliable classification through rank-based conformal prediction sets. *Pattern Recognition*, 112330.
- Meinshausen, N. 2006. Quantile regression forests. *Journal of Machine Learning Research*, 7: 983–999.
- Rana, P. 2019. Physicochemical properties of protein tertiary structure data set. <https://archive.ics.uci.edu/ml/datasets/Physicochemical+Properties+of+Protein+Tertiary+Structure>. Accessed: July, 2019.
- Romano, Y.; Patterson, E.; and Candès, E. 2019. Conformalized quantile regression. In *Advances in Neural Information Processing Systems*, 3538–3548.
- Romano, Y.; Sesia, M.; and Candès, E. 2020. Classification with valid and adaptive coverage. *Advances in Neural Information Processing Systems*, 33: 3581–3591.
- 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.
- Singh, K. 2016. Facebook Comment Volume Dataset. UCI Machine Learning Repository. DOI: 10.24432/C5Q886.

- Su, X.; Zhou, Z.; and Luo, R. 2024. Adaptive Conformal Inference by Particle Filtering under Hidden Markov Models. *arXiv preprint arXiv:2411.01558*.
- Tang, L.; Luo, R.; Zhou, Z.; and Colombo, N. 2025. Enhanced route planning with calibrated uncertainty set. *Machine Learning*, 114(5): 1–16.
- 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.
- Wang, T.; Zhang, K.; Luo, R.; Zhou, Z.; and Li, W. J. 2025. Trustworthy Drug-Drug Interaction Prediction Model via Weighted Aggregating Conformal Prediction. In *The 23rd Asia Pacific Bioinformatics Conference*.
- Wang, T.; Zhou, Z.; and Luo, R. 2025. Enhancing Trustworthiness of Graph Neural Networks with Rank-Based Conformal Training. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(20): 21261–21268.
- Zhang, Z.; Bao, J.; Zhou, Z.; Cheng, L.; Luo, R.; et al. 2025. Residual Reweighted Conformal Prediction for Graph Neural Networks. In *The 41st Conference on Uncertainty in Artificial Intelligence*.