

Training-conditional coverage for distribution-free predictive inference

Michael Bian and Rina Foygel Barber

Department of Statistics,
University of Chicago
e-mail: mybian99@gmail.com; rina@uchicago.edu

Abstract: The field of distribution-free predictive inference provides tools for provably valid prediction without any assumptions on the distribution of the data, which can be paired with any regression algorithm to provide accurate and reliable predictive intervals. The guarantees provided by these methods are typically marginal, meaning that predictive accuracy holds on average over both the training data set and the test point that is queried. However, it may be preferable to obtain a stronger guarantee of training-conditional coverage, which would ensure that most draws of the training data set result in accurate predictive accuracy on future test points. This property is known to hold for the split conformal prediction method. In this work, we examine the training-conditional coverage properties of several other distribution-free predictive inference methods, and find that training-conditional coverage is achieved by some methods but is impossible to guarantee without further assumptions for others.

MSC2020 subject classifications: Primary 62G08, 62G15, 62F40.

Keywords and phrases: Conformal prediction, jackknife+, distribution-free inference.

Received June 2022.

1. Introduction

Distribution-free predictive inference provides a set of methods for constructing predictive confidence intervals with minimal assumptions about the underlying distribution. Specifically, in the case of regression, suppose we are given i.i.d. training points $(X_i, Y_i) \in \mathcal{X} \times \mathbb{R}$, $i = 1, \dots, n$, and a regression algorithm \mathcal{A} that maps training points to prediction rules $\hat{\mu} : \mathcal{X} \rightarrow \mathbb{R}$. Given a new feature vector X_{n+1} , we would like to predict the unseen response Y_{n+1} . A predictive interval \hat{C}_n , trained on these n training data points using this regression algorithm \mathcal{A} , returns an interval (or more generally, a subset) $\hat{C}_n(X_{n+1}) \subseteq \mathbb{R}$, with the goal that $\hat{C}_n(X_{n+1})$ should contain the response value Y_{n+1} for this test point. We say that \hat{C}_n is a *distribution-free predictive interval* if, for every distribution P on $\mathcal{X} \times \mathbb{R}$ it holds that

$$\mathbb{P}_{P^{n+1}} \left\{ Y_{n+1} \in \hat{C}_n(X_{n+1}) \right\} \geq 1 - \alpha. \quad (1)$$

Here the notation $\mathbb{P}_{P^{n+1}} \{ \cdot \}$ denotes that the probability is computed with respect to $(X_1, Y_1), \dots, (X_n, Y_n), (X_{n+1}, Y_{n+1}) \stackrel{\text{iid}}{\sim} P$.

In practice, we are often interested in the coverage rate for test points once we fit a regression algorithm to a particular training set. However, the guarantee in (1) does not directly address this. Rather, it bounds the miscoverage rate *on average* over possible sets of training data and test points. As a result, if there is high variability in the coverage rate as a function of the training data, the test coverage rate may be substantially below $1 - \alpha$ for a particular training set. In this case, while (1) is satisfied on average, after fitting on the realized draw of the training set distribution the practitioner may be left with prediction intervals which drastically undercover.

To formalize this intuition, let $\mathcal{D}_n = ((X_1, Y_1), \dots, (X_n, Y_n))$ be the training data set. Then, define the miscoverage rate as a function of the training data:

$$\alpha_P(\mathcal{D}_n) = \mathbb{P}_P \left\{ Y_{n+1} \notin \widehat{C}_n(X_{n+1}) \mid \mathcal{D}_n \right\},$$

where the probability is now only with respect to the test point (X_{n+1}, Y_{n+1}) drawn from P . Then, the guarantee in (1) can be re-written as

$$\mathbb{E}_{P^n} [\alpha_P(\mathcal{D}_n)] \leq \alpha,$$

where the expectation is with respect to the training data $\mathcal{D}_n \sim P^n$ (and, in order to be distribution-free, this bound is again required to hold for every distribution P on $\mathcal{X} \times \mathbb{R}$).

While this expectation is bounded, $\alpha_P(\mathcal{D}_n)$ may have high variance over the training data. In particular, we can consider a worst-case scenario where

$$\mathbb{P}_{P^n} \{ \alpha_P(\mathcal{D}_n) = 1 \} = \alpha, \quad \mathbb{P}_{P^n} \{ \alpha_P(\mathcal{D}_n) = 0 \} = 1 - \alpha, \tag{2}$$

which trivially satisfies the marginal coverage guarantee (1) since $\mathbb{E}_{P^n} [\alpha_P(\mathcal{D}_n)] = \alpha$. In other words, in this worst-case scenario, a nonnegligible proportion of training sets might result in 0% training-conditional coverage even though the average coverage is still $1 - \alpha$, which may be highly problematic in practice. On the other hand, if we instead had $\alpha_P(\mathcal{D}_n) \approx \alpha$ with high probability over $\mathcal{D}_n \sim P^n$, this would be ideal, since it ensures that for nearly every possible draw of the training data, the resulting coverage over future test points should be $\approx 1 - \alpha$.

The variability of training-conditional miscoverage level $\alpha_P(\mathcal{D}_n)$ will in general depend on the distribution P , the regression algorithm \mathcal{A} , and the particular distribution-free method that is used to generate $\widehat{C}_n(X_{n+1})$. In this paper, we examine the variability of coverage for popular distribution-free methods for arbitrary P and \mathcal{A} . In particular, we seek to provide guarantees of the form

$$\mathbb{P}_{P^n} \{ \alpha_P(\mathcal{D}_n) > \alpha + o(1) \} \leq o(1), \tag{3}$$

also known as a ‘‘Probably Approximately Correct’’ (PAC) predictive interval. This type of guarantee turns out to be possible for some distribution-free methods and impossible for others. In addition, we demonstrate empirically that methods without guarantees of this form can exhibit highly variable training-conditional miscoverage rates in low stability regimes.

2. Background

In this section, we will briefly review four related methods for distribution-free predictive inference, to introduce the methods that we will study in this work. Consider an algorithm \mathcal{A} that maps datasets (consisting of (X, Y) pairs, with features $X \in \mathcal{X}$ and a real-valued response $Y \in \mathbb{R}$), to fitted regression functions $\hat{\mu} : \mathcal{X} \rightarrow \mathbb{R}$.

For a new data point whose features $X_{n+1} \in \mathcal{X}$ are observed, we would like to predict the unseen response $Y_{n+1} \in \mathbb{R}$. Given a model $\hat{\mu}$ obtained by training some algorithm \mathcal{A} on the available training data $(X_1, Y_1), \dots, (X_n, Y_n)$, can we construct a prediction interval for Y_{n+1} around the estimate $\hat{\mu}(X_{n+1})$? In many practical settings, the distribution of the data is likely unknown, and the regression algorithm \mathcal{A} may be a complex “black box” methods whose theoretical properties are not well understood, and therefore it may be challenging to guarantee a particular error bound for $\hat{\mu}(X_{n+1})$ as an estimator of the unseen response Y_{n+1} .

2.1. Distribution-free methods

2.1.1. Conformal prediction

The conformal prediction framework [19], which includes the full and split conformal methods (also called “transductive” and “inductive” conformal, respectively), provides a mechanism for constructing prediction intervals in this challenging setting, with distribution-free coverage guarantees. (See also [9] for additional background on these methods.)

To run the split conformal method, we first partition the n available labeled data points into a training set of size n_0 and a holdout set of size $n_1 = n - n_0$. After running the regression algorithm on the training data to obtain the fitted model $\hat{\mu}_{n_0} = \mathcal{A}((X_1, Y_1), \dots, (X_{n_0}, Y_{n_0}))$, the prediction interval is defined as

$$\hat{C}_n(X_{n+1}) = \hat{\mu}_{n_0}(X_{n+1}) \pm \hat{Q}_{n_1}, \quad (4)$$

where \hat{Q}_{n_1} is defined as the $\lceil (1 - \alpha)(n_1 + 1) \rceil$ -th smallest value of the holdout residuals $|Y_{n_0+1} - \hat{\mu}_{n_0}(X_{n_0+1})|, \dots, |Y_n - \hat{\mu}_{n_0}(X_n)|$. This method satisfies the marginal distribution-free predictive coverage guarantee (1) [19].

While split conformal offers both computational efficiency and distribution-free coverage, its precision may suffer from the loss of sample size incurred by splitting the data set. In contrast, full conformal uses all the available training data for model fitting, but comes at a high computational cost. Specifically, for every $y \in \mathbb{R}$, define $\hat{\mu}_{n+1}^y = \mathcal{A}((X_1, Y_1), \dots, (X_n, Y_n), (X_{n+1}, y))$, the fitted model obtained by running algorithm \mathcal{A} on the training data together with the hypothesized test point (X_{n+1}, y) . Then construct the prediction set

$$\hat{C}_n(X_{n+1}) = \left\{ y \in \mathbb{R} : |y - \hat{\mu}_{n+1}^y(X_{n+1})| \leq \hat{Q}_{n+1}^y \right\}, \quad (5)$$

where \widehat{Q}_{n+1}^y is defined as the $\lceil(1 - \alpha)(n + 1)\rceil$ -th smallest value of the residuals $|Y_1 - \widehat{\mu}_{n+1}^y(X_1)|, \dots, |Y_n - \widehat{\mu}_{n+1}^y(X_n)|, |y - \widehat{\mu}_{n+1}^y(X_{n+1})|$. Full conformal prediction also offers the distribution-free coverage guarantee (1) [19], under one additional assumption—the algorithm \mathcal{A} needs to be *symmetric* in the training data points, meaning that for any $m \geq 1$, any permutation σ on $[m] := \{1, \dots, m\}$, and any data points $(x_1, y_1), \dots, (x_m, y_m) \in \mathcal{X} \times \mathbb{R}$,

$$\mathcal{A}((x_1, y_1), \dots, (x_m, y_m)) = \mathcal{A}((x_{\sigma(1)}, y_{\sigma(1)}), \dots, (x_{\sigma(m)}, y_{\sigma(m)})). \tag{6}$$

Full conformal prediction is generally more statistically efficient than split conformal (i.e., will provide narrower prediction intervals) since we do not need to split the training data. On the other hand, the computational cost is high—aside from special cases (e.g., choosing \mathcal{A} to be the Lasso [8]), the prediction interval can only be calculated by running the regression algorithm \mathcal{A} for every possible $y \in \mathbb{R}$, or in practice, for a very fine grid of y values (theoretical guarantees for this discretized setting can also be obtained, as shown by [4]).

2.1.2. Jackknife+ and CV+

The jackknife+ and CV+ methods proposed by [1] offer a compromise between the computational efficiency of split conformal and the statistical efficiency of full conformal. These methods, which are closely related to the cross-conformal procedure of [18, 20], use a cross-validation type approach.

For jackknife+, let $\widehat{\mu}_{[n] \setminus \{i\}}$ denote the model fitted to the training data with data point i removed,

$$\widehat{\mu}_{[n] \setminus \{i\}} = \mathcal{A}((X_1, Y_1), \dots, (X_{i-1}, Y_{i-1}), (X_{i+1}, Y_{i+1}), \dots, (X_n, Y_n)).$$

Then the jackknife+ prediction interval is defined as

$$\widehat{C}_n(X_{n+1}) = \left[\begin{array}{l} \text{the } \lceil(1 - \alpha)(n + 1)\rceil\text{-th largest of } \{\widehat{\mu}_{[n] \setminus \{i\}}(X_{n+1}) - R_i\}_{i \in [n]}, \\ \text{the } \lceil(1 - \alpha)(n + 1)\rceil\text{-th smallest of } \{\widehat{\mu}_{[n] \setminus \{i\}}(X_{n+1}) + R_i\}_{i \in [n]} \end{array} \right], \tag{7}$$

where $R_i = |Y_i - \widehat{\mu}_{[n] \setminus \{i\}}(X_i)|$ for $i \in [n] := \{1, \dots, n\}$. The jackknife+ method offers a weaker distribution-free coverage guarantee [1, Theorem 1]: for every distribution P on $\mathcal{X} \times \mathbb{R}$, assuming \mathcal{A} is symmetric as in (6),

$$\mathbb{P}_{P^{n+1}} \left\{ Y_{n+1} \in \widehat{C}_n(X_{n+1}) \right\} \geq 1 - 2\alpha. \tag{8}$$

Note that, in this theoretical guarantee, noncoverage may be as high as 2α , rather than the target level α . However, empirically the method typically achieves coverage at level $1 - \alpha$, and indeed, under algorithmic stability assumptions, e.g.,

$$\mathbb{P}_{P^{n+1}} \left\{ \left| \widehat{\mu}_n(X_{n+1}) - \widehat{\mu}_{[n] \setminus \{i\}}(X_{n+1}) \right| \leq \epsilon \right\} \geq 1 - \nu, \tag{9}$$

the predictive coverage guarantee can be improved to $1 - \alpha - o(1)$ [1, Theorem 5].

While jackknife+ requires only n many calls to the regression algorithm \mathcal{A} (in contrast to full conformal, which in theory requires infinitely many calls), for a large sample size n this computational cost may still be too high. CV+ extends the jackknife+ method to K -fold cross-validation (where we can view jackknife+ as n -fold cross-validation, i.e., $K = n$). Let $[n] = S_1 \cup \dots \cup S_K$ be a partition of the training data into K subsets of size n/K , and write $\hat{\mu}_{[n] \setminus S_k}$ as the fitted model when the k -th fold S_k is removed from the n training data points. The CV+ prediction interval is defined as

$$\hat{C}_n(X_{n+1}) = \left[\begin{array}{l} \text{the } \lceil (1 - \alpha)(n + 1) \rceil\text{-th largest of } \{\hat{\mu}_{[n] \setminus S_{k(i)}}(X_{n+1}) - R_i\}_{i \in [n]}, \\ \text{the } \lceil (1 - \alpha)(n + 1) \rceil\text{-th smallest of } \{\hat{\mu}_{[n] \setminus S_{k(i)}}(X_{n+1}) + R_i\}_{i \in [n]}, \end{array} \right], \quad (10)$$

where now $R_i = |Y_i - \hat{\mu}_{[n] \setminus S_{k(i)}}(X_i)|$ for $i \in [n]$, and where $k(i)$ denotes the fold to which data point i belongs, i.e., $i \in S_{k(i)}$. The CV+ method's coverage guarantee is given by [1, Theorem 4] (see also [20] for a partial version of this result): for every distribution P on $\mathcal{X} \times \mathbb{R}$, assuming \mathcal{A} is symmetric as in (6),

$$\mathbb{P}_{P^{n+1}} \left\{ Y_{n+1} \in \hat{C}_n(X_{n+1}) \right\} \geq 1 - 2\alpha - \sqrt{2/n}. \quad (11)$$

As for jackknife+, the CV+ method typically achieves coverage near or above the target level $1 - \alpha$ in practice.

2.1.3. A note on randomized algorithms

The background given above implicitly treats the algorithm \mathcal{A} as a *deterministic* function of the training data—that is, we view \mathcal{A} as a function $((X_1, Y_1), \dots, (X_n, Y_n)) \mapsto \hat{\mu}$. In many settings, however, it is common to use a *randomized* regression algorithm—for instance, stochastic gradient descent. In this setting, we can formally view \mathcal{A} as a function $((X_1, Y_1), \dots, (X_n, Y_n), \xi) \mapsto \hat{\mu}$, where the term ξ introduces stochastic noise (effectively, a random seed). All the results described above hold for both the deterministic and randomized settings. (For results that assume \mathcal{A} is symmetric, the symmetry condition (6) should be understood in the distributional sense—that is, the training data points are treated symmetrically with respect to the randomized training procedure. For example, for stochastic gradient descent, if data points are drawn uniformly at random during the training epochs, then symmetry is satisfied.)

2.2. Marginal or conditional validity

The predictive coverage bound (1) achieved by split and full conformal, or the bounds (8) and (11) for the jackknife+ and CV+ methods, are all *marginal*

guarantees. This means that the probability is calculated over a random draw of both the training and test data. However, this may be unsatisfactory for practical purposes, in several ways.

Training-conditional coverage First, as discussed in Section 1 above, we may be interested in *training-conditional coverage*, which ensures that the predictive coverage guarantees hold (at least approximately) even after conditioning on the training data set $\mathcal{D}_n = ((X_1, Y_1), \dots, (X_n, Y_n))$. For the split conformal method described in (4), [17, Proposition 2a] establishes training-conditional coverage through a Hoeffding bound:

Theorem 1 ([17, Proposition 2a]). *Consider the split conformal method defined in (4) with sample size $n = n_0 + n_1$, where $n_0 \geq 1$ many data points are used for training the fitted model $\hat{\mu}_{n_0}$ (with an arbitrary algorithm) while the remaining $n_1 \geq 1$ data points are used as the holdout set. Then, for any distribution P and any $\delta \in (0, 0.5]$,*

$$\mathbb{P}_{P^n} \left\{ \alpha_P(\mathcal{D}_n) \leq \alpha + \sqrt{\frac{\log(1/\delta)}{2n_1}} \right\} \geq 1 - \delta.$$

This result holds for both deterministic and randomized algorithms \mathcal{A} . Note that it is not necessary to assume that \mathcal{A} is symmetric.

In other words, the probability that a training set results in a significantly higher training-conditional miscoverage rate than the nominal rate, is vanishingly small under the split conformal method. Of course, by running split conformal at a modified value $\alpha' := \alpha - \sqrt{\frac{\log(1/\delta)}{2n_1}}$, we would obtain a slightly more conservative prediction interval that would then satisfy

$$\mathbb{P}_{P^n} \{ \alpha_P(\mathcal{D}_n) \leq \alpha \} \geq 1 - \delta.$$

This type of guarantee (i.e., with probability at least $1 - \delta$, we obtain at least $1 - \alpha$ coverage, where α and δ are specified by the user) is often referred to as a *probably approximately correct* (PAC) guarantee. This style of inference guarantee dates back to the work of [23, 22] on setting “tolerance limits”, i.e., a prediction interval (in the univariate case) or prediction region (in the multivariate case), for a random variable $Y \sim P$, given n i.i.d. draws $Y_1, \dots, Y_n \stackrel{\text{iid}}{\sim} P$ (that is, a prediction region \hat{C}_n for Y without any covariate X , such that $P(\hat{C}_n) \geq 1 - \alpha$ holds with probability at least $1 - \delta$, for user-specified parameters α and δ). More recent results offering PAC-style training-conditional coverage guarantees for the regression setting, via split conformal and related methods, can be found in the work of [7, 3, 24, 13]; see also [12, 14, 25] for training-conditional coverage under covariate shift.

No analogous finite-sample results are known for distribution-free prediction methods beyond split conformal, although [16] analyze asymptotic training-conditional validity for the jackknife and for cross-validation under algorithmic

stability type assumptions such as (9). In this work, our goal will be to examine the finite-sample training-conditional coverage properties of distribution-free methods beyond split conformal.

Object-conditional or label-conditional coverage As a second way in which marginal coverage may not be sufficient for practical utility, we may also be interested in coverage at a particular new test feature vector X_{n+1} (referred to in [17] as *object-conditional coverage*)—for instance, if the data points correspond to individual patients in a clinical setting, is it true that a given patient with a particular feature vector $X_{n+1} = x$ has a $1 - \alpha$ probability of a correct predictive interval? That is, we would like to show that the conditional coverage probability $\mathbb{P}_{P^{n+1}} \{Y_{n+1} \in \widehat{C}_n(X_{n+1}) \mid X_{n+1} = x\}$ is $\geq 1 - \alpha$, at least approximately. However, [17, 10] show that this type of guarantee is impossible under any distribution P for which X is nonatomic (i.e., $\mathbb{P}_P \{X = x\} = 0$ for all $x \in \mathcal{X}$ —for instance, this is satisfied by any continuous distribution on \mathbb{R}^d); see also [2]. A third type of conditional guarantee is that of *label-conditional coverage* [17, 11] for the setting where the response Y is categorical, requiring accuracy conditional on the class, i.e., $\mathbb{P}_{P^{n+1}} \{Y_{n+1} \in \widehat{C}_n(X_{n+1}) \mid Y_{n+1} = y\} \geq 1 - \alpha$ for each category y . Both of these type of conditional guarantees are fundamentally very different from training-conditional coverage, and we will not address these further in this work.

3. Theoretical results

As shown in Theorem 1 above, a training-conditional guarantee of the form (3) was established by [17] for the split conformal method. In our work, we find that a guarantee of the form (3) can also be shown for the K -fold CV+ method (as long as n/K , the number of data points in each fold, is sufficiently large), but no such guarantees are possible for the full conformal or jackknife+ methods. In this section, we present the main results for each of the three previously unstudied methods. The proofs will be given in Section 4 below.

First, we consider the full conformal prediction method. In contrast to split conformal, it is impossible to guarantee training-conditional coverage for the full conformal method without further assumptions.

Theorem 2. *For any sample size $n \geq 2$ and any distribution P for which the marginal P_X is nonatomic, there exists a symmetric and deterministic regression algorithm \mathcal{A} such that the full conformal prediction method defined in (5) satisfies*

$$\mathbb{P}_{P^n} \{ \alpha_P(\mathcal{D}_n) \geq 1 - n^{-2} \} \geq \alpha - 6 \sqrt{\frac{\log n}{n}}.$$

In other words, without placing assumptions on the distribution P and/or the algorithm \mathcal{A} (beyond the standard symmetry assumption), we cannot avoid the worst-case scenario (2), where the marginal guarantee of $1 - \alpha$ coverage stated

in (1) is achieved only because the training data set yields $\approx 100\%$ coverage with probability $\approx 1 - \alpha$, and $\approx 0\%$ coverage with probability $\approx \alpha$. (Our result holds only for distributions P where X is nonatomic, i.e., $\mathbb{P}_P\{X = x\} = 0$ for all $x \in \mathcal{X}$ —this condition appears also in the impossibility results for object-conditional coverage as described earlier in Section 2.2.)

Next, for the jackknife+, the same worst-case result holds.

Theorem 3. *For any sample size $n \geq 2$ and any distribution P for which the marginal P_X is nonatomic, there exists a symmetric and deterministic regression algorithm \mathcal{A} such that the jackknife+ prediction interval defined in (7) satisfies*

$$\mathbb{P}_{P^n} \{ \alpha_P(\mathcal{D}_n) \geq 1 - n^{-2} \} \geq \alpha - 6\sqrt{\frac{\log n}{n}}.$$

Thus, as for full conformal, without placing assumptions on P and/or \mathcal{A} (beyond the standard symmetry assumption), we cannot ensure that the jackknife+ method will avoid the worst-case scenario (2).

In contrast, for CV+, we will now see that the lower bound on marginal coverage, which is $\gtrsim 1 - 2\alpha$ as shown in (11), can also be obtained as a training-conditional guarantee.

Theorem 4. *For any integers $K \geq 2$ and $m \geq 1$, and let $n = Km$. Suppose CV+ is run with K folds each of size m . Then, for any regression algorithm \mathcal{A} and any distribution P , the K -fold CV+ method (10) satisfies*

$$\mathbb{P}_{P^n} \left\{ \alpha_P(\mathcal{D}_n) \leq 2\alpha + \sqrt{\frac{2\log(K/\delta)}{m}} \right\} \geq 1 - \delta$$

for any $\delta > 0$.

As long as the size of each fold, $m = n/K$, is large, the bound on $\alpha_P(\mathcal{D}_n)$ is approximately 2α . Comparing to the marginal result (11) for the CV+ method, we see that the conditional coverage guarantee (for “most” training data sets \mathcal{D}_n) essentially matches the marginal coverage guarantee, and thus could not be improved. Note also that, as in Theorem 1 for split conformal, we do not need to assume \mathcal{A} is symmetric, and the result holds regardless of whether \mathcal{A} is deterministic or randomized.

4. Proofs

Before proceeding to the proofs, we give some brief intuition for why the split conformal and CV+ methods offer training-conditional coverage guarantees, while full conformal and jackknife+ do not. For split conformal, the n_1 residuals on the holdout set

$$\{|Y_i - \widehat{\mu}_{n_0}(X_i)| : i = n_0 + 1, \dots, n\}$$

are i.i.d. after conditioning on the fitted model $\widehat{\mu}_{n_0}$, and therefore, for large n_1 , their sample quantiles concentrate around the corresponding population

quantiles. Similarly, for CV+, for each fold $k = 1, \dots, K$ we have $m = n/K$ many residuals

$$\{|Y_i - \widehat{\mu}_{[n] \setminus S_k}(X_i)| : i \in S_k\}$$

that are again i.i.d. conditional on the k -th fitted model $\widehat{\mu}_{[n] \setminus S_k}$, and thus again their sample quantiles concentrate as long as the fold size $m = n/K$ is large. This concentration of the sample quantiles (which we formalize in Lemma 1 below) is the key ingredient for establishing training-conditional coverage. On the other hand, for both full conformal and jackknife+, there is no independence among residuals calculated in each method—for example, for jackknife+, in the leave-one-out residuals $R_i = |Y_i - \widehat{\mu}_{[n] \setminus \{i\}}(X_i)|$, for two data points $i \neq j$, data point (X_i, Y_i) is used for training when computing the j -th residual R_j , and vice versa.

4.1. Proofs for split conformal and CV+

We begin by considering the split conformal and CV+ methods, which both achieve training-conditional coverage. Both the split conformal result, Theorem 1 [17, Proposition 2a], and the CV+ result, Theorem 4, can be proved as consequences of the following lemma.

Lemma 1. *Let $n \geq 2$ and choose a holdout set A with $\emptyset \subsetneq A \subsetneq [n]$. Let $\widehat{\mu}_{[n] \setminus A} = \mathcal{A}((X_i, Y_i) : i \in [n] \setminus A)$, where \mathcal{A} is any algorithm and may be deterministic or randomized. Define*

$$p_A(x, y) = \frac{1}{|A|} \sum_{i \in A} \mathbb{1} \{|Y_i - \widehat{\mu}_{[n] \setminus A}(X_i)| \geq |y - \widehat{\mu}_{[n] \setminus A}(x)|\},$$

and

$$p_A^*(x, y) = \mathbb{P}_P \{|Y - \widehat{\mu}_{[n] \setminus A}(X)| \geq |y - \widehat{\mu}_{[n] \setminus A}(x)| \mid \widehat{\mu}_{[n] \setminus A}\}.$$

Then $p_A^*(X_{n+1}, Y_{n+1})$ is a valid p -value conditional on the training data, i.e.,

$$\mathbb{P}_P \{p_A^*(X_{n+1}, Y_{n+1}) \leq a \mid \mathcal{D}_n\} \leq a \text{ for all } a \in [0, 1], \text{ almost surely over } \mathcal{D}_n. \quad (12)$$

Moreover, for any $\Delta \geq \sqrt{\frac{\log 2}{2|A|}}$,

$$\mathbb{P}_{P^n} \left\{ \sup_{(x, y) \in \mathcal{X} \times \mathbb{R}} (p_A^*(x, y) - p_A(x, y)) > \Delta \right\} \leq e^{-2|A|\Delta^2}, \quad (13)$$

Next we will see how this lemma implies the two theorems. First, for split conformal, Theorem 1 is proved by [17, Proposition 2a], but here we reformulate the proof in terms of the above lemma, to set up intuition for our CV+ proof later on.

Proof of Theorem 1 [17, Proposition 2a]. By definition of the split conformal method (4), we have

$$\widehat{C}_n(X_{n+1}) = \left\{ y \in \mathbb{R} : \widehat{Q}_{n_1} \geq |y - \widehat{\mu}_{n_0}(X_{n+1})| \right\}$$

$$\begin{aligned}
 &= \left\{ y \in \mathbb{R} : \sum_{i=n_0+1}^n \mathbb{1} \{ |Y_i - \hat{\mu}_{n_0}(X_i)| < |y - \hat{\mu}_{n_0}(X_{n+1})| \} < \lceil (1 - \alpha)(n_1 + 1) \rceil \right\} \\
 &\supseteq \left\{ y \in \mathbb{R} : \sum_{i=n_0+1}^n \mathbb{1} \{ |Y_i - \hat{\mu}_{n_0}(X_i)| \geq |y - \hat{\mu}_{n_0}(X_{n+1})| \} > \alpha n_1 \right\} \\
 &= \{ y \in \mathbb{R} : p_{[n] \setminus [n_0]}(X_{n+1}, y) > \alpha \}.
 \end{aligned}$$

where $p_{[n] \setminus [n_0]}(x, y)$ is defined as in Lemma 1 by choosing the holdout set $A = [n] \setminus [n_0]$. Therefore, we have

$$\alpha_P(\mathcal{D}_n) = \mathbb{P}_P \left\{ Y_{n+1} \notin \hat{C}_n(X_{n+1}) \mid \mathcal{D}_n \right\} \leq \mathbb{P}_P \left\{ p_{[n] \setminus [n_0]}(X_{n+1}, Y_{n+1}) \leq \alpha \mid \mathcal{D}_n \right\}.$$

Next, fixing any $\Delta > 0$, consider the event that

$$\sup_{(x,y) \in \mathcal{X} \times \mathbb{R}} \left(p_{[n] \setminus [n_0]}^*(x, y) - p_{[n] \setminus [n_0]}(x, y) \right) \leq \Delta,$$

which is a function of the training data \mathcal{D}_n . On this event, we have

$$\alpha_P(\mathcal{D}_n) \leq \mathbb{P}_P \left\{ p_{[n] \setminus [n_0]}^*(X_{n+1}, Y_{n+1}) \leq \alpha + \Delta \mid \mathcal{D}_n \right\} \leq \alpha + \Delta,$$

where the last step holds by (12). In other words, so far we have shown that

$$\begin{aligned}
 &\mathbb{P}_{P^n} \{ \alpha_P(\mathcal{D}_n) > \alpha + \Delta \} \\
 &\leq \mathbb{P}_{P^n} \left\{ \sup_{(x,y) \in \mathcal{X} \times \mathbb{R}} \left(p_{[n] \setminus [n_0]}^*(x, y) - p_{[n] \setminus [n_0]}(x, y) \right) > \Delta \right\}.
 \end{aligned}$$

Finally, applying (13), this probability is bounded by δ when we choose $\Delta = \sqrt{\frac{\log(1/\delta)}{2n_1}}$. \square

Next, we prove Theorem 4 for the CV+ method.

Proof of Theorem 4. As in the definition of the CV+ method (10), we let $R_i = |Y_i - \hat{\mu}_{[n] \setminus S_{k(i)}}(X_i)|$ for each $i \in [n]$. Following [1, Proof of Theorem 4], the CV+ prediction interval defined in (10) deterministically satisfies

$$\begin{aligned}
 \hat{C}_n(X_{n+1}) &\supseteq \left\{ y \in \mathbb{R} : \sum_{i=1}^n \mathbb{1} \{ |y - \hat{\mu}_{[n] \setminus S_{k(i)}}(X_{n+1})| > R_i \} < (1 - \alpha)(n + 1) \right\} \\
 &\supseteq \left\{ y \in \mathbb{R} : \sum_{i=1}^n \mathbb{1} \{ R_i \geq |y - \hat{\mu}_{[n] \setminus S_{k(i)}}(X_{n+1})| \} > \alpha n \right\} \\
 &= \left\{ y \in \mathbb{R} : \sum_{k=1}^K \sum_{i \in S_k} \mathbb{1} \{ R_i \geq |y - \hat{\mu}_{[n] \setminus S_k}(X_{n+1})| \} > \alpha n \right\} \\
 &= \left\{ y \in \mathbb{R} : \frac{1}{K} \sum_{k=1}^K p_{S_k}(X_{n+1}, y) > \alpha \right\},
 \end{aligned}$$

where for each fold k , $p_{S_k}(x, y)$ is defined as in Lemma 1 by choosing the holdout set $A = S_k$. Therefore, we have

$$\begin{aligned}\alpha_P(\mathcal{D}_n) &= \mathbb{P}_P \left\{ Y_{n+1} \notin \widehat{C}_n(X_{n+1}) \mid \mathcal{D}_n \right\} \\ &\leq \mathbb{P}_P \left\{ \frac{1}{K} \sum_{k=1}^K p_{S_k}(X_{n+1}, Y_{n+1}) \leq \alpha \mid \mathcal{D}_n \right\}.\end{aligned}$$

Next, fixing any $\Delta > 0$, consider the event that

$$\max_k \sup_{(x,y) \in \mathcal{X} \times \mathbb{R}} (p_{S_k}^*(x, y) - p_{S_k}(x, y)) \leq \Delta,$$

which is a function of the training data \mathcal{D}_n . On this event, we have

$$\alpha_P(\mathcal{D}_n) \leq \mathbb{P}_P \left\{ \frac{1}{K} \sum_{k=1}^K p_{S_k}^*(X_{n+1}, Y_{n+1}) \leq \alpha + \Delta \mid \mathcal{D}_n \right\} \leq 2\alpha + 2\Delta,$$

where the last step holds because each $p_{S_k}^*(X_{n+1}, Y_{n+1})$ is a valid p-value conditional on \mathcal{D}_n by (12), and the average of valid p-values is itself a p-value up to a factor of 2 [15, 21]. Combining everything so far, we have shown that

$$\begin{aligned}\mathbb{P}_{P^n} \{ \alpha_P(\mathcal{D}_n) > 2\alpha + 2\Delta \} &\leq \mathbb{P}_{P^n} \left\{ \max_k \sup_{(x,y) \in \mathcal{X} \times \mathbb{R}} (p_{S_k}^*(x, y) - p_{S_k}(x, y)) > \Delta \right\} \\ &\leq \sum_{k=1}^K \mathbb{P}_{P^n} \left\{ \sup_{(x,y) \in \mathcal{X} \times \mathbb{R}} (p_{S_k}^*(x, y) - p_{S_k}(x, y)) > \Delta \right\},\end{aligned}$$

where for the last step we take a union bound. Finally, applying (13) to bound this probability for each fold k , the above quantity is bounded by δ when we choose $\Delta = \sqrt{\frac{\log(K/\delta)}{2m}}$. \square

To conclude this section, we now prove the supporting lemma.

Proof of Lemma 1. First, for any fixed function $\mu : \mathcal{X} \rightarrow \mathbb{R}$, define

$$\bar{F}_\mu(t) = \mathbb{P}_P \{ |Y - \mu(X)| \geq t \}.$$

In other words, \bar{F}_μ is right-tailed CDF of $|Y - \mu(X)|$ under $(X, Y) \sim P$. We can therefore write

$$p_A^*(X_{n+1}, Y_{n+1}) = \bar{F}_{\widehat{\mu}_{[n] \setminus A}}(|Y_{n+1} - \widehat{\mu}_{[n] \setminus A}(X_{n+1})|).$$

Since $(X_{n+1}, Y_{n+1}) \sim P$ (and is independent of $\widehat{\mu}_{[n] \setminus A}$), this is clearly a valid p-value by definition of $\bar{F}_{\widehat{\mu}_{[n] \setminus A}}$, and so we have proved (12).

Next, for any $(x, y) \in \mathcal{X} \times \mathbb{R}$, we can calculate

$$p_A^*(x, y) - p_A(x, y)$$

$$\begin{aligned}
 &= \bar{F}_{\hat{\mu}_{[n]\setminus A}}(|y - \hat{\mu}_{[n]\setminus A}(x)|) - \frac{1}{|A|} \sum_{i \in A} \mathbb{1} \{ |Y_i - \hat{\mu}_{[n]\setminus A}(X_i)| \geq |y - \hat{\mu}_{[n]\setminus A}(x)| \} \\
 &\leq \sup_{t \in \mathbb{R}} \left(\bar{F}_{\hat{\mu}_{[n]\setminus A}}(t) - \frac{1}{|A|} \sum_{i \in A} \mathbb{1} \{ |Y_i - \hat{\mu}_{[n]\setminus A}(X_i)| \geq t \} \right).
 \end{aligned}$$

Finally, since $\{(X_i, Y_i)\}_{i \in A}$ are drawn i.i.d. from P and are independent from $\hat{\mu}_{[n]\setminus A}$, for any $\Delta \geq \sqrt{\frac{\log 2}{2|A|}}$ the Dvoretzky–Kiefer–Wolfowitz inequality implies that, conditional on $\hat{\mu}_{[n]\setminus A}$,

$$\sup_{t \in \mathbb{R}} \left(\bar{F}_{\hat{\mu}_{[n]\setminus A}}(t) - \frac{1}{|A|} \sum_{i \in A} \mathbb{1} \{ |Y_i - \hat{\mu}_{[n]\setminus A}(X_i)| \geq t \} \right) \leq \Delta$$

holds with probability at least $1 - e^{-2|A|\Delta^2}$. The same bound therefore holds marginally as well. This proves (13). \square

4.2. Proofs for full conformal and jackknife+

Next, we turn to the results for full conformal and for jackknife+, where we show that training-conditional coverage cannot be guaranteed without further assumptions. The proofs for the two methods are closely related and share the same structure.

First, fix some large integer M (which we will specify later), and partition \mathcal{X} into M sets, $\mathcal{X} = A_0 \cup A_1 \cup \dots \cup A_{M-1}$, where $\mathbb{P}_P \{X \in A_m\} = \frac{1}{M}$ for each $m = 0, \dots, M - 1$ (since we have assumed X is nonatomic under the distribution P , such a partition exists by [5, Proposition A.1]). Define a map $a : \mathcal{X} \rightarrow \{0, \dots, M - 1\}$,

$$a(x) = \begin{cases} 0, & x \in A_0, \\ 1, & x \in A_1, \\ \dots & \\ M - 1, & x \in A_{M-1}, \end{cases}$$

assigning each $x \in \mathcal{X}$ to a particular set in the partition. Then, by our choice of the A_m 's, we see that

$$a(X) \sim \text{Unif}\{0, \dots, M - 1\}$$

under the distribution P . By extension, $\text{mod}(\sum_{i=1}^n a(X_i), M) \sim \text{Unif}\{0, \dots, M - 1\}$. For the sake of illustration, consider a “clock” partitioned into M segments and a hand whose position represents the value of the modulo. In this case, the hand moves forward by $a(X_i)$ segments when the i -th term is added inside the modulo—see Figure 1 for an illustration.

Next, we will define three events, \mathcal{E}_{\max} , \mathcal{E}_{mod} , and $\mathcal{E}_{\text{unif}}$, which are all functions of the training data \mathcal{D}_n . Define

$$y_* = \text{the } (1 - n^{-2})\text{-quantile of } |Y| \text{ under distribution } P,$$

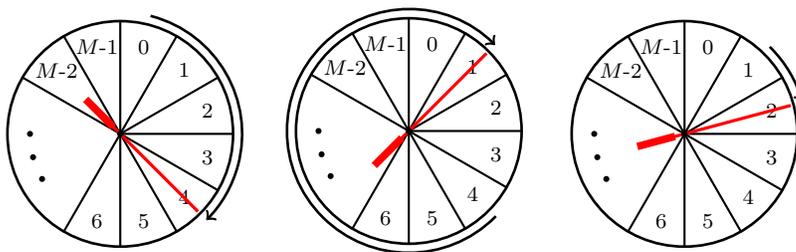


FIG 1. Representation of $\text{mod}(\sum_{i=1}^n a(X_i), M)$ for $n = 3$, when $a(X_1) = 4$, $a(X_2) = M - 3$, and $a(X_3) = 1$. The left plot shows moving from 0 to $\text{mod}(a(X_1), M)$, the middle plot shows moving from $\text{mod}(a(X_1), M)$ to $\text{mod}(a(X_1) + a(X_2), M)$, and the right plot shows moving from $\text{mod}(a(X_1) + a(X_2), M)$ to $\text{mod}(a(X_1) + a(X_2) + a(X_3), M)$.

and let

$$M_1 = \left\lfloor M \left(\alpha - \sqrt{\frac{2 \log n}{n}} - \frac{2}{n} \right) \right\rfloor \approx \alpha M.$$

(Note that we can assume n is sufficiently large so that $\alpha - \sqrt{\frac{2 \log n}{n}} - \frac{2}{n} > 0$, and thus $M_1 \geq 0$, since if this does not hold then the results of the theorems hold trivially.) With these values fixed, the three events are defined as follows:

- Let \mathcal{E}_{\max} be the event that $\max_{i \in [n]} |Y_i| < y_*$.
- Let \mathcal{E}_{mod} be the event that $\text{mod}(\sum_{i=1}^n a(X_i), M) < M_1$.
- Let $\mathcal{E}_{\text{unif}}$ be the event that $\sum_{i=1}^n \mathbb{1} \{ \text{mod}(a(X_i) + m, M) < M - M_1 \} \geq \lceil (1 - \alpha)(n + 1) \rceil$ for all integers m .

Figures 2 and 3 illustrate the events \mathcal{E}_{mod} and $\mathcal{E}_{\text{unif}}$, respectively.

The following result shows that, with probability at least $\approx \alpha$, all three events occur:

Lemma 2. *Under the definitions and notation above, for $(X_1, Y_1), \dots, (X_n, Y_n) \stackrel{\text{iid}}{\sim} P$, we have*

$$\mathbb{P}_{P^n} \{ \mathcal{E}_{\text{mod}} \} \geq \alpha - \sqrt{\frac{2 \log n}{n}} - \frac{2}{n} - \frac{1}{M}, \quad \mathbb{P}_{P^n} \{ \mathcal{E}_{\max} \} \geq 1 - \frac{1}{n}, \quad \mathbb{P}_{P^n} \{ \mathcal{E}_{\text{unif}} \} \geq 1 - \frac{2}{n},$$

and therefore,

$$\mathbb{P}_{P^n} \{ \mathcal{E}_{\text{mod}} \cap \mathcal{E}_{\max} \cap \mathcal{E}_{\text{unif}} \} \geq \alpha - \sqrt{\frac{2 \log n}{n}} - \frac{1}{M} - \frac{5}{n}.$$

By choosing M to be sufficiently large, then, we obtain

$$\mathbb{P}_{P^n} \{ \mathcal{E}_{\text{mod}} \cap \mathcal{E}_{\max} \cap \mathcal{E}_{\text{unif}} \} \geq \alpha - 6\sqrt{\frac{\log n}{n}}. \tag{14}$$

Now it remains to be shown that, for both full conformal prediction and for jackknife+, we can find an algorithm \mathcal{A} such that, if the events \mathcal{E}_{\max} , \mathcal{E}_{mod} , and

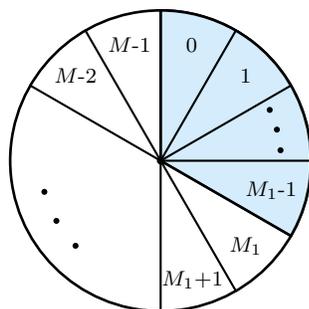


FIG 2. An illustration of the event \mathcal{E}_{mod} , which is the event that $\text{mod}(\sum_{i=1}^n a(X_i), M) < M_1$ for $M_1 \approx \alpha M$. In the figure, the event \mathcal{E}_{mod} holds if and only if the value $\text{mod}(\sum_{i=1}^n a(X_i), M)$ lands in the shaded region of the “clock”.

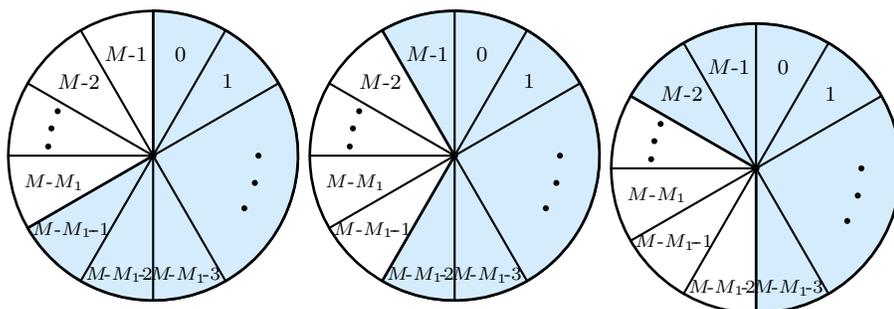


FIG 3. An illustration of the event $\mathcal{E}_{\text{unif}}$, which is the event that $\text{mod}(a(X_i) + m, M) < M - M_1$ holds for at least $\lceil (1 - \alpha)(n + 1) \rceil$ many training data points i , for every integer m . In the figure, the shaded region $\{a \in \{0, \dots, M - 1\} : \text{mod}(a + m, M) < M - M_1\}$ is shown for $m = 0$ (left figure), $m = 1$ (center figure), and $m = 2$ (right figure). The event holds if and only if at least $\lceil (1 - \alpha)(n + 1) \rceil$ many training data points $i \in [n]$ have $a(X_i)$ lying in the shaded region, for each integer m (i.e., for each of the three displayed figures, as well as all other possible values of m).

$\mathcal{E}_{\text{unif}}$ all hold, then the training-conditional miscoverage rate $\alpha_P(\mathcal{D}_n)$ is close to 1.

Proof of Theorem 2. For full conformal, we define a symmetric regression algorithm \mathcal{A} that maps a data set $\{(x_1, y_1), \dots, (x_{n+1}, y_{n+1})\}$ to the fitted function

$$\hat{\mu}(x) = \begin{cases} 2y_*, & \text{if } \text{mod}(-a(x) + \sum_{i=1}^{n+1} a(x_i), M) < M_1, \\ 0, & \text{otherwise.} \end{cases}$$

Below, we will show that, for any training data set \mathcal{D}_n ,

$$\text{If } \mathcal{E}_{\text{max}} \cap \mathcal{E}_{\text{mod}} \cap \mathcal{E}_{\text{unif}} \text{ holds, then } \hat{C}_n(X_{n+1}) \subseteq (y_*, \infty) \text{ almost surely over } X_{n+1}. \tag{15}$$

By definition of y_* , we therefore have

$$\begin{aligned} \mathbb{P}_P \left\{ Y_{n+1} \in \widehat{C}_n(X_{n+1}) \mid \mathcal{D}_n \right\} &\leq \mathbb{P}_P \left\{ Y_{n+1} \in (y_*, \infty) \mid \mathcal{D}_n \right\} \\ &= \mathbb{P}_P \left\{ Y_{n+1} > y_* \right\} \leq n^{-2}, \end{aligned}$$

for any \mathcal{D}_n such that $\mathcal{E}_{\max} \cap \mathcal{E}_{\text{mod}} \cap \mathcal{E}_{\text{unif}}$ holds. Combining this with the bound (14), we have proved that $\mathbb{P}_{P^n} \left\{ \alpha_P(\mathcal{D}_n) \geq 1 - n^{-2} \right\} \geq \alpha - 6\sqrt{\frac{\log n}{n}}$, as desired.

To complete the proof, we now verify (15). Condition on the training data \mathcal{D}_n , and assume $\mathcal{E}_{\max} \cap \mathcal{E}_{\text{mod}} \cap \mathcal{E}_{\text{unif}}$ holds. First, by \mathcal{E}_{mod} , we have

$$\text{mod} \left(-a(X_{n+1}) + \sum_{i=1}^{n+1} a(X_i), M \right) = \text{mod} \left(\sum_{i=1}^n a(X_i), M \right) < M_1$$

for any value of X_{n+1} , and therefore, for any $y \in \mathbb{R}$,

$$\widehat{\mu}_{n+1}^y(X_{n+1}) = 2y_*$$

where as before, $\widehat{\mu}_{n+1}^y = \mathcal{A}((X_1, Y_1), \dots, (X_n, Y_n), (X_{n+1}, y))$. On the other hand, for any $i \in [n]$, define

$$\begin{aligned} m_i &= \text{mod} \left(-a(X_i) + \sum_{j=1}^{n+1} a(X_j), M \right) \\ &= M - 1 - \text{mod} \left(a(X_i) - \sum_{j=1}^{n+1} a(X_j) - 1, M \right), \end{aligned}$$

where the last step holds since $\text{mod}(k, M) = M - 1 - \text{mod}(-k - 1, M)$ for all integers k . By the event $\mathcal{E}_{\text{unif}}$ (applied with $m = -\sum_{j=1}^{n+1} a(X_j) - 1$), we see that $m_i \geq M_1$ for at least $\lceil (1 - \alpha)(n + 1) \rceil$ many $i \in [n]$. Therefore, for all $y \in \mathbb{R}$, $\widehat{\mu}_{n+1}^y(X_i) = 0$ for at least $\lceil (1 - \alpha)(n + 1) \rceil$ many $i \in [n]$. Next, by \mathcal{E}_{\max} , we have $|Y_i| < y_*$ for all $i \in [n]$, and therefore, for all $y \in \mathbb{R}$, $R_i^y < y_*$ for at least $\lceil (1 - \alpha)(n + 1) \rceil$ many $i \in [n]$.

Returning to the definition of full conformal prediction given in (5), we therefore have $\widehat{Q}_{n+1}^y < y_*$. Therefore, $y \in \widehat{C}_n(X_{n+1})$ can hold only if $|y - \widehat{\mu}_{n+1}^y(X_{n+1})| < y_*$, which implies $\widehat{C}_n(X_{n+1}) \subseteq (y_*, 3y_*) \subseteq (y_*, \infty)$. This verifies (15), and thus completes the proof of the theorem. \square

Proof of Theorem 3. For jackknife+, we define a symmetric regression algorithm \mathcal{A} that maps a data set $\{(x_1, y_1), \dots, (x_{n-1}, y_{n-1})\}$ to the fitted function

$$\widehat{\mu}(x) = \begin{cases} 0, & \text{if } \text{mod}(a(x) + \sum_{i=1}^{n-1} a(x_i), M) < M_1, \\ 2y_*, & \text{otherwise.} \end{cases}$$

As in the proof of Theorem 2, it is sufficient to verify that (15) again holds in this case.

Condition on the training data \mathcal{D}_n , and assume $\mathcal{E}_{\max} \cap \mathcal{E}_{\text{mod}} \cap \mathcal{E}_{\text{unif}}$ holds. First, by \mathcal{E}_{mod} , for all $i \in [n]$, we have

$$\text{mod} \left(a(X_i) + \sum_{j \in [n] \setminus \{i\}} a(X_j), M \right) = \text{mod} \left(\sum_{j=1}^n a(X_j), M \right) < M_1,$$

and therefore

$$\widehat{\mu}_{[n] \setminus \{i\}}(X_i) = 0.$$

By \mathcal{E}_{\max} , we have $|Y_i| < y_*$ for all $i \in [n]$, and therefore, $R_i = |Y_i - \widehat{\mu}_{[n] \setminus \{i\}}(X_i)| < y_*$ for all $i \in [n]$.

On the other hand, for any $i \in [n]$, define

$$m_i = \text{mod} \left(a(X_{n+1}) + \sum_{j \in [n] \setminus \{i\}} a(X_j), M \right) = \text{mod} \left(-a(X_i) + \sum_{j=1}^{n+1} a(X_j), M \right).$$

Exactly as in the proof of Theorem 2, by the event $\mathcal{E}_{\text{unif}}$ we see that $m_i \geq M_1$ for at least $\lceil (1 - \alpha)(n + 1) \rceil$ many $i \in [n]$. Therefore, $\widehat{\mu}_{[n] \setminus \{i\}}(X_{n+1}) = 2y_*$ for at least $\lceil (1 - \alpha)(n + 1) \rceil$ many $i \in [n]$.

Combining these calculations, we see that $\widehat{\mu}_{[n] \setminus \{i\}}(X_{n+1}) - R_i > y_*$ for at least $\lceil (1 - \alpha)(n + 1) \rceil$ many $i \in [n]$. Thus, by definition of the jackknife+ predictive interval given in (7), we have $\widehat{C}_n(X_{n+1}) \subseteq (y_*, \infty)$. This verifies (15), and thus completes the proof of the theorem. \square

Finally, we need to prove Lemma 2.

Proof of Lemma 2. First, since y_* is chosen to be the $(1 - n^{-2})$ quantile of $|Y|$ under the distribution P , we have

$$\mathbb{P}_{P^n} \left\{ \max_{i \in [n]} |Y_i| \geq y_* \right\} \leq \sum_{i=1}^n \mathbb{P}_P \{ |Y_i| \geq y_* \} \leq n \cdot n^{-2} = \frac{1}{n},$$

and thus $\mathbb{P}_{P^n} \{ \mathcal{E}_{\max} \} \geq 1 - \frac{1}{n}$ as desired.

Next, since $a(X_i) \stackrel{\text{iid}}{\sim} \text{Unif}\{0, \dots, M - 1\}$ for $i \in [n]$, it follows immediately that $\text{mod}(a(X_1) + \dots + a(X_n), M) \sim \text{Unif}\{0, \dots, M - 1\}$ also, and so by definition of M_1 we have

$$\mathbb{P}_{P^n} \{ \mathcal{E}_{\text{mod}} \} = \frac{M_1}{M} = \frac{\left\lfloor M \left(\alpha - \sqrt{\frac{2 \log n}{n} - \frac{2}{n}} \right) \right\rfloor}{M} \geq \alpha - \sqrt{\frac{2 \log n}{n}} - \frac{2}{n} - \frac{1}{M}.$$

Finally, we turn to $\mathcal{E}_{\text{unif}}$. This is the event that

$$\sum_{i=1}^n \mathbb{1} \{ \text{mod}(a(X_i) + m, M) < M - M_1 \} \geq \lceil (1 - \alpha)(n + 1) \rceil$$

for all integers m , but by definition of the modulo function, it is equivalent to requiring that this bound holds only for all integers $m = 0, \dots, M - 1$, i.e.,

$$\mathcal{E}_{\text{unif}} = \bigcap_{m=0}^{M-1} \left\{ \sum_{i=1}^n \mathbb{1} \{ \text{mod}(a(X_i) + m, M) < M - M_1 \} \geq \lceil (1 - \alpha)(n + 1) \rceil \right\}.$$

Now let $U_1, \dots, U_n \stackrel{\text{iid}}{\sim} \text{Unif}[0, 1]$. Then $\lfloor MU_i \rfloor \stackrel{\text{iid}}{\sim} \text{Unif}\{0, \dots, M - 1\}$, and so

$$\begin{aligned} & \mathbb{P}_{P^n} \{ \mathcal{E}_{\text{unif}} \} \\ &= \mathbb{P} \left\{ \bigcap_{m=0}^{M-1} \left\{ \sum_{i=1}^n \mathbb{1} \{ \text{mod}(\lfloor MU_i \rfloor + m, M) < M - M_1 \} \geq \lceil (1 - \alpha)(n + 1) \rceil \right\} \right\}. \end{aligned} \tag{16}$$

Next, suppose it holds that

$$\sup_{s \in [0,1]} \left| \sum_i \mathbb{1} \{ U_i < s \} - ns \right| \leq \sqrt{\frac{n \log n}{2}}. \tag{17}$$

Then, for each $m \in \{0, 1, \dots, M - M_1 - 1\}$, we have

$$\begin{aligned} & \sum_{i=1}^n \mathbb{1} \{ \text{mod}(\lfloor MU_i \rfloor + m, M) < M - M_1 \} \\ &= \sum_{i=1}^n \mathbb{1} \{ \lfloor MU_i \rfloor < M - M_1 - m \} + \sum_{i=1}^n \mathbb{1} \{ \lfloor MU_i \rfloor \geq M - m \} \\ &= \sum_{i=1}^n \mathbb{1} \left\{ U_i < 1 - \frac{M_1 + m}{M} \right\} + n - \sum_{i=1}^n \mathbb{1} \left\{ U_i < 1 - \frac{m}{M} \right\} \\ &\geq \left(n \cdot \left(1 - \frac{M_1 + m}{M} \right) - \sqrt{\frac{n \log n}{2}} \right) + n - \left(n \cdot \left(1 - \frac{m}{M} \right) + \sqrt{\frac{n \log n}{2}} \right) \\ &= n \cdot \left(1 - \frac{M_1}{M} \right) - \sqrt{2n \log n} \\ &\geq \lceil (1 - \alpha)(n + 1) \rceil, \end{aligned}$$

where the last step holds by definition of M_1 . Next, for each $m \in \{M - M_1, \dots, M - 1\}$, we have

$$\begin{aligned} & \sum_{i=1}^n \mathbb{1} \{ \text{mod}(\lfloor MU_i \rfloor + m, M) < M - M_1 \} \\ &= \sum_{i=1}^n \mathbb{1} \{ M - m \leq \lfloor MU_i \rfloor < 2M - M_1 - m \} \\ &= \sum_{i=1}^n \mathbb{1} \left\{ U_i < 2 - \frac{M_1 + m}{M} \right\} - \sum_{i=1}^n \mathbb{1} \left\{ U_i < 1 - \frac{m}{M} \right\} \end{aligned}$$

$$\begin{aligned}
 &\geq \left(n \cdot \left(2 - \frac{M_1 + m}{M} \right) - \sqrt{\frac{n \log n}{2}} \right) - \left(n \cdot \left(1 - \frac{m}{M} \right) + \sqrt{\frac{n \log n}{2}} \right) \\
 &= n \cdot \left(1 - \frac{M_1}{M} \right) - \sqrt{2n \log n} \\
 &\geq \lceil (1 - \alpha)(n + 1) \rceil.
 \end{aligned}$$

Therefore, returning to (16), we have

$$\mathbb{P}_{P^n} \{ \mathcal{E}_{\text{unif}} \} \geq \mathbb{P} \left\{ \sup_{s \in [0,1]} \left| \sum_i \mathbb{1} \{ U_i \leq s \} - ns \right| \leq \sqrt{\frac{n \log n}{2}} \right\} \geq 1 - \frac{2}{n},$$

where the last step holds by the Dvoretzky–Kiefer–Wolfowitz inequality. This completes the proof. \square

5. Empirical results

The theoretical results above suggest that we should be concerned about the training conditional coverage of the full conformal and jackknife+ prediction intervals. However, the algorithms used as counterexamples in the proof are extremely unrealistic. In particular, the constructions appearing in the proofs display an extremely high amount of instability, since the inclusion of a single training point can greatly impact the output of the regression function.

Therefore, a natural question is how large the variability of $\alpha_P(\mathcal{D}_n)$ is in practice, particularly in unstable environments. In our simulation, we will examine the empirical performance of the training-conditional miscoverage rate $\alpha_P(\mathcal{D}_n)$ for the four distribution-free predictive inference tools studied in this work, for a linear regression task where high dimensionality may cause some instability in the regression algorithm.¹

5.1. Setting

We choose a target coverage rate of 90%, i.e., $\alpha = 0.1$, and will compare the performance of split conformal (with $n_0 = n_1 = n/2$), full conformal, jackknife+, and CV+ (with $K = 20$ folds).

We use sample size $n = 500$ for the training set, and $n_{\text{test}} = 1000$ for the test set. For each trial, we generate i.i.d. data points (X_i, Y_i) , $i = 1, \dots, n + n_{\text{test}}$, from the following distribution:

$$X_i \sim \mathcal{N}(0, I_d) \text{ and } Y_i \mid X_i \sim \mathcal{N}(X_i^\top \beta, 1),$$

where $\beta = \sqrt{10} \cdot U$ for a random unit vector U drawn uniformly from the unit sphere in $\mathcal{X} = \mathbb{R}^d$. We repeat the experiment for each dimension $d = 125, 250, 500, 1000$, with 200 independent trials for each dimension.

¹Code to reproduce the experiment is available at http://rinafb.github.io/code/training_conditional.zip.

Our algorithm \mathcal{A} is given by ridge regression with penalty parameter $\lambda = 0.0001$, i.e., for any data set $(x_1, y_1), \dots, (x_N, y_N)$, the fitted model $\hat{\mu} = \mathcal{A}((x_1, y_1), \dots, (x_N, y_N))$ is given by

$$\hat{\mu}(x) = x^\top \hat{\beta} \text{ where } \hat{\beta} = \arg \min_{\beta \in \mathbb{R}^d} \left\{ \sum_{i=1}^N (y_i - x_i^\top \beta)^2 + \lambda \|\beta\|_2^2 \right\}.$$

Finally, we estimate the training-conditional miscoverage rate $\alpha_P(\mathcal{D}_n)$ for each of the four methods, by computing the empirical coverage over the n_{test} many test points:

$$\alpha_P(\mathcal{D}_n) \approx \frac{1}{n_{\text{test}}} \sum_{i=1}^{n_{\text{test}}} \mathbb{1} \left\{ Y_{n+i} \in \hat{C}_n(X_{n+i}) \right\}.$$

Instability of the algorithm In this simulation, we apply the ridge regression algorithm to a training set where the x_i 's are standard Gaussian, with a penalty parameter $\lambda \approx 0$. This optimization problem is extremely poorly conditioned when the training set size is $\approx d$, but is well-behaved if the number of training points is *either* sufficiently large or sufficiently small relative to d (see [6] for an analysis of “ridgeless” regression, i.e., taking $\lambda \rightarrow 0$, in the over-parametrized setting). As a result, if the number of training points is $\approx d$, the outcome of the algorithm may be highly unstable—the stability assumption (9) will not hold, and in general, predictions $\hat{\mu}(x)$ will vary greatly with a new draw of the training set. However, instability will not occur if the training set size is substantially smaller than d or larger than d .

For split conformal, since the model is trained on $n_0 = n/2 = 250$ many data points, this instability will be high for $d = 250$ (but not for $d = 125, 500, 1000$). In contrast, when running full conformal or jackknife+ or CV+, the models are trained on $n + 1 = 501$ or $n - 1 = 499$ or $n - n/K = 475$ many data points, respectively. Therefore, for these three methods, we expect instability to be high for dimension $d = 500$ (but not for $d = 125, 250, 1000$).

5.2. Results

The results of the simulations are displayed in Figure 4. For both split conformal and CV+, as the theory suggests, the training-conditional miscoverage rate $\alpha_P(\mathcal{D}_n)$ is consistently near or below the nominal level $\alpha = 0.1$. This is the case even for dimensions where the trained models for split conformal or for CV+ are likely to exhibit instability, as discussed above. Specifically, for split conformal, the $\alpha_P(\mathcal{D}_n)$ values concentrate around α for each choice of dimension d . For CV+, the same is true for dimensions $d = 125, 250, 1000$ where the algorithm is fairly stable, while in the unstable regime $d = 500$, CV+ appears to be highly conservative, with $\alpha_P(\mathcal{D}_n)$ values consistently much lower than α . (The same outcome occurs if we repeat the experiment with dimension $d = 475$, where instability for CV+ is highest, but for brevity we do not show results for this value of d .)

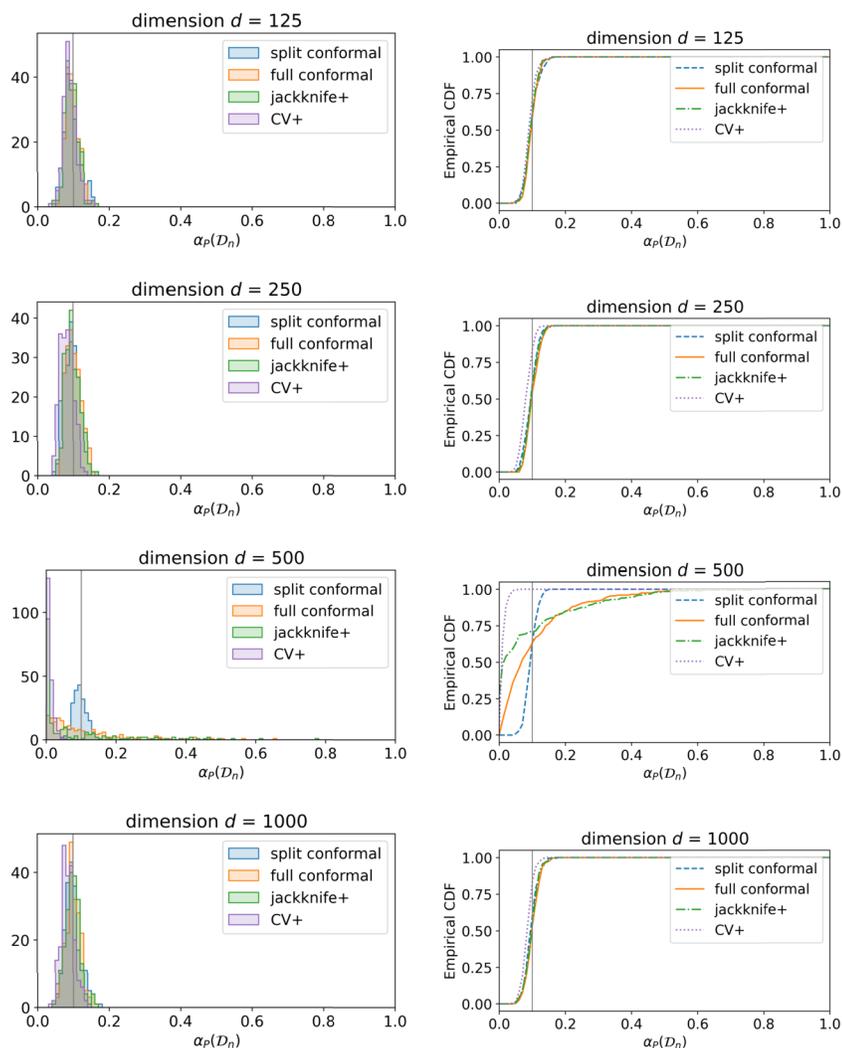


FIG 4. Plots of the estimated training-conditional miscoverage level $\alpha_P(\mathcal{D}_n)$ for 200 independent trials of the simulation, at each dimension $d = 125, 250, 500, 1000$. The gray vertical lines indicates the target miscoverage level $\alpha = 0.1$. For each d , the empirical distribution of $\alpha_P(\mathcal{D}_n)$ is displayed as a histogram on the left, and as an empirical CDF on the right.

In contrast, for full conformal and for jackknife+, we see that at $d = 500$ (where the trained models for these two methods are likely to be unstable), the training-conditional miscoverage rate $\alpha_P(\mathcal{D}_n)$ is highly variable—specifically, we see that $\alpha_P(\mathcal{D}_n)$ is substantially higher than nominal level $\alpha = 0.1$ for a large fraction of the trials. On the other hand, the training-conditional miscoverage rate $\alpha_P(\mathcal{D}_n)$ concentrates near $\alpha = 0.1$ for both methods, for all other values of

d. This is true both for a low-dimensional setting when $d = 125, 250$ and a high-dimensional (i.e., overparameterized) setting when $d = 1000$, suggesting that algorithmic stability may play a key role in understanding training-conditional coverage, as we discuss further below.

6. Conclusion

In this paper, we examine one form of conditional validity for methods of distribution-free predictive inference: training-conditional validity. While this form of validity has been previously established for the split conformal prediction method, here we examined whether this property holds for other distribution-free prediction tools. We showed that training conditional coverage guarantees can be ensured for the CV+ method, but are not possible for either the full conformal or jackknife+ methods without additional assumptions. In addition, we demonstrated empirically that training-conditional miscoverage rates far above the nominal level α can occur in realistic data sets with the latter two methods.

6.1. The role of algorithmic stability

An interesting open question is whether there are any mild assumptions that would ensure training-conditional coverage for full conformal and/or for jackknife+. One possibility is to consider algorithmic stability assumptions such as (9). In particular, our empirical results show that poor training-conditional coverage for these two methods is observed exactly in those settings where the behavior of the regression algorithm \mathcal{A} is highly unstable (specifically, when $d \approx n$, in our linear regression simulation). This suggests that assuming stability of \mathcal{A} could potentially be sufficient to ensure training-conditional coverage for these methods. We leave this open question for future work.

Acknowledgments

R.F.B. was supported by the National Science Foundation via grants DMS-1654076 and DMS-2023109, and by the Office of Naval Research via grant N00014-20-1-2337.

References

- [1] BARBER, R. F., CANDÈS, E. J., RAMDAS, A. and TIBSHIRANI, R. J. (2021). Predictive inference with the jackknife+. *The Annals of Statistics* **49** 486–507. [MR4206687](#)
- [2] BARBER, R. F., CANDÈS, E. J., RAMDAS, A. and TIBSHIRANI, R. J. (2021). The limits of distribution-free conditional predictive inference. *Information and Inference: A Journal of the IMA* **10** 455–482. [MR4270755](#)

- [3] BATES, S., ANGELOPOULOS, A., LEI, L., MALIK, J. and JORDAN, M. (2021). Distribution-free, risk-controlling prediction sets. *Journal of the ACM (JACM)* **68** 1–34. [MR4402354](#)
- [4] CHEN, W., CHUN, K.-J. and BARBER, R. F. (2018). Discretized conformal prediction for efficient distribution-free inference. *Stat* **7** e173. [MR3769053](#)
- [5] DUDLEY, R. M. and NORVAIŠA, R. (2011). *Concrete functional calculus*. Springer. [MR2732563](#)
- [6] HASTIE, T., MONTANARI, A., ROSSET, S. and TIBSHIRANI, R. J. (2022). Surprises in high-dimensional ridgeless least squares interpolation. *The Annals of Statistics* **50** 949–986. [MR4404925](#)
- [7] 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.
- [8] LEI, J. (2019). Fast exact conformalization of the lasso using piecewise linear homotopy. *Biometrika* **106** 749–764. [MR4031197](#)
- [9] LEI, J., G’SSELL, M., RINALDO, A., TIBSHIRANI, R. J. and WASSERMAN, L. (2018). Distribution-free predictive inference for regression. *Journal of the American Statistical Association* **113** 1094–1111. [MR3862342](#)
- [10] LEI, J. and WASSERMAN, L. (2014). Distribution-free prediction bands for non-parametric regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **76** 71–96. [MR3153934](#)
- [11] LÖFSTRÖM, T., BOSTRÖM, H., LINUSSON, H. and JOHANSSON, U. (2015). Bias reduction through conditional conformal prediction. *Intelligent Data Analysis* **19** 1355–1375.
- [12] PARK, S., DOBRIBAN, E., LEE, I. and BASTANI, O. (2022). PAC prediction sets under covariate shift. In *The 10th International Conference on Learning Representations, ICLR 2022*.
- [13] PARK, S., LI, S., LEE, I. and BASTANI, O. (2021). PAC confidence predictions for deep neural network classifiers. In *The 9th International Conference on Learning Representations, ICLR 2021*.
- [14] QIU, H., DOBRIBAN, E. and TCHETGEN, E. T. (2023). Prediction sets adaptive to unknown covariate shift. To appear in *Journal of the Royal Statistical Society Series B: Statistical Methodology*.
- [15] RÜSCHENDORF, L. (1982). Random variables with maximum sums. *Advances in Applied Probability* **14** 623–632. [MR0665297](#)
- [16] STEINBERGER, L. and LEEB, H. (2023). Conditional predictive inference for stable algorithms. *The Annals of Statistics* **51**(1) 290–311. [MR4564857](#)
- [17] VOVK, V. (2012). Conditional validity of inductive conformal predictors. In *Asian conference on machine learning* 475–490. PMLR.
- [18] VOVK, V. (2015). Cross-conformal predictors. *Annals of Mathematics and Artificial Intelligence* **74** 9–28. [MR3353894](#)
- [19] VOVK, V., GAMMERMAN, A. and SHAFER, G. (2005). *Algorithmic learning in a random world*. Springer Science & Business Media. [MR2161220](#)
- [20] VOVK, V., NOURETDINOV, I., MANOKHIN, V. and GAMMERMAN, A. (2018). Cross-conformal predictive distributions. In *Conformal and Probabilistic Prediction and Applications* 37–51. PMLR.

- [21] VOVK, V. and WANG, R. (2020). Combining p-values via averaging. *Biometrika* **107** 791–808. [MR4186488](#)
- [22] WALD, A. (1943). An extension of Wilks' method for setting tolerance limits. *The Annals of Mathematical Statistics* **14** 45–55. [MR0007965](#)
- [23] WILKS, S. S. (1941). Determination of sample sizes for setting tolerance limits. *The Annals of Mathematical Statistics* **12** 91–96. [MR0004451](#)
- [24] YANG, Y. and KUCHIBHOTLA, A. K. (2021). Finite-sample efficient conformal prediction. *arXiv preprint arXiv:2104.13871*. [MR4460565](#)
- [25] YANG, Y., KUCHIBHOTLA, A. K. and TCHETGEN, E. T. (2022). Doubly Robust Calibration of Prediction Sets under Covariate Shift. *arXiv preprint arXiv:2203.01761*.