



Root-finding approaches for computing conformal prediction set

Eugene Ndiaye¹ · Ichiro Takeuchi²

Received: 25 June 2021 / Revised: 30 June 2022 / Accepted: 12 August 2022 /
Published online: 23 November 2022

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2022

Abstract

Conformal prediction constructs a confidence set for an unobserved response of a feature vector based on previous identically distributed and exchangeable observations of responses and features. It has a coverage guarantee at any nominal level without additional assumptions on their distribution. Its computation deplorably requires a refitting procedure for all replacement candidates of the target response. In regression settings, this corresponds to an infinite number of model fits. Apart from relatively simple estimators that can be written as pieces of linear function of the response, efficiently computing such sets is difficult, and is still considered as an open problem. We exploit the fact that, *often*, conformal prediction sets are intervals whose boundaries can be efficiently approximated by classical root-finding algorithms. We investigate how this approach can overcome many limitations of formerly used strategies; we discuss its complexity and drawbacks.

Keywords Prediction Set · Uncertainty Quantification · Distribution-Free · Inference · Conformal Prediction · Reliability

1 Introduction

Gamerman et al. (1998), Vovk et al. (2005), Shafer and Vovk (2008) introduced Conformal Prediction (CP) as a general method for predicting a confidence set of a random variable from its point prediction. Given an observed data set $\mathcal{D}_n = \{(x_1, y_1), \dots, (x_n, y_n)\}$ sampled from a distribution \mathbb{P} , it constructs a $100(1 - \alpha)\%$ confidence set that contains the unobserved response y_{n+1} of a new instance x_{n+1} . In this way, it equips traditional

Editor: Henrik Boström

✉ Eugene Ndiaye
ndiayeeugene@gmail.com

Ichiro Takeuchi
ichiro.takeuchi@mae.nagoya-u.ac.jp

¹ H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, USA

² Faculty of Engineering, Nagoya University / RIKEN AIP, Furo-cho, Nagoya, Japan

statistical learning algorithms with a confidence value when predicting the response of a new test example. The general idea is to learn a predictive model on the augmented database $\mathcal{D}_{n+1}(z) = \mathcal{D}_n \cup (x_{n+1}, z)$ where z replaces the unknown response y_{n+1} . We can therefore define a prediction loss for each observation, and rank them. A candidate z will be considered as conformal or typical if the rank of its loss is sufficiently small. The conformal prediction set will merely collect the most typical z as a confidence set for y_{n+1} . As long as the sequence $\{(x_i, y_i)\}_{i=1}^{n+1}$ is exchangeable¹, and the predictive model is invariant with respect to permutation of the data, this method benefits from a strong coverage guarantee without any assumption on the distribution. This holds for any finite sample size n .

Several extensions, and applications of conformal prediction have been developed for designing uncertainty sets in active learning (Ho & Wechsler, 2008), anomaly detection (Laxhammar & Falkman, 2015; Bates et al., 2021b), image classification (Angelopoulos et al., 2020), few shot learning (Fisch et al., 2021), time series (Xu & Xie, 2021) or to infer the performance guarantee for statistical learning algorithms (Holland, 2020; Cella & Ryan, 2020; Bates et al., 2021a). We refer to the extensive reviews in Balasubramanian et al. (2014) for other applications to artificial intelligence.

Despite these attractive properties, the computation of conformal prediction sets is challenging for regression problems since an infinite number of models must be fitted with an augmented training set $\mathcal{D}_{n+1}(z)$, for all possible $z \in \mathbb{R}$. This is not only expensive, it is simply impossible in most cases. In general, efficiently computing conformal sets with the full data remains an open problem. The current successful approaches for calculating the set of conformal predictions are twofold.

- Exhaustive search with a *homotopy continuation*.² The fundamental idea is to rely on the fact that the typicalness function that maps each candidate with the rank of its prediction loss is piecewise constant. As follows, if we carefully manage to list all its transition points, we can find exactly where it is above the prescribed confidence level. For estimators that have a closed-form formula, e.g., Ridge (Hoerl, 1962) or Lasso (Tibshirani, 1996), it is possible to draw the solutions curve *w.r.t.* the input candidate z . They are often pieces of linear function which enable the exhaustive listing of the change points of the rank function; see Nourtdinov et al. (2001), and Lei (2019).
- *Inductive confidence machine* also called *Splitting* (Papadopoulos et al., 2002; Lei et al., 2018). The observed dataset is divided into two parts. A proper training set to fit the regression model, and an independent calibration set to calculate prediction losses, and ranks. This method is the most computationally efficient because it requires only a single model fit on a sub-part of the data. Separating the roles of the data to build the model, and to evaluate its performance avoids refitting without loss of coverage guarantees. The use of splitting techniques in statistics can be dated at least to Cox (1975)

These strategies have some noticeable limitations. The homotopy methods rely on strong assumptions on the model fit, and are numerically unstable due to multiple matrix inversions that are potentially poorly conditioned. They can suffer from exponential complexity

¹ Their joint probability distribution is invariant *w.r.t.* permutation of the data.

² Also called numerical continuation or path following methods (Allgower & Georg, 2012) are techniques for numerical approximation of a solution curve implicitly defined by a system of equations. In our context, we are interested in solving optimization problems, and we need to describe the solution curve *w.r.t.* parameter changes of the first order optimality conditions.

in the worst cases, and must frequently be abandoned because of extremely small step sizes (Gärtner et al., 2012; Mairal & Yu, 2012). The data splitting approach does not use all the data in the training phase. It generally results in a wider confidence region, *i.e.*, of wider size. As an alternative, a common heuristic unduly restricts the function evaluations to an arbitrary discrete grid of trial values z , and select the most typical one among them. These strategies might lose the coverage guarantee, and are still computationally inefficient. As a viable alternative, one relaxes the exact computation of the regression model at every step, and then approximately follows the homotopy continuation path by tightly controlling the optimization error as in Giesen et al. (2010), Ndiaye et al. (2019). Ndiaye and Takeuchi (2019) has shown this is a safer discretization strategy, and that it can cope with more general nonlinear regressions. Still, it is so far limited to convex problems with strong regularity assumptions on the model fit, and fails to be applicable to most machine learning prediction methods.

1.1 Summary of the contributions

We build on the striking remark that for common practical situations, the conformal prediction set is a bounded interval of the real line. Its boundaries are the roots of coverage level α minus the typicalness function, and these can be efficiently computed by a root-finding algorithm such as bisection search, with high precision, and without suffering from the limitations mentioned above. Despite its simplicity, it overcomes the limitations of the aforementioned strategies, and significantly improves, and extends the applicability of full conformal prediction to problems where it was considered intractable so far.

We highlight some advantages of our approach.

- *Efficiency* We demonstrate that computing a full conformal prediction set is tractable under mild assumptions. Relying on a bisection search, approximating the boundaries of the full exact conformal set at a prescribed accuracy $\epsilon > 0$, requires about $O(\log_2(1/\epsilon))$ number of model fit. The latter, trained on the whole data, allows to obtain a more informative confidence set than splitting methods. Accordingly, we maintain both statistical, and computational efficiency.
- *Flexibility* Our strategy offers considerable freedom on the choice of the regression estimator. For example, it can be defined as an output of a gradient descent process that maximizes a likelihood. It can be terminated when the norm of the gradient is smaller than a tolerance ϵ_0 or after 100 iterations of the algorithm. Consequently, the estimator can be parameterized by the number of iterations or the optimization error resulting from an iterative process as long as the symmetry of the data is preserved. The proposed root-finding approaches are easily applicable to more sophisticated recent machine learning techniques, such as deep neural networks or models involving a non-convex regularization.
- *Simplicity* The proposed methods are straightforward to implement. One substantially benefits from freely available scientific computing software packages like `scikit-learn` (Pedregosa et al., 2011) or `scipy` (Virtanen et al., 2020) to adjust models, and find the endpoints of the conformal set.

We also introduce an interpolation point of view of grid based approaches that properly justifies how the coverage guarantee can be maintained along with reduced computational time. In the case where a piecewise linear (or constant) interpolation scheme, and a simple

conformity score (for example the absolute value) is used, the assumption that the conformal set is an interval is not required. The computations can be easily carried following a homotopy strategy. To further reduce the number of model evaluations, we additionally provide a differentiable approximation of the rank function which effectively improves the computational efficiency of the root-finding solvers. We carefully analyze its coverage guarantee, and point out the trade-off between calibration, and number of model evaluations when such smoothing techniques are used. Such smoothing is mainly beneficial when a high precision is required.

Notation. For a non zero integer n , we denote $[n]$ the set $\{1, \dots, n\}$. We denote by $Q_{1-\alpha}$, the $(1 - \alpha)$ -quantile of a real valued sequence $(U_i)_{i \in [n+1]}$, defined as the variable $Q_{1-\alpha} = U_{(\lceil (n+1)(1-\alpha) \rceil)}$, where $U_{(i)}$ are the i -th order statistics. The interval $[a - \tau, a + \tau]$ will be denoted $[a \pm \tau]$. For an index j in $[n + 1]$, the rank of U_j among U_1, \dots, U_{n+1} is defined as $\text{Rank}(U_j) = \sum_{i=1}^{n+1} \mathbb{1}_{U_i \leq U_j}$.

2 Conformal prediction

We recall the arguments presented in Vovk et al. (2005), Shafer and Vovk (2008), Lei et al. (2018) while underlining in details the intuitions, and principles that sustain the construction, and validity of conformal prediction. Let us consider an input random variable X , and output Y . The goal is to construct a confidence set for the variable Y , *i.e.*, find a set $\mathcal{C}(X)$ such that

$$\mathbb{P}(Y \in \mathcal{C}(X)) \geq 1 - \alpha, \quad \forall \alpha \in (0, 1). \tag{1}$$

Given a prediction function $\mu(\cdot)$ that maps the input to the output space, and a loss measure S , one can assess the prediction error as $E = S(Y, \mu(X))$. It is a random variable with cumulative distribution function F and quantile Q defined as $F(z) = \mathbb{P}(E \leq z)$ and $Q(\delta) = \inf\{z \in \mathbb{R} : F(z) \geq \delta\}$. The main tool for building a set $\mathcal{C}(X)$ that satisfies the probabilistic guarantee in Equation (1), is the following classical result³:

$$\forall \delta \in (0, 1), \quad \mathbb{P}(F(E) \leq \delta) \geq \delta. \tag{2}$$

It implies $F(E) = F(S(Y, \mu(X))) \leq 1 - \alpha$ with probability larger than or equal to the confidence level $\delta = 1 - \alpha$. One then defines a confidence set for Y as the collection of candidate z that satisfy the same inequality, *i.e.*,

$$\mathcal{C}(X) = \{z : F(S(z, \mu(X))) \leq 1 - \alpha\}.$$

It turns out that the same principle can be applied to compute a confidence set for sequential observations. To do so, the coverage bound in Equation (2) can be extended to empirical cumulative distribution and empirical quantile functions defined as:

$$F_{n+1}(z) = \frac{1}{n+1} \sum_{i=1}^{n+1} \mathbb{1}_{E_i \leq z}, \quad Q_{n+1}(\delta) = \inf\{z \in \mathbb{R} : F_{n+1}(z) \geq \delta\}.$$

³ By definition, we have $Q(\delta)$ is the smallest real value z such that $F(z) \geq \delta$. Thus $\delta \leq F(Q(\delta)) = \mathbb{P}(E \leq Q(\delta)) = \mathbb{P}(F(E) \leq \delta)$.

Lemma 1 For a sequence of exchangeable random variables E_1, \dots, E_{n+1} , it holds $\mathbb{P}(F_{n+1}(E_{n+1}) \leq \delta) \geq \delta$, for any $\delta \in (0, 1)$.

Proof We follow the proof in Romano et al. (2019). By definition of the empirical quantile, we have

$$\delta \leq F_{n+1}(Q_{n+1}(\delta)) = \frac{1}{n+1} \sum_{i=1}^{n+1} \mathbb{1}_{E_i \leq Q_{n+1}(\delta)}.$$

Taking the expectation on both side, we have

$$\delta \leq \frac{1}{n+1} \sum_{i=1}^{n+1} \mathbb{P}(E_i \leq Q_{n+1}(\delta)) = \frac{1}{n+1} \sum_{i=1}^{n+1} \mathbb{P}(F_{n+1}(E_i) \leq \delta).$$

Moreover, we have for any i in $[n]$, $\mathbb{P}(F_{n+1}(E_i) \leq \delta) = \mathbb{P}(F_{n+1}(E_{n+1}) \leq \delta)$ by exchangeability. Hence the result. \square

Using Lemma 1, we have $F_{n+1}(E_{n+1}) \leq 1 - \alpha$ with probability larger than or equal to the confidence level $\delta = 1 - \alpha$. Given the n previous observations, one can define a confidence set for an unobserved variable E_{n+1} as the random set

$$\{z : F_{n+1}(z) \leq 1 - \alpha\}.$$

In supervised statistical learning problems, where we observe both the responses, and the features, one can apply this principle while taking benefits of an underlying model trained on the observed data. For the augmented dataset $\mathcal{D}_{n+1}(z) = \mathcal{D}_n \cup \{(x_{n+1}, z)\}$ for $z \in \mathbb{R}$, an example of predictive model is given as $\mu_z(x) = \Phi(x, \hat{\beta}(z))$, where Φ is a regression model, e.g., a kernel machine or a Deep Neural Network with parameter $\hat{\beta}(z)$ adjusted on the data. For example, by using empirical risk minimization principle, one defines

$$\hat{\beta}(z) \in \arg \min_{\beta \in \mathbb{R}^p} L(\beta \mid \mathcal{D}_{n+1}(z)) + \lambda \Omega(\beta), \quad (3)$$

where $\lambda > 0$ and $L(\beta \mid \mathcal{D}_{n+1}(z)) = \sum_{i=1}^n \ell(y_i, \Phi(x_i, \beta)) + \ell(z, \Phi(x_{n+1}, \beta))$ is the data fitting term and the regularization function Ω enforces structured solutions, e.g., sparsity.

Examples A popular example of an instance-wise loss function found in the literature is the *power norm*, where $\ell(a, b) = |a - b|^q$. When $q = 2$, this corresponds to classical linear regression. Cases where $q \in (0, 2)$ are common in robust statistics. In particular, $q = 1$ is known as least absolute deviation. The *logcosh* loss $\ell(a, b) = \gamma \log(\cosh(a - b)/\gamma)$ is a differentiable alternative to the ℓ_∞ -norm. One can also have the *Linex* loss function (Gruber, 2010; Chang & Hung, 2007) which provides an asymmetric loss $\ell(a, b) = \exp(\gamma(a - b)) - \gamma(a - b) - 1$, for $\gamma \neq 0$. The regularization functions Ω , e.g., Ridge (Hoerl & Kennard, 1970) or sparsity inducing norms (Bach et al., 2012; Obozinski & Bach, 2016) can be considered as well as non convex penalties (Xie & Huang, 2009).

Given the fitted model $\mu_z(\cdot)$ and a loss measure S , let us define the sequence of instance-wise prediction errors as:

$$\forall i \in [n], E_i(z) = S(y_i, \mu_z(x_i)), \text{ and } E_{n+1}(z) = S(z, \mu_z(x_{n+1})).$$

The sequence $\{E_1(y_{n+1}), \dots, E_n(y_{n+1}), E_{n+1}(y_{n+1})\}$ is exchangeable as long as the data $\{(x_i, y_i)\}_{i=1}^{n+1}$ is exchangeable, and the model fit $\mu_z(\cdot)$ is invariant *w.r.t.* permutation of the data. We can then apply Lemma 1 to obtain a coverage guarantee.

Definition 1 The full conformal prediction set is formally defined as

$$\Gamma^{(\alpha)}(x_{n+1}) = \{z : F_{n+1}(E_{n+1}(z)) \leq 1 - \alpha\}, \tag{4}$$

where

$$F_{n+1}(E_{n+1}(z)) = \frac{1}{n+1} \sum_{i=1}^{n+1} \mathbb{1}_{E_i(z) \leq E_{n+1}(z)}. \tag{5}$$

The term “full” refers to the fact that the entire data set is used to fit the regression model; in contrast to the splitting approach presented in detail below.

Lemma 1 implies that the set $\Gamma^{(\alpha)}(x_{n+1})$ is a valid confidence set for y_{n+1} in the sense of Equation 1, *i.e.*, $\mathbb{P}(y_{n+1} \in \Gamma^{(\alpha)}(x_{n+1})) \geq 1 - \alpha$ for any α in $(0, 1)$. Somehow, the refitting procedure with the extended dataset $\mathcal{D}_{n+1}(z)$, puts all the variables on equal feet, and preserves the exchangeability of the sequence of prediction errors. Using the rank function, we have $(n+1)F_{n+1}(E_{n+1}(z)) = \text{Rank}(E_{n+1}(z))$, and one can rewrite the CP set as (which is in fact the traditional notation)

$$\Gamma^{(\alpha)}(x_{n+1}) = \{z : \pi(z) \geq \alpha\},$$

where $z \mapsto \pi(z)$ is the typicalness function that measures how conformal a candidate is. It is defined as

$$\pi(z) = 1 - \frac{1}{n+1} \text{Rank}(E_{n+1}(z)) = 1 - F_{n+1}(E_{n+1}(z)). \tag{6}$$

Lemma 1 reads $\mathbb{P}(\pi(y_{n+1}) \leq \alpha) \leq \alpha$, *i.e.*, the random variable $\pi(y_{n+1})$ takes small values with small probability. Thus, it is unlikely that y_{n+1} will take the value z when $\pi(z)$ is small. More precisely, $\pi(y_{n+1})$ is (sub) uniformly distributed as usual for classical statistics for hypothesis testing. For example p -value function satisfies such a property under the null hypothesis; see (Lehmann & Romano, 2006, Lemma 3.3.1). One can then interpret the typicalness $\pi(\cdot)$ as a p -value function for testing the null hypothesis $H_0 : y_{n+1} = z$ against the alternative $H_1 : y_{n+1} \neq z$, for z in \mathbb{R} . The conformal prediction set merely corresponds to the collection of candidate z for which the null hypothesis H_0 is not rejected.

3 Computing conformal prediction set

For regression problem where y_{n+1} lies in a subset of \mathbb{R} , one need to evaluate $\pi(z)$ in Equation (6), and so refitting the model $\mu_z(\cdot)$ for infinitely many candidate z . This merely renders the overall computation challenging, and leaves the problem open in general. Nevertheless, some peculiar regularity structure of the typicalness function $\pi(\cdot)$ can be exploited. For example, by utilizing the fact that it is piecewise constant, it is sufficient to enumerate the transition points (when they are finite) to compute the conformal set. This is possible for

a limited number of cases, *e.g.*, Ridge or Lasso where the map $z \mapsto \mu_z(\cdot)$ can be explicitly described. Unfortunately, only a very small class of statistical learning problems has such nice regularity structure. For other estimators, the computation of CP set when y_{n+1} can take countless number of values, is unclear.

3.1 Splitting

To overcome this issue, the split conformal prediction set introduced in Papadopoulos et al. (2002), separates the model fitting, and the score ranking step. Let us define

- the training set $\mathcal{D}_{\text{tr}} = \{(x_1, y_1), \dots, (x_m, y_m)\}$ with $m < n$,
- the calibration set $\mathcal{D}_{\text{cal}} = \{(x_{m+1}, y_{m+1}), \dots, (x_n, y_n)\}$.

Then the model is fitted on the training set \mathcal{D}_{tr} to get $\mu_{\text{tr}}(\cdot)$, and define the score function on the calibration set \mathcal{D}_{cal} :

$$\forall i \in [m + 1, n], E_i^{\text{cal}} = S(y_i, \mu_{\text{tr}}(x_i)), \text{ and } E_{n+1}^{\text{cal}}(z) = S(z, \mu_{\text{tr}}(x_{n+1})).$$

Thus, we obtain the split typicalness function as

$$\pi_{\text{split}}(z) = 1 - F_{\text{split}}(E_{n+1}^{\text{cal}}(z)), \text{ where}$$

$$F_{\text{split}}(E_{n+1}^{\text{cal}}(z)) = \frac{1}{n - m + 1} \sum_{i=m+1}^{n+1} \mathbb{1}_{E_i^{\text{cal}} \leq E_{n+1}^{\text{cal}}(z)}.$$

The latter is proportional to the rank of the $(n + 1)$ th score on the calibration set. Finally, we define

$$\Gamma_{\text{split}}^{(\alpha)}(x_{n+1}) = \{z : \pi_{\text{split}}(z) \geq \alpha\} = \{z : E_{n+1}^{\text{cal}}(z) \leq Q_{1-\alpha}^{\text{cal}}\},$$

where $Q_{1-\alpha}^{\text{cal}}$ is the $(1 - \alpha)$ quantile of the calibration scores $\{E_{m+1}^{\text{cal}}, \dots, E_{n+1}^{\text{cal}}\}$. When the score function is the absolute value $S(a, b) = |a - b|$, the split CP set is the interval $\Gamma_{\text{split}}^{(\alpha)}(x_{n+1}) = [\mu_{\text{tr}}(x_{n+1}) \pm Q_{1-\alpha}^{\text{cal}}]$. While this approach avoids the computational bottleneck, the statistical efficiency of the model can be reduced due to a significantly smaller sample available during the training, and calibration phase. Moreover, the length of the split conformal set tends to have a higher variance. In general, the proportion of training vs calibration set is a hyperparameter that requires appropriate tuning: a small calibration set leads to highly variable conformity scores, and a small training set leads to poor model fitting. In all our experiments, we set the splitting proportion to 2, which means that the two sets play symmetric roles. Since the sequence of scores $\{E_{m+1}^{\text{cal}}, \dots, E_n^{\text{cal}}, E_{n+1}^{\text{cal}}(z)\}$ is exchangeable, the Lemma 1 implies that $\mathbb{P}(y_{n+1} \in \Gamma_{\text{split}}^{(\alpha)}(x_{n+1})) \geq 1 - \alpha$.

3.2 Cross-conformal predictors

The trade-off mentioned above is very recurrent in machine learning, and often appears in the debate between bias reduction, and variance reduction. It is often decided by the cross-validation method with several folds (Arlot & Celisse, 2010). *Cross-conformal predictors* (Vovk, 2015) follow the same ideas, and exploit the full dataset for calibration, and significant proportions for training the model. The dataset is partitioned into K folds, and

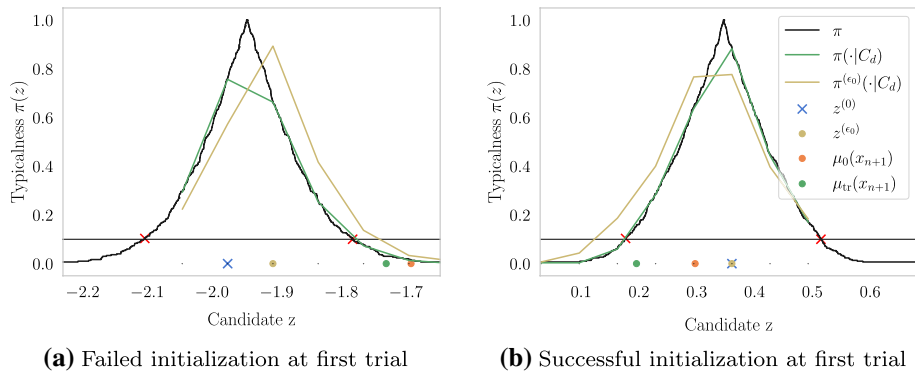


Fig. 1 Illustration of the initialization steps when both the initial prediction based on observed data $z_0 = \mu_{D_o}(x_{n+1})$, and midpoint of split conformal interval $\mu_{tr}(x_{n+1})$ fails to be in the conformal prediction set whose boundaries are delimited by the red crosses. The synthetic data are generated with `sklearn` as $X, y = \text{make_regression}(n = 300, p = 50)$. We choose an optimization accuracy of $\epsilon_0 = \|(y_1, \dots, y_n)\|_2^2 / 10^4$ for approximating the ridge estimator. The trial points are $C_d = \{z_1, \dots, z_d\}$ with $d = 10$, and we denoted $z^{(\epsilon)} \in \arg \max_{z \in C_d} \pi^{(\epsilon)}(z)$ is the most conformal trial candidate at precision $\epsilon \geq 0$. To be more explicit, the approximated conformity functions obtained from the C_d grid are denoted $\pi^{(\epsilon_0)}(\cdot | C_d)$ when early stopping at optimization accuracy ϵ_0 is used, and $\pi(\cdot | C_d)$ when the exact solution is used

one performs a split conformal set by sequentially defining the k th fold as calibration set, and the remaining as training set for $k \in \{1, \dots, K\}$. However, aggregating the different pi-values is not straightforward, and the validity of the method might be jeopardized without stronger assumptions on the score function, see (Carlsson et al., 2014; Linusson et al., 2017). More precisely, it can be shown that the confidence level is inflated by a factor of 2, i.e., the (not improvable) theoretical coverage level is $1 - 2\alpha$ instead of $1 - \alpha$, see (Barber et al., 2021). Under additional stability assumption, Cross-conformal predictors can only *approximately* achieve the target coverage $1 - \alpha$. Otherwise, without approximation, in order to remove the factor 2, one can consider an overly conservative set whose extremity are defined as the smallest, and largest residual over all leave-one-out residuals. The leave-one-out (also called Jackknife) CP set, will require $K = n$ model fit which is prohibitive even when n is moderately large. On the other hand, the K -fold version will require K model fit but will come at the cost of fitting on a lower sample size, and will leads to an additional excess coverage of $O(\sqrt{2/n})$. A Bootstrap version (Vovk, 2015, Appendix B) will suffer from the same inflation (Kim et al., 2020). In all cases, we are not aware of a (variant of) cross-conformal predictors that simultaneously achieves $1 - \alpha$ provable coverage guarantee, and a non-conservative prediction set. Nevertheless, the practical performance is fairly acceptable both computationally, and statistically. In this paper, we only compare with the methods that provably achieve the prescribed $1 - \alpha$ confidence level, namely the splitting method, and the oracle conformal prediction described in Sect. 4.

3.3 Approximation to a prescribed accuracy

In this paper, we promptly take advantage of the remarkable fact that the conformal regions are *often* intervals. We subsequently take an alternative direction which carefully avoids tracking the integral path of all model fit, and also avoids any data splitting. When the $(1 - \alpha)$ -level set of the function $z \mapsto F_{n+1}(E_{n+1}(z))$ is convex, e.g., Fig. 1, we propose to

employ a numerical root-finding solver to approximate the endpoints of the interval. The statistical validity is automatic, we obtain simultaneously an upper, and lower bound on each extremity of the confidence set, and the approximation error ϵ can be made arbitrarily small at the cost of $O(\log(1/\epsilon))$ number of model fit; without inflation of the confidence level.

3.4 Outline of the algorithm: rootCP

Assuming that the conformal set is a non empty interval of finite length, we denote

$$I^{(\alpha)}(x_{n+1}) = [\ell_\alpha(x_{n+1}), u_\alpha(x_{n+1})].$$

Given a tolerance $\epsilon > 0$, we proceed as follows:

1. find $z_{\min} < z_0 < z_{\max}$ such that

$$\pi(z_{\min}) < \alpha < \pi(z_0) \text{ and } \alpha > \pi(z_{\max}). \tag{7}$$

2. Perform a bisection search in $[z_{\min}, z_0]$. It will output a point $\hat{\ell}$ such that $\ell_\alpha(x_{n+1})$ belongs to $[\hat{\ell} \pm \epsilon]$ after at most $\log_2(\frac{z_0 - z_{\min}}{\epsilon})$ iterations.
3. Perform a bisection search in $[z_0, z_{\max}]$. It will output a point \hat{u} such that $u_\alpha(x_{n+1})$ belongs to $[\hat{u} \pm \epsilon]$ after at most $\log_2(\frac{z_{\max} - z_0}{\epsilon})$ iterations.

3.5 Initialization

For the initial lower, and upper bounds, we suggest

$$z_{\min} = \min_{i \in [n]} y_i \text{ and } z_{\max} = \max_{i \in [n]} y_i.$$

For most of the situations encountered in our numerical experiments, we consistently get $\pi(z_{\min})$, and $\pi(z_{\max})$ both smaller than the threshold level α . Otherwise, we can always take values even farther apart without affecting the complexity thanks to the logarithmic dependence in the length of the initialization brackets. This is especially necessary when the total number of samples n is small⁴. The most crucial part is to choose z_0 so that $\pi(z_0) > \alpha$. It is equivalent to get a point in the interior of the conformal set itself. In the ideal case where the length of the conformal set is extremely small, finding an initialization point might be notoriously hard. Indeed, it corresponds to a *rare event* equivalent to sampling a point in a low probability region. We adopted a simple strategy which consists in estimating y_{n+1} with the observed data \mathcal{D}_n . We subsequently denote it

$$z_0 = \mu_{\mathcal{D}_n}(x_{n+1}).$$

In our sequence of repetitive numerical experiments, this choice rarely fails. Naturally, its success depends on the prediction capabilities of the model fit. In the rare cases where it

⁴ Indeed, we have $\mathbb{P}(y_{n+1} \in [z_{\min}, z_{\max}]) \geq 1 - \frac{1}{n+1}$. Hence when n is sufficiently large, *i.e.*, $n \geq 1 + 1/\alpha$, then $[z_{\min}, z_{\max}]$ is a $(1 - \alpha)$ confidence set.

fails, we propose to test the initialization condition on some query points selected on an initial estimation $[z_\alpha^-, z_\alpha^+]$ of the CP set. This localization step aims to exploit additional problem structure, and can be interpreted as an iterative importance sampling to maintain a reasonably low computational cost.

1. *Localization* Given an easy to compute estimate set $[z_\alpha^-, z_\alpha^+]$ that is potentially larger⁵ than the targeted conformal set, we select its mid point

$$z_0 = \frac{z_\alpha^+ + z_\alpha^-}{2}.$$

If z_0 satisfies $\pi(z_0) > \alpha$, then we have a valid initialization by paying only a single model fit. Otherwise, we run the next step on the bracket $[z_\alpha^-, z_\alpha^+]$.

For instance, one can use the interval obtained from the splitting approach $[z_\alpha^-, z_\alpha^+] = \Gamma_{\text{split}}^{(\alpha)}(x_{n+1})$ or a rough approximation $[z_\alpha^-, z_\alpha^+] = \{z : \pi_{\mathcal{D}_n}(z) > \alpha\}$ where $\pi_{\mathcal{D}_n}(\cdot)$ is an unsafe estimation of $\pi(\cdot)$ with $\mu_z(x)$ replaced by $\mu_{\mathcal{D}_n}(x)$ for any candidate z , and any input feature x .

2. *Sampling* For a small number d , e.g., $d = 5$, and given a bracket search $[z^-, z^+]$, select candidates $C_d = \{z_1, \dots, z_d\}$ uniformly. If there is $\pi(z_0) > \alpha$ for a z_0 in C_d , we have a valid initialization. Otherwise, we use these query points to interpolate the model fit as in eq. (9). Thus, by selecting additional points that have a higher typicalness according to the interpolated model, one can refine the sampling set C_d , and repeat the process.

For completeness, we summarize the procedure in Algorithm 1.

Algorithm 1 rootCP using bisection search Algorithm 2

Input: data $\mathcal{D}_n = \{(x_1, y_1), \dots, (x_n, y_n)\}$, and x_{n+1}

Coverage level $\alpha \in (0, 1)$, accuracy $\epsilon > 0$

Output: approximation of $\Gamma_{\text{up}}^{(\alpha)}(x_{n+1}) = [\ell_\alpha(x_{n+1}), u_\alpha(x_{n+1})]$ up to accuracy ϵ .

Initialization

Set $z_{\min} = y_{(1)}$ and $z_{\max} = y_{(n)}$

Set $z_0 = \mu_0(x_{n+1})$ where we fitted a model μ_0 on the training data \mathcal{D}_n

if $z_{\min} < z_0 < z_{\max}$ such that

$$\pi(z_{\min}) < \alpha < \pi(z_0) \text{ and } \alpha > \pi(z_{\max}) \tag{8}$$

then

Approximation of Boundary Points

1. $\hat{\ell} = \text{bisection_search}(\pi - \alpha, z_{\min}, z_0, \epsilon)$

2. $\hat{u} = \text{bisection_search}(\pi - \alpha, z_0, z_{\max}, \epsilon)$

Return: $[\hat{\ell}, \hat{u}]$

else

Initialization Failure

Perform a heuristic search in a grid point $C_d = \{z_1, \dots, z_d\}$ until Equation (8) is met

Go to root-finding steps 1. and 2.

end if

Return: *Global Failure*: refer to alternative methods (e.g., splitCP or interp1CP)

⁵ One can slightly enlarge the set by taking $[z_\alpha^- - 0.5(z_\alpha^+ - z_\alpha^-), z_\alpha^+ + 0.5(z_\alpha^+ - z_\alpha^-)]$.

Algorithm 2 `bisection_search`

Input: Function f , scalars $a < b$ such that $\text{sign}(f(a)) \neq \text{sign}(f(b))$, accuracy $\epsilon > 0$
Set $\text{max_iter} = \log_2(\frac{b-a}{\epsilon})$
for $i = 1$ **to** max_iter **do**
 $c = \frac{a+b}{2}$
 if $f(c) = 0$ or $c - a \leq \epsilon$ **then**
 Return: c
 end if
 if $f(c)f(a) < 0$ **then**
 $b = c$
 else
 $a = c$
 end if
end for
Return: the algorithm did not converge, increase max_iter .

Additionally, we explain below how this model fit interpolation can be used to obtain an alternative CP set. Note that its midpoint can also be used as a candidate for initialization. For computational efficiency, one can rely predominantly on the fact that for the usual prediction problems in machine learning, it is unneeded to optimize below the inevitable statistical error (Bousquet & Bottou, 2008). This means that a high optimization accuracy in the model fit might be unnecessary to achieve better generalization performances. Therefore, with a coarse optimization tolerance, we can preview the final shape of the conformity function. Whence, one can replace $\pi(z)$ by $\pi^{(\epsilon_0)}(z)$ which is computed with a rough optimization error ϵ_0 in order to guess the shape of the function $\pi(\cdot)$. Similarly, if $\pi(z_0)$ fails to be a valid initialization, then decrease ϵ_0 , and repeat the process. In all our experiments, it works fine after a very few number of iterations. Nonetheless, we do not have strategies to avoid worst-case situations or any mathematical guarantee in the total number of iterations needed to find a valid initialization. We illustrate this strategy in Fig. 1 in both situations of failure, and success.

3.6 Further complexity reduction

In cases where the regression map $z \mapsto \mu_z(x)$ for any feature x , can be traced with homotopy as in Ridge (Noureddinov et al., 2001), and Lasso (Lei, 2019), it takes $O(n^2)$ to compute the exact conformal set. This can be reduced to $O(n \log n)$ by sorting the roots of the instance-wise scores $E_i(z) - E_{n+1}(z)$ for i in $[n]$, and cleverly flattening the double loop when evaluating the ranks of the score functions (Vovk et al., 2005, Chapter 2.3). By relaxing the exactness, none of these two steps is needed in our approach. We obtain an asymptotic improvement to $O(n \log_2(1/\epsilon))$, and an easier to implement algorithm.

When the model fit is parameterized by the solution of optimization problem in Equation (3), the regularity of the loss function, and penalty terms play a major role in the computational tractability of the full conformal prediction set. Leveraging smoothness, and convexity assumptions on the loss or penalty functions, it has been shown in Ndiaye and Takeuchi (2019) that approximate solutions can be used without refitting the model for close candidates z . The resulting conformal set is $\bar{\Gamma}^{(\alpha, \epsilon)} = \{z \in \mathbb{R} : \bar{\pi}(z, \epsilon) > \alpha\}$ where the corresponding typicalness function incorporates the optimization error, *i.e.*,

$$\bar{\pi}(z, \epsilon) = \frac{1}{n + 1} \sum_{i=1}^{n+1} \mathbb{1}_{E_i(z) \geq E_{n+1}(z) - 2\sqrt{2ve}}$$

One can further show that the typicalness function based on exact solution $\hat{\pi}(\cdot)$ is uniformly upper bounded by $\bar{\pi}(\cdot, \epsilon)$, and then $\hat{F}^{(\alpha)} \subset \bar{F}^{(\alpha, \epsilon)}$. This can be equally used to reduce the number of model fit, and also to wrap the CP set based on exact solution by applying `rootCP` directly to $\bar{\pi}(\cdot, \epsilon)$ instead of computing a whole approximation path.

Compared with the homotopy approach, `rootCP` will always make a smaller number of model fits. By merely storing each model evaluation, it benefits from warm-start boosting by employing the solutions of the previous function call. Alas, the homotopy approaches require either an exact solution or an approximate one with a strict control of the optimization error. This control is not always available if one does not provide a computable upper bound of it, for example by precisely evaluating the duality gap. Such bounds are hardly available in non-convex settings which greatly reduce its applicability in modern machine learning techniques. Meanwhile, the complexity of the approximate homotopy algorithm is $O(\frac{z_{\max} - z_{\min}}{\sqrt{\epsilon_0}})$ where ϵ_0 is the optimization error of the model fit. Additionally to the linear dependence in the initial interval length, it cannot be launched for small ϵ_0 whereas the number of model fit in the root-finding approach is not degraded. In a nutshell, the proposed method avoids the computation of the whole path. Hence, it enjoys an exponential improvement over the homotopy approach *w.r.t.* to the initial interval length $z_{\max} - z_{\min}$, and an overall complexity that is independent of ϵ_0 . It can then be used with highly optimized model fit where the homotopy method cannot even be launched.

3.7 Drawbacks

Full conformal prediction set is *not always* an interval. When it is a union of few well separated intervals, our proposed method cannot be applied without finely bracketing these intervals. One can include a human in the loop. The discrete function $\pi^{(\epsilon_0)}$ offers a cheap pre-visualization of the landscape of the conformity function that allows to detect these situations, and infer a proper bracketing. At this point, efficiently enumerating all the roots remains a challenging task that we leave as an open problem. In the following proposition, we provide a sufficient condition so that the conformal set is an interval. It essentially consists of a simple condition so that the conformity function is monotonically increasing until it reaches its maximum value, and then monotonically decreasing.

Proposition 1 *If for any i in $[n]$, the difference of instance-wise error function $z \mapsto E_i(z) - E_{n+1}(z)$ is quasi-concave, and has two zeros $a_i \leq b_i$ such that*

$$\max_{i \in [n]} a_i \leq \min_{i \in [n]} b_i,$$

then the typicalness function $\pi(\cdot)$ is quasi-concave, and the conformal prediction set at a level $\alpha \in (0, 1)$ is either empty or an interval.

Proof The function $\psi_i(z) := E_i(z) - E_{n+1}(z)$ is quasi-concave implies that its 0-level set is convex, and $\{z : \psi(z) \leq 0\} = (-\infty, a_i] \cup [b_i, +\infty)$ or empty. Whence,

$$F_{n+1}(E_{n+1}(z)) = \frac{1}{n+1} \sum_{i=1}^{n+1} \mathbb{1}_{(-\infty, a_i] \cup [b_i, +\infty)}(z),$$

where, without loss of generality, we assume that all intervals are non-empty. Those that are not have a zero contribution to the sum and can therefore be deleted. Now, the condition $\max_{i \in [n]} a_i \leq \min_{i \in [n]} b_i$ implies that the function $z \mapsto \pi(z) = 1 - F_{n+1}(E_{n+1}(z))$ is monotonically increasing in $(-\infty, a_{(n+1)}]$, and decreasing in $[b_{(1)}, +\infty)$. Hence the result. \square

Proposition 1 generalizes (Lei, 2019, Theorem 3.3) which provides a sufficient condition so that the Lasso conformal set is an interval. Unfortunately, such sufficient conditions are not testable for most problems. Indeed, it requires knowing all the zero crossing points of the function $z \mapsto E_i(z) - E_{n+1}(z)$ for all indices $i \in [n]$ which is as hard as computing the whole function $z \mapsto \pi(z)$.

In the literature, similar assumptions are made to obtain a conformal predictive distribution by enforcing a monotonicity on the score functions, see Vovk et al. (2017). We remind that even when the typicalness function $z \mapsto \pi(z)$ is not quasi-concave, our algorithm is still valid as long as the conformal set is an interval; which is a much weaker assumption than quasi-concavity. It would be interesting to study in more detail how to characterize the class of score function that systematically leads to a CP set being an interval. This is not necessarily obvious since it is easy to build an adversarial example. Indeed, consider a sinusoidal score function $S(a, b) = |\sin(a - b)|$. It treats the data symmetrically across the instances, and then satisfies all the assumptions. The corresponding CP set is a union of an infinite number of intervals even for very simple regression models such as least squares. Which is a reminiscence of No-Free-Lunch: without any assumption, computations of a CP set is impossible, even with splitting approaches!

3.8 Interpolated conformal prediction

The full conformal prediction set is computationally expensive since it requires knowing exactly the map $z \mapsto \mu_z(\cdot)$. The splitting approach does not use all the data in the learning phase but is computationally efficient since it requires a single model fit. Alternatively, it was proposed in Lei et al. (2018) to use an arbitrary discretization, and its theoretical analysis in Chen et al. (2018) unfortunately failed to preserve the coverage guarantee. In this section, we argue that grid based strategy with an interpolation point of view, stands as an "in-between" strategy that exploits full data with a restricted computational time while preserving the coverage guarantee. We propose to compute a conformal prediction set based on an interpolation of the model fit map given a finite number of query points. The main insight is that the underlying model fit plays a minor role in the coverage guarantee; the only requirement is to be symmetric with respect to permutation of the data. As such, the model path $z \mapsto \hat{\mu}_z(\cdot)$ can be replaced by an interpolated map $z \mapsto \tilde{\mu}_z(\cdot)$ based on query

points z_1, \dots, z_d . It reads to a valid prediction set as long as the interpolation preserves the symmetry⁶. For instance, one can rely on a piecewise linear interpolation

$$\tilde{\mu}_z = \begin{cases} \frac{z_1-z}{z_1-z_{\min}} \hat{\mu}_{z_{\min}} + \frac{z_{\min}-z}{z_{\min}-z_1} \hat{\mu}_{z_1} & \text{if } z \leq z_{\min}, \\ \frac{z_1-z_{\min}}{z_{t+1}-z_{\min}} \hat{\mu}_{z_t} + \frac{z-z_t}{z_{t+1}-z_t} \hat{\mu}_{z_{t+1}} & \text{if } z \in [z_t, z_{t+1}], \\ \frac{z_t-z_{t+1}}{z_{\max}-z_d} \hat{\mu}_{z_{\max}} + \frac{z_{\max}-z}{z_{\max}-z_d} \hat{\mu}_{z_d} & \text{if } z \geq z_{\max}, \end{cases} \tag{9}$$

where $\hat{\mu}_z(x)$ is a prediction map trained on the augmented dataset $\mathcal{D}_{n+1}(z)$ as in Equation (3).

As before, one defines the instance-wise score functions

$$\forall i \in [n], \tilde{E}_i(z) = S(y_i, \tilde{\mu}_z(x_i)) \text{ and } \tilde{E}_{n+1}(z) = S(z, \tilde{\mu}_z(x_{n+1})).$$

The conformal set based on interpolated model fit is then defined as

$$\tilde{I}^{(\alpha)}(x_{n+1}) = \{z : \tilde{\pi}(z) \geq \alpha\}, \text{ where}$$

$$\tilde{\pi}(z) = 1 - \frac{1}{n+1} \sum_{i=1}^{n+1} \mathbf{1}_{\tilde{E}_i(z) \leq \tilde{E}_{n+1}(z)}.$$

Since the map $z \mapsto \tilde{\mu}_z(\cdot)$ is (or can be made) symmetric, it is immediate to see that $\tilde{I}^{(\alpha)}(x_{n+1})$ is a valid conformal set following the same proof technique.

We remind that the conformal set can be highly concentrated around its midpoint, and the typicalness of most candidates is close to zero. Whence, we suggest restricting the query points around an estimate of the conformal set provided by a localization step. Also, it could be interesting to evaluate the performance of more sophisticated interpolation methods like splines in order to have more symbiosis between the interpolation, and the smoothing of the rank function introduced in the next section. Nevertheless, the simplicity of linear interpolation allows an exact calculation of the conformal prediction set because one can easily enumerate the change points of the rank function. This is not necessarily preserved with higher order interpolation, and requires further investigation.

Remark 1 (Interpolation of the typicalness map) Given the query points, and their corresponding typicalness $(z_1, \pi(z_1)), \dots, (z_d, \pi(z_d))$, one can also directly learn a function that approximate the typicalness $z \mapsto \pi(z)$. However, in this case, we could not establish the theoretical coverage guarantee of this method. Moreover, when the conformal set is highly localized, most of the $\pi(z_i)$ might be close to zero leading to a flat, and poorly interpolated typicalness map.

Previous discretization approaches did not preserve the coverage guarantee or did it at expensive cost by approximating a model fit path on a wide range, with a high precision at every step, and was restricted to convex problems. The interpolation point of view that we provided allows us to compute a valid conformal prediction set with arbitrary discretization without loss in the coverage guarantee, and without restriction to convex problems. Also note that depending on the interpolation used, there is no need to assume that the

⁶ Otherwise one can always perform a symmetrization, e.g., using a model parameter $\tilde{\beta}(z) = \frac{1}{(n+1)!} \sum_{\sigma \in \Sigma_{n+1}} \beta(w_{\sigma(1)}, \dots, w_{\sigma(n)}, w_{\sigma(n+1)})$, where $w_i = (x_i, y_i)$ if i in $[n]$, $w_{n+1} = (x_{n+1}, z)$, and Σ_{n+1} is the group of permutation of $[n+1]$.

conformal prediction set is an interval. Indeed, in the case below of piecewise linear interpolation, one can easily enumerate all the change points of the conformal function as in homotopy methods. However, in general we also recommend the use of `rootCP`. Finally, note that in the case of the ridge estimator (which is linear in z), the exact conformal prediction coincides with that of the interpolation.

3.9 Smoothed conformal prediction

Conformal prediction sets rely on rank computations. The latter function is piecewise constant, and has no useful first order information in the sense that it is either null or undefined. We propose a smooth approximation of the typicalness function to reduce the number of query points. In addition to exchangeability, we merely use the fact that F_{n+1} is increasing, and the linearity of the sum to obtain the coverage guarantee. Likewise, one should be able to replace $\mathbb{1}_{E_i - z \leq 0}$ with a continuously differentiable, and increasing function $\phi_\gamma(E_i - z)$. Hence, replacing the function $z \mapsto F_{n+1}(E_{n+1}(z))$ by a smoother one allows the use of more efficient gradient or quasi-Newton-based root finding methods. We further investigate the influence of such smoothing on the coverage guarantee. In practice, we simply choose the sigmoid function $\phi_\gamma(x) = \frac{e^{-\gamma x}}{1 + e^{-\gamma x}}$ as in Qin et al. (2010). We have

$$\text{Rank}(u_j) := \sum_{i=1}^{n+1} \mathbb{1}_{u_i - u_j \leq 0} \approx \sum_{i=1}^{n+1} \phi_\gamma(u_i - u_j) =: \text{sRank}(u_j, \gamma).$$

The main advantage is that the map $\text{sRank}(\cdot, \gamma)$ improves the regularity of $\text{Rank}(\cdot)$, and allows faster convergence.

The smooth approximation of the typicalness function is then defined as

$$\pi(z) \approx \pi(z, \gamma) := 1 - \frac{1}{n+1} \text{sRank}(E_{n+1}(z), \gamma),$$

and the smoothed conformal prediction set (illustrated in Fig. 2) as

$$\Gamma^{(\alpha, \gamma)}(x_{n+1}) = \{z : \pi(z, \gamma) > \alpha\}.$$

Now computing an approximation of the conformal prediction set is equivalent to finding the smallest, and largest solution of the equation $\pi(z, \gamma) = \alpha$ which is often easier to solve than $\pi(z) = \alpha$. Using different root-finding solvers, we illustrate the computational advantages by displaying the reduction of the number of model fit in Fig. 2.

Remark 2 (Gradient Based Solvers) When, for any feature x , the regression map $z \mapsto \mu_z(x)$ is differentiable, the solutions of equation $\pi(z, \gamma) = \alpha$ could be approximated with more efficient gradient based root-finding algorithm. However, the function $z \mapsto \pi(z, \gamma)$ is mostly flat except at a tiny vicinity of the conformal set which makes the convergence difficult unless a good initialization is found. One could also rely on a regularized version by minimizing $(\pi(z, \gamma) - \alpha)^2 + \tau z^2$ which requires a proper tuning of the hyper parameter τ . Both of these strategies turn out to be less stable, and need further investigations.

Remark 3 Here it is clear that the term “smooth” refers to the differentiability of the conformity function; and was introduced for computational reasons. This is not to be confused with the “Smoothed Conformal Predictors” introduced in Vovk et al. (2005, Page 27) where the borderline cases $E_i(z) = E_{n+1}(z)$ are treated more carefully, *i.e.*, smoothly

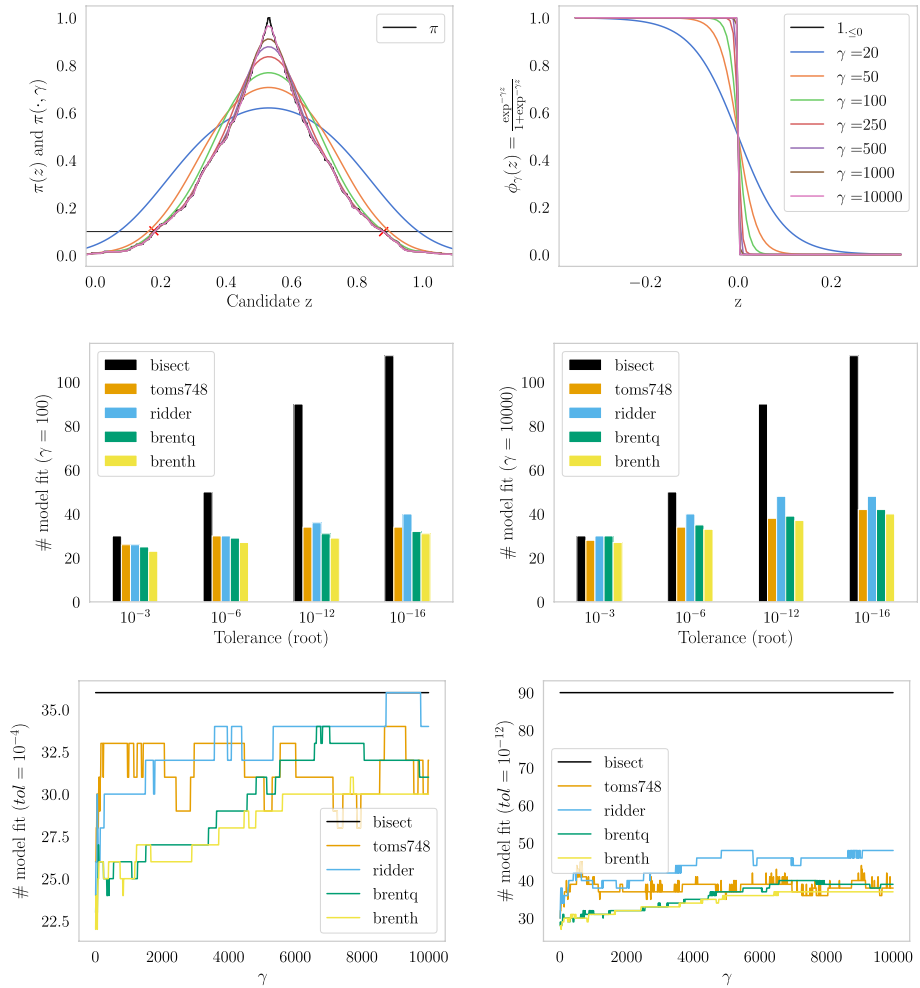


Fig. 2 Illustration of the smoothed conformal set with data generated from sklearn as $X, y = \text{make_regression}(n = 300, p = 50)$. The smoothed typicalness function $\pi(\cdot, \gamma)$ is evaluated with several values for the hyperparameter γ . The underlying estimator is the ridge regressor with parameter $\lambda = p / \|\beta_{LS}\|^2$ where β_{LS} is the Least-squares estimator on the observed dataset \mathcal{D}_n

penalized with a random parameter between 0, and 1 instead of increasing the rank with 1. This randomization essentially breaks the ties to ensure an exact coverage guarantee. All the computational methods we introduce in this article apply immediately to this case.

We analyze the statistical consequences of using a continuous version of the indicator function. We recall the definition of the smoothed version of the empirical cumulative distribution, and empirical quantile:

$$\tilde{F}_{n+1}(z) = \frac{1}{n+1} \sum_{i=1}^{n+1} \phi_\gamma(E_i - z), \quad \tilde{Q}_{n+1}(\alpha) = \inf\{z \in \mathbb{R} : \tilde{F}_{n+1}(z) \geq \alpha\}.$$

Proposition 2 (Coverage guarantee of the smooth relaxation) *For a sequence of exchangeable random variable E_1, \dots, E_{n+1} , it holds for any $\tilde{\alpha}$ in $(0, 1)$,*

$$\mathbb{P}(\tilde{F}_{n+1}(E_{n+1}) \leq \tilde{\alpha}) \geq \tilde{\alpha} - \Delta(\gamma),$$

where $\Delta(\gamma) = \sup_x(\phi_\gamma - \mathbb{1}_{\leq 0})(x)$.

Proof By definition of \tilde{F}_{n+1} , and \tilde{Q}_{n+1} , we have

$$\begin{aligned} \tilde{\alpha} \leq \tilde{F}_{n+1}(\tilde{Q}_{n+1}(\tilde{\alpha})) &= \frac{1}{n+1} \sum_{i=1}^{n+1} \phi_\gamma(E_i - \tilde{Q}_{n+1}(\tilde{\alpha})) \\ &= \frac{1}{n+1} \sum_{i=1}^{n+1} \mathbb{1}_{E_i \leq \tilde{Q}_{n+1}(\tilde{\alpha})} + (\phi_\gamma - \mathbb{1}_{\leq 0})(E_i - \tilde{Q}_{n+1}(\tilde{\alpha})) \\ &\leq \frac{1}{n+1} \sum_{i=1}^{n+1} \mathbb{1}_{\tilde{F}_{n+1}(E_i) \leq \tilde{\alpha}} + \Delta(\gamma). \end{aligned}$$

We conclude by taking the expectation on both sides along with exchangeability. □

In order to display a probabilistic statement, one needs to maintain the indicator function when defining the typicalness function. Replacing it with a continuous version will distort the coverage guarantee as described in Proposition 2. To obtain a well α -calibrated confidence set, one must take into account such approximation error by choosing $(\tilde{\alpha}, \gamma)$ such that

$$\tilde{\alpha} - \Delta(\gamma) \geq \alpha.$$

If $\tilde{\alpha}$ is fixed, one needs to be careful when choosing γ . Otherwise, we obtain a vacuous upper bound, and all the coverage guarantee is lost. Meanwhile, if γ is chosen such that ϕ_γ is a lower approximation of the indicator function, then $\tilde{\alpha}$ can be taken as α , and there is no calibration loss. However, when $\Delta(\gamma)$ is close to zero, ϕ_γ will be flat almost everywhere, and we will not get useful first order information. This brings a *trade-off* between number of model fitting (which influences the computational time), and efficiency, *i.e.*, length of the interval (wider $\tilde{\alpha}$ -level set).

Building a gap. To finely assess how the vanilla conformal, and smoothed conformal set can be related in practice, one can simply design both a lower, and upper approximation of the indicator function, *i.e.*, ϕ_γ^+ and ϕ_γ^- .

In that case, it is easy to see that

$$\ell_\gamma^+ \leq \ell_\alpha(x_{n+1}) \leq \ell_\gamma^- \text{ and } u_\gamma^- \leq u_\alpha(x_{n+1}) \leq u_\gamma^+,$$

which is equivalent to

$$\Gamma_\gamma^{(\alpha,-)}(x_{n+1}) \subset \Gamma^{(\alpha)}(x_{n+1}) \subset \Gamma_\gamma^{(\alpha,+)}(x_{n+1}).$$

The overall complexity is moderately expanded (we now need to compute two different conformal prediction sets), and not too time consuming as long as the underlying model fit is reasonably computable.

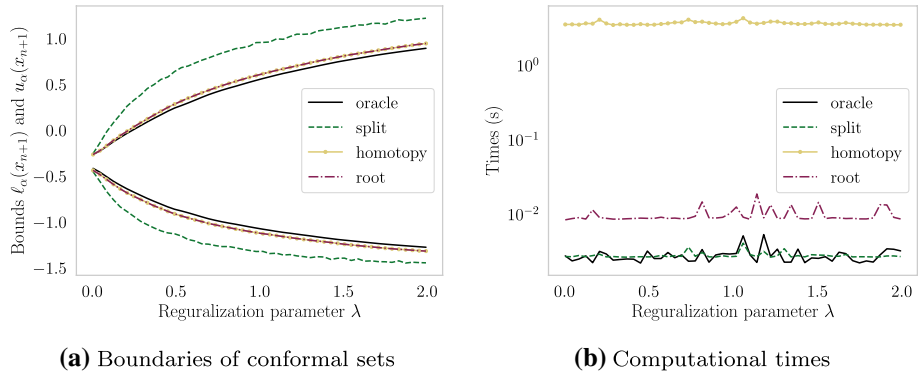


Fig. 3 Benchmark on ridge regression. Conformal prediction set computed with various regularization parameter on synthetic dataset generated from `sklearn` as $X, y = \text{make_regression}(n = 1000, p = 100)$ with 90 informative features. For the splitting method, we average the results of 100 independent run. For the proposed root finding method, we approximate the boundaries of the exact set at precision 10^{-12}

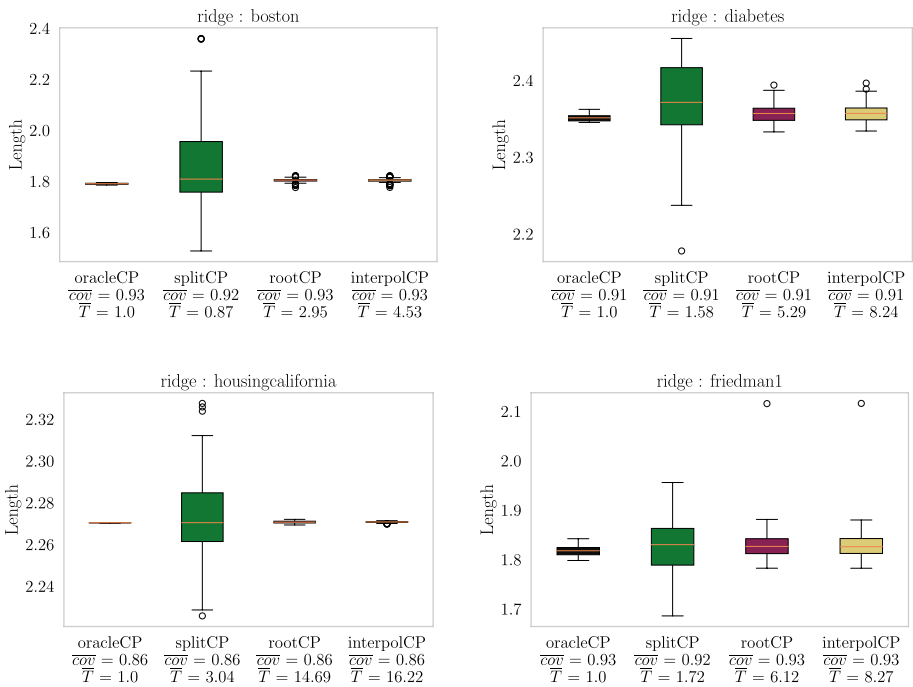


Fig. 4 Benchmarking conformal sets for ridge regression models on real datasets. We display the lengths of the confidence sets over 100 random permutation of the data. We denoted $\overline{\text{cov}}$ the average coverage, and \overline{T} the average computational time normalized with the average time for computing `oracleCP` which requires a single model fit on the whole data

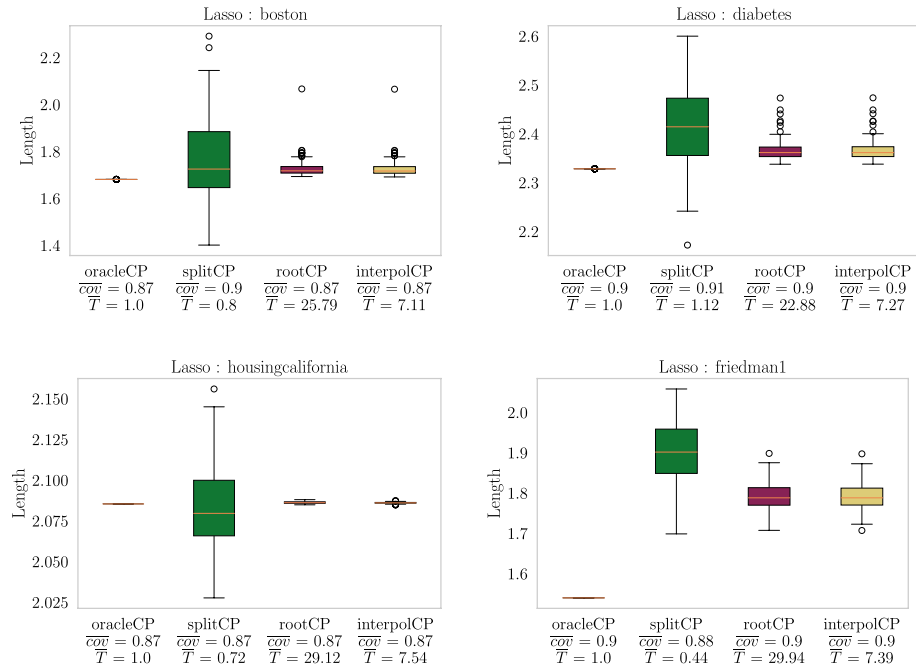


Fig. 5 Benchmarking conformal sets for Lasso regression models on real datasets. We display the lengths of the confidence sets over 100 random permutation of the data. We denoted \overline{cov} the average coverage, and \overline{T} the average computational time normalized with the average time for computing oracleCP which requires a single model fit on the whole data

4 Experiments

We numerically examine the performance of the root-finding methods to compute various conformal prediction sets for regression problems on both synthetic, and real databases. We summarize the datasets in Table 1.

The experiments were conducted with a coverage level of 0.9, *i.e.*, $\alpha = 0.1$. For comparisons, we run the evaluations on 100 repetitions of examples, and display the average of the following performance statistics for different methods: 1) the empirical coverage, *i.e.*, the percentage of times the prediction set contains the held-out target y_{n+1} ; 2) the length of the confidence intervals; 3) the execution time. For each run, we randomly select a couple of input/output (x_i, y_i) to constitute the targeted variables for which we will compute the conformal prediction set, and the rest is considered as observed data \mathcal{D}_n . Similar experimental setting was considered in Lei (2019).

From Lemma 1, we have $\pi(y_{n+1}) \geq \alpha$ with probability larger than $1 - \alpha$. Whence one can define the OracleCP as $\pi^{-1}([\alpha, +\infty))$ where π is obtained with a model fit optimized on the oracle data $\mathcal{D}_{n+1}(y_{n+1})$. In the case where the conformity function is the absolute value, we obtain the reference prediction set as in Ndiaye and Takeuchi (2019)

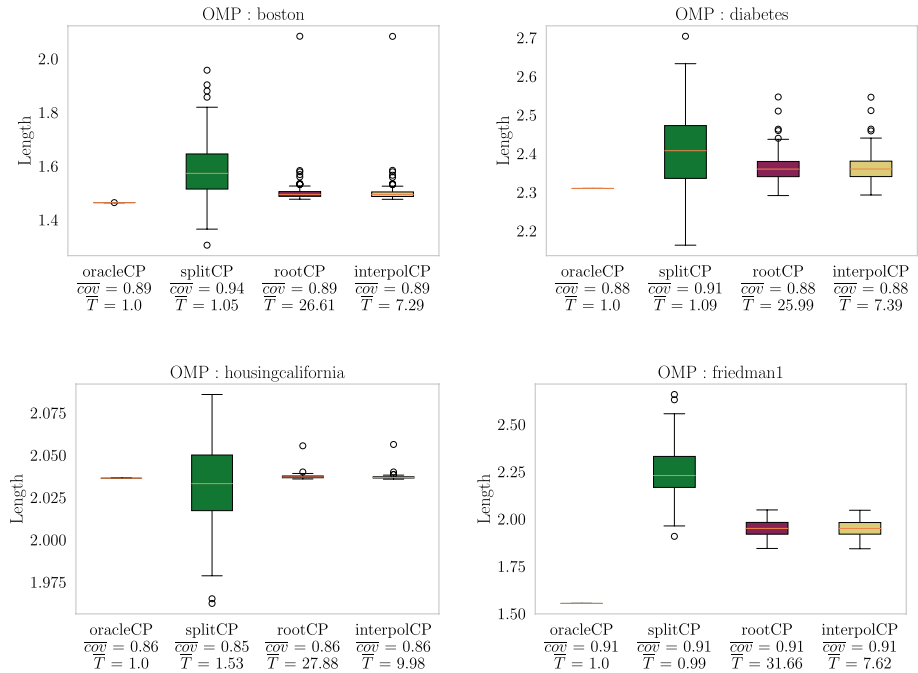


Fig. 6 Benchmarking conformal sets for orthogonal Matching Pursuit regression models on real datasets. We display the lengths of the confidence sets over 100 random permutation of the data. We denoted \overline{cov} the average coverage, and \overline{T} the average computational time normalized with the average time for computing `oracleCP` which requires a single model fit on the whole data

$$\text{oracleCP} : [\mu_{y_{n+1}}(x_{n+1}) \pm Q_{1-\alpha}(y_{n+1})].$$

We remind that the target variable y_{n+1} is not available in practice.

In the case of Ridge regression, *exact* conformal prediction sets can be computed by homotopy without data splitting, and without additional assumptions (Noureddinov et al., 2001). This allows us to finely assess the precisions of the proposed approaches, and illustrate the speed up benefit in Fig. 3.

We also illustrate the performance of our approach compared to the approximate homotopy method for the Lasso problem on a real data set from climate measurements. We first note that `splitCP` has a strictly, and significantly larger confidence set while the other approaches are quite close to the oracle performance. The approximate homotopy method uses all the data, and does not lose statistical efficiency if the model tolerance error is moderately small. However, as already noted in Ndiaye and Takeuchi (2019), it becomes unusable when the accuracy of the optimization becomes low. This is because its complexity depends directly on the accuracy of the model optimization, see the discussion in Sect. 3.3. We have shown that this does not affect our method because the number of model fits does not increase as the optimization error decreases. In particular, we can observe in Table 2 that `rootCP` is two to fifteen times faster when the tolerance errors range from 10^{-2} to 10^{-6} . This is mainly because the complexity of the root finding approach, i.e., the number of times it calls the model,

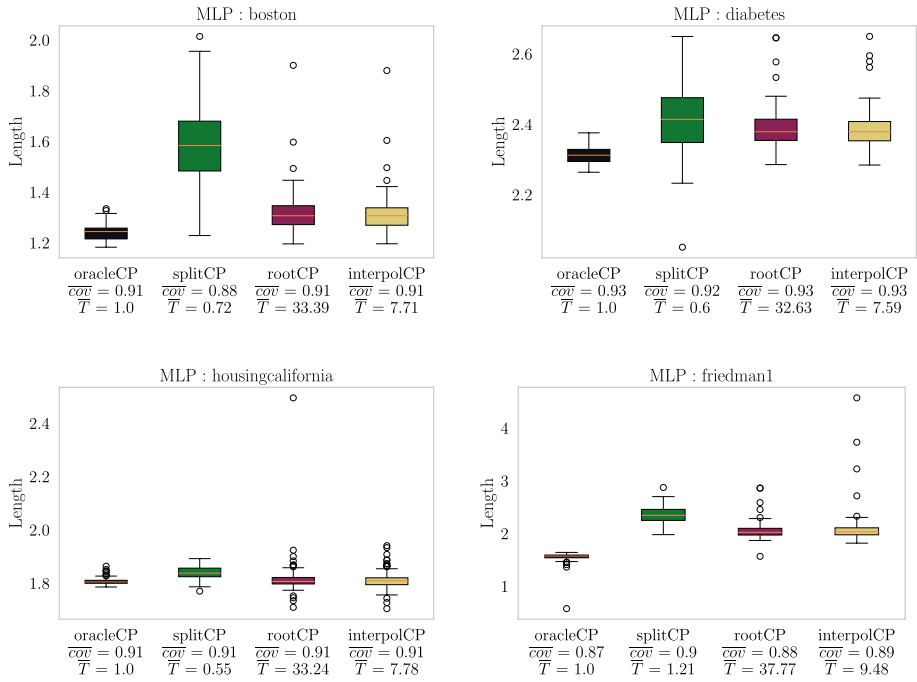


Fig. 7 Benchmarking conformal sets for Multi-layer Perceptron regression models on real datasets. We display the lengths of the confidence sets over 100 random permutation of the data. We denoted \overline{cov} the average coverage, and \overline{T} the average computational time normalized with the average time for computing `oracleCP` which requires a single model fit on the whole data

is independent of the optimization error of the underlying model fit. Whence, it allows the use of highly accurate estimators while maintaining feasible computations.

We run experiments on more complex regression models such as ridge in Fig. 4, Lasso in Fig. 5, orthogonal matching pursuit (OMP) in Fig. 6, feedforward neural network also called Multi-Layer Perceptron (MLP) in Fig. 7, Random Forest in Fig. 8, and Gradient Boosting in Fig. 9 (with warm start). In most of these settings, the estimator is obtained by approximating a solution of a non-convex optimization problem where none of the homotopy methods are available. We can observe in Figs. 4, 5, 6, 7 and 8 that the root-finding approach computes a full conformal prediction set while maintaining a reasonable computational time. In the worst cases observed, it costs about 30 times a single model fit which roughly corresponds to the 15 model fit for each root as predicted by the complexity (e.g., $\epsilon = 10^{-4}$). Moreover, this computational time is significantly reduced by the `interpolationCP` while achieving a statistical performance almost identical to the vanilla one in all our simulations. In all our experiments, we have chosen $d = 8$ number of query points for the interpolation method.

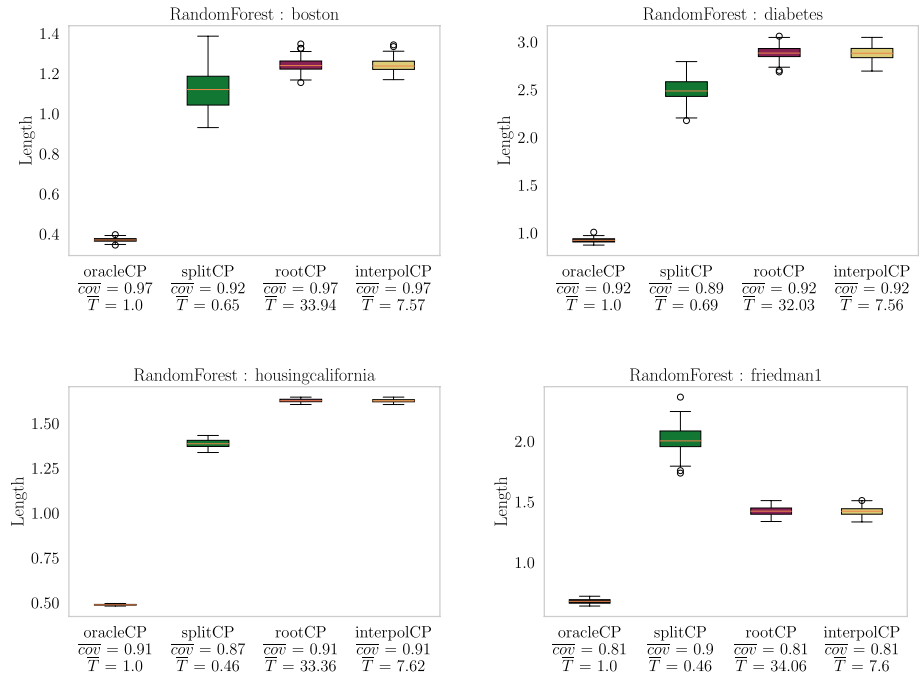


Fig. 8 Benchmarking conformal sets for Random Forest regression models on real datasets. We display the lengths of the confidence sets over 100 random permutation of the data. We denoted \overline{cov} the average coverage, and \overline{T} the average computational time normalized with the average time for computing `oracleCP` which requires a single model fit on the whole data

5 Conclusion

Since its introduction, the computation of a confidence region with full conformal prediction methods has been a major weakness to its adoption by a broader audience. The algorithms available until now were based on too strong assumptions that limited them to estimators whose map $z \mapsto \mu_z(\cdot)$ can be traced, *e.g.*, by using homotopy methods. We have shown that the limitations of the previous methods can be overcome by directly estimating the endpoints of the α -level set of the typicalness function with a root finding algorithm. Therefore, it is unnecessary to train the regression estimator an infinite number of times nor to make strong additional assumptions on the prediction model. As long as the conformal set represents an interval containing the point prediction obtained from the observed data, it can easily be estimated with only about ten numbers of model fits. The proposed approach can be readily applied to recent generalizations of the conformal prediction set beyond the exchangeability assumption as in Chernozhukov et al. (2018), Chernozhukov et al. (2021).

Nevertheless, we insist that a full conformal prediction set is not always an interval. In this case, our approach fails without properly bracketing these intervals. Both, testing the interval assumption or finding a proper bracketing are still difficult. Another severe,

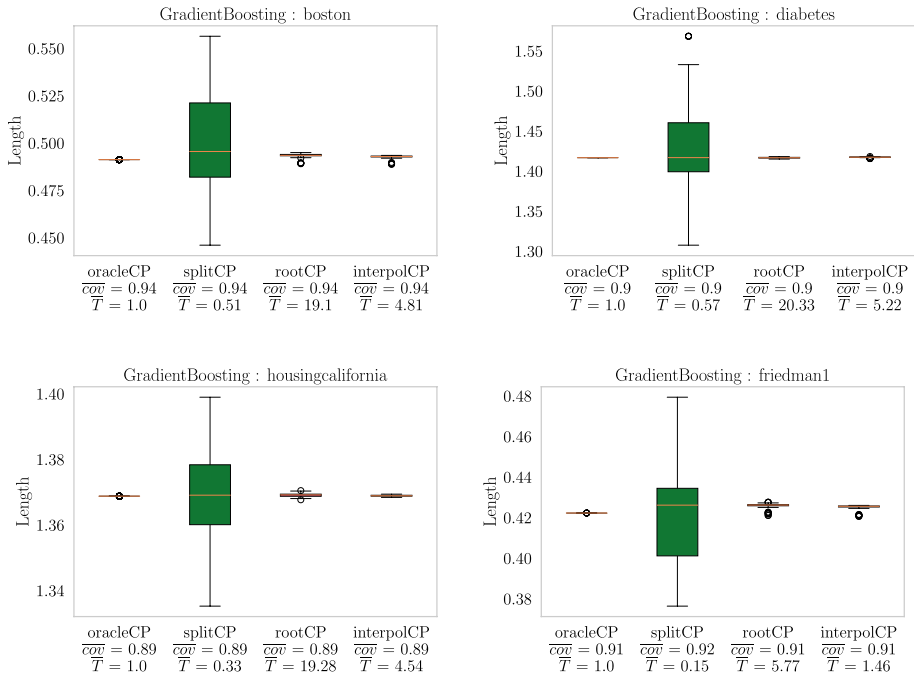


Fig. 9 Benchmarking conformal sets for Gradient Boosting regression models on real datasets. We display the lengths of the confidence sets over 100 random permutation of the data. We denoted \overline{cov} the average coverage, and \overline{T} the average computational time normalized with the average time for computing `oracleCP` which requires a single model fit on the whole data

and silent disadvantage is that the conformal set itself might be ill-defined when applied to regression estimators that depend on solving a non-convex optimization problem or a stochastic scheme (e.g., stochastic gradient descent). Indeed, in these cases, given a fixed candidate z , $\pi(z)$ can take multiple values depending on the initialization or the random seed. Moreover, the symmetry assumptions might be violated if the instances are not used evenly (e.g., when using stochastic gradient descent with importance sampling). As a future work, it would therefore be interesting to understand how these points can negatively affect the coverage guarantee, and the computational complexity.

Table 1 The first four datasets used in our experiments are available in `sklearn`

	Boston	Diabetes	Housingcalifornia	Friedman1	Climate
n	506	442	20,640	500	814
p	13	10	8	100	73,570

Here, n is the number of samples, and p is the number of features. The noise level in the Friedman dataset is fixed to 1. The climate dataset is from NCEP/NCAR Reanalysis Kalnay et al. (1996)

Table 2 Computing a conformal set for a Lasso regression problem on a climate data set with $n = 814$ observations, and $p = 73,570$ features

	oracleCP	splitCP	Approximate homotopy/rootCP		
			1e-2	1e-4	1e-6
Coverage	0.94	0.90	0.95/0.93	0.94/0.94	0.94/0.94
Length	0.693	1.054	0.747/0.699	0.711/0.708	0.707/0.707
Time (s)	3.088	1.409	4.289/2.478	29.987/9.173	316.67/20.92

We display the coverage, length, and execution times for different methods averaged over 100 randomly held-out validation data sets

Author contributions EN introduced, and formalized the contributions of the paper. He wrote, and analyzed the technical details, implemented the algorithms, designed, and performed the numerical simulations. IT contributed to the conceptualization, formulation, and design of the numerical experiments in this study.

Funding IT was partially supported by MEXT KAKENHI (20H00601, 16H06538), JST CREST (JPM-JCR21D3), JST Moonshot R &D (JPMJMS2033-05), JST AIP Acceleration Research (JPMJCR21U2), NEDO (JPNP18002, JPNP20006), and RIKEN Center for Advanced Intelligence Project.

Data availability All the data, and software used are available in open source.

Code availability The source of our implementation is available at <https://github.com/EugeneNdiaye/rootCP>.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Ethics approval Not applicable.

Consent to participate All authors agreed to participate as authors of this article.

Consent for publication All authors consent for submission, and publication.

References

- Allgower, E. L., & Georg, K. (2012). *Numerical continuation methods: An introduction*. Springer.
- Angelopoulos, A., Bates, S., Malik, J., & Jordan, M. (2020). Uncertainty sets for image classifiers using conformal prediction. In *ICLR*.
- Arlot, S., & Celisse, A. (2010). A survey of cross-validation procedures for model selection. *Statistics Surveys*, 4, 40–79.
- Bach, F., Jenatton, R., Mairal, J., & Obozinski, G. (2012). Optimization with sparsity-inducing penalties. *Foundations and Trends in Machine Learning*, 4, 1–106.

- Balabasubramanian, V., Ho, S.-S., & Vovk, V. (2014). *Conformal prediction for reliable machine learning: Theory, adaptations and applications*. Elsevier.
- Barber, R. F., Candès, E. J., Ramdas, A., & Tibshirani, R. J. (2021). Predictive inference with the jackknife+. *The Annals of Statistics*, 49, 486–507.
- Bates, S., Angelopoulos, A., Lei, L., Malik, J., & Jordan, M. (2021a). Distribution-free, risk-controlling prediction sets. arXiv preprint [arXiv:2101.02703](https://arxiv.org/abs/2101.02703)
- Bates, S., Candès, E., Lei, L., Romano, Y., & Sesia, M. (2021b). Testing for outliers with conformal p -values. arXiv preprint [arXiv:2104.08279](https://arxiv.org/abs/2104.08279)
- Bousquet, O., & Bottou, L. (2008). The tradeoffs of large scale learning. In *NeurIPS*.
- Carlsson, L., Eklund, M., & Norinder, U. (2014). Aggregated conformal prediction.
- Cella, L., & Ryan, R. (2020). Valid distribution-free inferential models for prediction. arXiv preprint [arXiv:2001.09225](https://arxiv.org/abs/2001.09225).
- Chang, Y.-C., & Hung, W.-L. (2007). LINEX loss functions with applications to determining the optimum process parameters. *Quality & Quantity*, 41, 291–301.
- Chen, W., Chun, K.-J., & Barber, R. F. (2018). Discretized conformal prediction for efficient distribution-free inference. *Stat*, 7, e173.
- Chernozhukov, V., Wüthrich, K., & Zhu, Y. (2018). Exact and robust conformal inference methods for predictive machine learning with dependent data. In *Conference On learning theory*.
- Chernozhukov, V., Wüthrich, K., & Zhu, Y. (2021). An exact and robust conformal inference method for counterfactual and synthetic controls. *Journal of the American Statistical Association*, 116, 1849–1864.
- Cox, D. R. (1975). A note on data-splitting for the evaluation of significance levels. *Biometrika*, 62, 441–444.
- Fisch, A., Schuster, T., Jaakkola, T., & Barzilay, R. (2021). Few-shot conformal prediction with auxiliary tasks. In *ICML*.
- Gamerman, A., Vovk, V., & Vapnik, V. (1998). Learning by transduction. In *Proceedings of the fourteenth conference on uncertainty in artificial intelligence*.
- Gärtner, B., Jaggi, M., & Maria, C. (2012). An exponential lower bound on the complexity of regularization paths. *Journal of Computational Geometry*.
- Giesen, J., Jaggi, M., & Laue, S. (2010). Approximating parameterized convex optimization problems. In *European Symposium on Algorithms*.
- Gruber, M. (2010). *Regression estimators: A comparative study*. JHU Press.
- Ho, S.-S., & Wechsler, H. (2008). Query by transduction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30, 1557–1571.
- Hoerl, A. E. (1962). Application of ridge analysis to regression problems. *Chemical Engineering Progress*, 58, 54–59.
- Hoerl, A. E., & Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12, 55–67.
- Holland, M.J. (2020). Making learning more transparent using conformalized performance prediction. arXiv preprint [arXiv:2007.04486](https://arxiv.org/abs/2007.04486).
- Kalnay, E., Kanamitsu, M., Kistler, R., Collins, W., Deaven, D., Gandin, L., Iredell, M., Saha, S., White, G., Woollen, J., et al. (1996). The NCEP/NCAR 40-year reanalysis project. *Bulletin of the American Meteorological Society*, 77, 437–472.
- Kim, B., Xu, C., & Barber, R. (2020). Predictive inference is free with the jackknife+ after bootstrap. *Advances in Neural Information Processing Systems*, 33, 4138–4149.
- Laxhammar, R., & Falkman, G. (2015). Inductive conformal anomaly detection for sequential detection of anomalous sub-trajectories. *Annals of Mathematics and Artificial Intelligence*, 74, 67–94.
- Lehmann, E.L., & Romano, J.P. (2006). *Testing statistical hypotheses*. Springer Science.
- Lei, J. (2019). Fast exact conformalization of Lasso using piecewise linear homotopy. *Biometrika*, 106, 749–764.
- Lei, J., G’Sell, M., Rinaldo, A., Tibshirani, R. J., & Wasserman, L. (2018). Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113, 1094–1111.
- Linusson, H., Norinder, U., Boström, H., Johansson, U., & Löfström, T. (2017). On the calibration of aggregated conformal predictors. In *Conformal and probabilistic prediction and applications*.
- Mairal, J., & Yu, B. (2012). Complexity analysis of the Lasso regularization path. In *ICML*.
- Ndiaye, E., Le, T., Fercocq, O., Salmon, J., & Takeuchi, I. (2019). Safe grid search with optimal complexity. In *ICML*.
- Ndiaye, E., & Takeuchi, I. (2019). Computing full conformal prediction set with approximate homotopy. In *NeurIPS*.
- Noureddinov, I., Melluish, T., & Vovk, V. (2001). Ridge regression confidence machine. In *ICML*.

- Obozinski, G., & Bach, F. (2016). A unified perspective on convex structured sparsity: Hierarchical, symmetric, submodular norms and beyond. In *HAL Id : hal-01412385, version 1*.
- Papadopoulos, H., Proedrou, K., Vovk, V., & Gammerman, A. (2002). Inductive confidence machines for regression. In *European conference on machine learning*.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research, 12*, 2825–2830.
- Qin, T., Liu, T.-Y., & Li, H. (2010). A general approximation framework for direct optimization of information retrieval measures. *Information Retrieval, 13*, 375–397.
- Romano, Y., Patterson, E., & Candes, E. J. (2019). Conformalized quantile regression. In *NeurIPS*.
- Shafer, G., & Vovk, V. (2008). A tutorial on conformal prediction. *Journal of Machine Learning Research*.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological), 58*, 267–288.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Larson, E., Carey, C. J.,...SciPy 1.0 Contributors (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods, 17*, 261–272.
- Vovk, V. (2015). Cross-conformal predictors. *Annals of Mathematics and Artificial Intelligence, 74*, 9–28.
- Vovk, V., Gammerman, A., & Shafer, G. (2005). *Algorithmic learning in a random world*. Springer.
- Vovk, V., Shen, J., Manokhin, V., & Xie, M. (2017). Nonparametric predictive distributions based on conformal prediction. In *Conformal and probabilistic prediction and applications*.
- Xie, H., & Huang, J. (2009). SCAD-penalized regression in high-dimensional partially linear models. *The Annals of Statistics, 37*, 673–696.
- Xu, C., & Xie, Y. (2021). Conformal prediction interval for dynamic time-series. In *ICML*.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.