

Valid prediction intervals for regression problems

Nicolas Dewolf · Bernard De Baets ·
Willem Waegeman

Abstract Over the last few decades, various methods have been proposed for estimating prediction intervals in regression settings, including Bayesian methods, ensemble methods, direct interval estimation methods and conformal prediction methods. An important issue is the validity and calibration of these methods: the generated prediction intervals should have a predefined coverage level, without being overly conservative. So far, no study has analysed this issue whilst simultaneously considering these four classes of methods. In this independent comparative study, we review the above four classes of methods from a conceptual and experimental point of view in the i.i.d. setting. Results on benchmark data sets from various domains highlight large fluctuations in performance from one data set to another. These observations can be attributed to the violation of certain assumptions that are inherent to some classes of methods. We illustrate how conformal prediction can be used as a general calibration procedure for methods that deliver poor results without a calibration step.

Keywords Prediction interval - calibration - regression - Bayesian network - ensemble theory - conformal prediction

Declarations

Funding: This research received funding from the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” programme.

Nicolas Dewolf
KERMIT, Department of Data Analysis and Mathematical Modelling
Ghent University, Belgium
E-mail: nicolas.dewolf@ugent.be

Bernard De Baets
KERMIT, Department of Data Analysis and Mathematical Modelling
Ghent University, Belgium
E-mail: bernard.debaets@ugent.be

Willem Waegeman
KERMIT, Department of Data Analysis and Mathematical Modelling
Ghent University, Belgium
E-mail: willem.waegeman@ugent.be

Conflicts of interest: None

Availability of data and material: All data sets are freely available (references are provided in the text).

1 Introduction

Machine learning methods, and in particular deep learning methods, often serve as work horses in artificial intelligence systems that have a strong impact on the daily life of humans, such as self-driving cars [10, 56], machine translation [61, 79] and medical diagnostics [16, 37]. However, such systems are only accepted by humans if they exhibit a sufficient degree of reliability. As a result, analyzing what systems “know” and what they “don’t know” has become an important topic of recent deep learning research, using “uncertainty quantification” and “uncertainty estimation” as prominent buzz words [52, 54, 83, 85]. Driven by popular application domains like computer vision and natural language processing, most of the recent literature heavily focuses on classification problems – see e.g. [31, 33, 58, 85]. Regression problems are often ignored, or at least less analyzed in such papers.

Methods for uncertainty quantification in classification and regression problems usually differ substantially. Many traditional classification methods produce probability estimates, which are used as a starting point for uncertainty quantification, out-of-distribution detection and open-set recognition [26]. In regression, most traditional methods are so-called point predictors; they only predict one summary statistic of the conditional distribution (in many cases the conditional mean). However, a point predictor cannot express how confident it is of a prediction, and typically a prediction interval is returned by more complicated methods to quantify uncertainty, i.e. the wider the interval, the larger the uncertainty. Such prediction intervals can be obtained by modelling the conditional distribution in an exact or approximate manner, using for example Bayesian methods [29, 92] or ensemble methods [40, 95]. Prediction intervals can also be estimated in a more direct manner, without modelling the conditional distribution entirely, using for example quantile regression [46] or conformal prediction methods [73, 89].

The estimation of prediction intervals for regression has received little attention recently, and the last general review predates the ongoing deep learning wave [42] (at the time of writing another review appeared with a strong focus on fuzzy methods [9]). By now, many of the older methods have been superseded or have become unwieldy due to the complexity of modern data sets. Moreover, recent advances in deep learning have led to the introduction of some novel frameworks, such as Bayesian neural networks [29], which are typically only applied to classification problems, while having also potential for regression problems. As a result, an up-to-date comparison of methods that generate prediction intervals is necessary. This paper intends to bridge that gap, by focusing on the following aspects:

- i) We give an overview of four general classes of methods that produce prediction intervals: Bayesian methods, ensemble methods, direct estimation methods and conformal prediction methods. We discuss some representative examples of each class, while paying attention to potential advantages or disadvantages.

- ii) We intend to elaborate on the calibration (or validity) of prediction intervals and relate it to data and model properties. For many applications it is indeed important to have well-calibrated prediction intervals (see below for a formal definition of calibration).
- iii) We show how *conformal prediction* can be applied as a general framework to obtain well-calibrated prediction intervals, starting from simple point predictors and methods that generate uncalibrated intervals.
- iv) We provide an in-depth experimental comparison of the four main classes of methods based on their performance across a wide range of data sets. We interpret the observed differences and discuss practical difficulties such as hyperparameter tuning and model selection based on prior knowledge.

The general structure of the paper is as follows. In Section 2 some general aspects of the estimation of prediction intervals for regression are discussed. Subsequently, in Section 3, the different classes of methods are reviewed. The setup of an experimental assessment for a selection of methods is presented in Section 4. A discussion of the main results of those experiments can be found in Section 5.

2 Formal problem statement

Before being able to formally state the problem of confidence estimation, a small note on the notations used in the remainder of the text is required. In this study, only regression problems involving real-valued functions $f : \mathcal{X} \rightarrow \mathbb{R}$ will be considered, where the feature space \mathcal{X} can contain both continuous and discrete “dimensions”. Generic models and model parameters will be denoted by \hat{y} and θ , respectively, while a generic interval estimator (to be introduced below) will be denoted by Γ . If emphasis needs to be placed on the quantity that is estimated, the estimator will be indicated by a caret, e.g. an estimator for the standard deviation σ will be denoted by $\hat{\sigma}$. The fixed confidence level will always be denoted by α . The Cartesian product of the feature and target spaces will be denoted by $\mathcal{Z} := \mathcal{X} \times \mathbb{R}$. A generic element of the instance space \mathcal{Z} will be denoted by $\mathbf{z} \equiv (\mathbf{x}, y)$, while a collection of such elements will be denoted by $\mathbf{Z} \equiv (\mathbf{X}, \mathbf{y})$, where \mathbf{X} denotes the design matrix obtained from stacking the feature arrays from the individual observations. Here, the common convention is used that tuples are indicated by boldface and matrices are written in boldface capitals. A new, previously unseen instance will be denoted by an asterisk: $\mathbf{z}^* \equiv (\mathbf{x}^*, y^*)$. General data sets will be denoted by \mathcal{D} . These can be divided into a training, validation (or calibration) and test set, which are denoted by \mathcal{T} , \mathcal{V} and \mathcal{D}^* respectively.

Given a tuple of features $\mathbf{x} \in \mathcal{X}$, it is not only interesting to obtain a point prediction $\hat{y}(\mathbf{x}) \in \mathbb{R}$, but often one is also interested in a region $\Gamma(\mathbf{x}) \subseteq \mathbb{R}$ such that the true response $y \in \mathbb{R}$ will be contained in $\Gamma(\mathbf{x})$ with a predetermined probability. Functions Γ that assigns such a region to every point $\mathbf{x} \in \mathcal{X}$ are generally called *confidence estimators*. In this study, only interval estimators will be considered, i.e. functions $\Gamma : \mathcal{X} \rightarrow [\mathbb{R}]$, where $[\mathbb{R}]$ denotes the set of all closed subintervals of \mathbb{R} . When constructing prediction intervals one has to specify the probability with which the estimator is allowed to make errors, since in general no realistic model, i.e. a model with a finite number of parameters inferred from

a finite number of data points, can be perfect. Defining the *coverage (probability)* of the interval estimator Γ as:

$$\mathcal{C}(\Gamma, P) := \mathbb{E}_{(\mathbf{x}, y) \sim P} [\mathbb{1}(y \in \Gamma(\mathbf{x}))] = \text{Prob}\{y \in \Gamma(\mathbf{x})\}, \quad (1)$$

where $\mathbb{1}$ denotes the indicator function and P denotes the joint distribution on \mathcal{Z} , a *significance* or *confidence level* α is chosen [17, 21] such that

$$\mathcal{C}(\Gamma, P) \geq 1 - \alpha \quad (2)$$

should hold. In some cases this condition will, however, not be satisfied. To distinguish between these situations, estimators satisfying the inequality are said to be (*conservatively*) *valid* [89] or *calibrated* (in accordance with the classification literature [31]). It should be clear that the best models are those that get as close as possible to saturating the inequality. Of course, one could argue that a perfect model would have a coverage of 100%, i.e. it should always be completely confident, but this cannot be a realistic requirement. In real-life situations only a finite number of data samples can be provided, which implies that there always exists some uncertainty.

Moreover, simply having an estimate of the uncertainty that satisfies condition (2) is not sufficient. A trivial example would be the situation where a model always outputs the full target space as a prediction region. It clearly satisfies the validity condition (2), but it is hard to extract any meaning from the result. The simplest measure of meaningfulness will, therefore, be the average length of the prediction intervals:

$$\mathcal{W}(\Gamma, P) := \mathbb{E}_{(\mathbf{x}, y) \sim P} [|u(\mathbf{x}) - l(\mathbf{x})|], \quad (3)$$

where the functions u, l denote the upper and lower bound of the prediction interval produced by Γ . Depending on the context (this includes both the type of data and the class of models) other, more informative, measures can be considered. However, since this study aims to be as high-level as possible, only this particular choice will be considered.

3 Overview of methods

In this section, four general classes of models are discussed. The first class has its roots in probability theory and, therefore, can be expected to have better theoretical guarantees for the validity and behaviour. The second class consists of methods that are built from a collection of estimators and generally have a superior predictive performance when compared to individual models. The third class contains the models that are specifically trained to yield a prediction interval, while the last class constitutes a framework that allows to turn any given point predictor into a valid interval estimator. Many different methods are nowadays being used to produce uncertainty estimates. However, to our opinion, all methods that construct prediction intervals in a regression setting can be assigned to one of these four classes.

Algorithm 1: Bayesian modellingframework Bayes ($\mathcal{A}, \mathcal{L}, p, \mathcal{D}$)**Input** : Model architecture \mathcal{A} , likelihood function $\mathcal{L}(Y | X, \theta)$, prior distribution $p(\theta)$, data set \mathcal{D} **Output**: Predictive distribution $p(Y | X, \mathcal{D})$

- 1 Construct a model with architecture \mathcal{A} , where the parameters are sampled from $p(\theta)$
- 2 Update the prior distribution through Bayes' rule (4)
- 3 Infer the predictive distribution $p(Y | X, \mathcal{D})$ using Eq. (5)
- 4 **return** $p(Y | X, \mathcal{D})$

3.1 Bayesian methods

In Bayesian inference one tries to model the distribution of interest by updating a prior estimate using a collection of observed data. The conditional distribution $p(Y | X, \mathcal{D})$ is inferred from a given parametric model or *likelihood* function $p(Y | X, \theta)$, a prior distribution $p(\theta)$ over the model parameters and a data set $\mathcal{D} \equiv (\mathbf{X}, \mathbf{y})$. The first step is to update the prior belief based on the data set using Bayes' rule:

$$p(\theta | \mathcal{D}) = \frac{p(\mathbf{y}, \theta | \mathbf{X})}{p(\mathbf{y} | \mathbf{X})} = \frac{p(\mathbf{y} | \mathbf{X}, \theta)p(\theta)}{\int p(\mathbf{y} | \mathbf{X}, \theta)p(\theta) d\theta}. \quad (4)$$

After this posterior distribution is computed, the posterior predictive distribution is calculated by marginalizing over the parameters:

$$p(y^* | \mathbf{x}^*, \mathcal{D}) = \int p(y^* | \mathbf{x}^*, \theta)p(\theta | \mathcal{D}) d\theta. \quad (5)$$

This process is summarized in Algorithm 1. Note that the algorithm can simply be repeated in an on-line fashion when more data becomes available. One simply takes the “old” posterior distribution $p(\theta | \mathcal{D})$ as the new prior distribution.

To obtain a point estimate for future predictions, the most popular choice is the conditional mean $E[y^* | \mathbf{x}^*, \mathcal{D}]$. For Bayesian methods the construction of a prediction interval reduces to the classical, but not always straightforward, problem of finding prediction intervals for a given distribution. One of the main advantages of Bayesian methods compared to other approaches is that one can incorporate domain knowledge into the prior distribution. By leveraging this knowledge to more accurately characterize the class of distributions underlying the data generating process, one can obtain improved uncertainty estimates, especially in settings with very few data.

Gaussian processes

One of the most popular probabilistic models for regression problems is the Gaussian process [92]. The main reason for its popularity is that it is one of the only Bayesian methods where the inference step (4) can be performed exactly, since the marginalization of multivariate normal distributions can be written in closed form and, hence, no approximations are needed (in principle). Formally, a Gaussian process (GP) is defined as a stochastic process for which the joint distribution of

any finite number of random variables is (multivariate) Gaussian. Gaussian processes are, therefore, characterized by only two functions: the mean and covariance functions $m : \mathcal{X} \rightarrow \mathbb{R}$ and $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$.

If one assumes that the population is distributed according to a Gaussian process with given mean and covariance functions, by definition, the responses for the training and test instances \mathbf{X}, \mathbf{x}^* are distributed according to a multivariate Gaussian as follows (functions are calculated elementwise on tuples and matrices):

$$\begin{pmatrix} \mathbf{y} \\ y^* \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} m(\mathbf{X}) \\ m(\mathbf{x}^*) \end{pmatrix}, \begin{pmatrix} k(\mathbf{X}, \mathbf{X}) & k(\mathbf{X}, \mathbf{x}^*) \\ k(\mathbf{x}^*, \mathbf{X}) & k(\mathbf{x}^*, \mathbf{x}^*) \end{pmatrix} \right) \equiv \mathcal{N} \left(\begin{pmatrix} \boldsymbol{\mu} \\ \mu^* \end{pmatrix}, \begin{pmatrix} \Sigma & \Sigma^* \\ (\Sigma^*)^\top & \Sigma^{**} \end{pmatrix} \right), \quad (6)$$

where the notations $\boldsymbol{\mu} := m(\mathbf{X})$, $\mu^* := m(\mathbf{x}^*)$, $\Sigma := k(\mathbf{X}, \mathbf{X})$, etc. are introduced for brevity. With this convention, $\boldsymbol{\mu}$ and Σ are, for example, a tuple and a matrix, respectively. By conditioning on the data set \mathcal{D} , one can obtain the posterior distribution:

$$(y^* | \mathbf{x}^*, \mathcal{D}) \sim \mathcal{N}(\mu^* + (\Sigma^*)^\top \Sigma^{-1}(\mathbf{y} - \boldsymbol{\mu}), \Sigma^{**} - (\Sigma^*)^\top \Sigma^{-1} \Sigma^*). \quad (7)$$

From the latter formula it is evident that the point estimate is given by

$$\mathbb{E}[y^* | \mathbf{x}^*, \mathcal{D}] = \mu^* + (\Sigma^*)^\top \Sigma^{-1}(\mathbf{y} - \boldsymbol{\mu}). \quad (8)$$

Although they are conceptually extremely simple, it is also immediately clear that Gaussian processes suffer from a major drawback. Due to the matrix inversion in the calculation of both the mean and covariance of the posterior distribution, they are computationally less efficient than neural network-based methods. The memory complexity of matrix inversion scales as $O(|\mathcal{D}|^2)$, while the computational complexity scales as $O(|\mathcal{D}|^3)$. For large data sets this gives rise to a considerable overhead. However, many approximations exist in the literature. Some of these try to reduce the complexity of the matrix operations by introducing stochastic approximations, while others belong to the variational class to be introduced in the next section [34, 93].

There is also an important remark to be made with respect to the validity condition (2). When using GPs, two important assumptions are made. By definition it is assumed that the data is (conditionally) normally distributed and a further choice of covariance function has to be made by the user. When these assumptions do not properly characterize the data, it can be expected that the resulting distribution will also not be optimal. In general this will lead to invalid models.

Bayesian neural networks

Bayesian modelling can easily be implemented using neural networks. Here, the likelihood function $p(Y | X, \theta)$ in Eq. (4) is given by a parametric function for which the parameters are estimated by a neural network [29, 53, 59]. A common example is a fixed-variance Gaussian likelihood [36], where the mean is estimated by a neural network with weights θ :

$$(y | \mathbf{x}, \theta) \sim \mathcal{N}(\hat{y}^\theta(\mathbf{x}), \sigma^2). \quad (9)$$

Although ordinary neural networks have the benefit that even for a large number of features and weights they can be implemented very efficiently, their Bayesian incarnation suffers from a problem. The nonlinearities in the activation functions and the sheer number of parameters, although they are the features that make traditional NNs so powerful, lead to the inference steps (4) and (5) becoming intractable.

Approximate Bayesian inference

Both the integral in the inference step (4) and in the prediction step (5) can in general not be computed exactly (conjugate priors [19], such as normal distributions, form an important exception). At inference there are two general classes of approximations available:

1. Variational inference (VI): Instead of computing the posterior distribution through Eq. (4), the problem is reformulated as a variational problem, i.e. the posterior distribution $p(\theta | \mathcal{D})$ is replaced by a parametric family of distributions $q(\theta; \lambda)$ and a divergence measure is optimized within this family [3]. For the Kullback-Leibler divergence D_{KL} , the loss function

$$\mathcal{L}_{\text{VI}}(\mathcal{D}; \lambda) := D_{\text{KL}}(q(\theta; \lambda) \| p(\theta)) - \mathbb{E}[\log p(\mathcal{D} | \theta)] \quad (10)$$

is called the *variational free energy* in statistical physics terminology or *evidence lower bound* (ELBO) in machine learning terminology. For Bayesian neural networks this function can be optimized using a backpropagation scheme similar to traditional neural networks [4].

2. Monte Carlo integration: The general idea of Monte Carlo (MC) integration [27] is that for a general distribution p , an integral of the form

$$I[f] := \int f(x)p(x) \, dx \quad (11)$$

is approximated by a finite sum

$$I[f] \approx \frac{1}{n} \sum_{i=1}^n f(x_i), \quad (12)$$

where the points x_i are drawn from the distribution p .

For making predictions, i.e. calculating the conditional mean, MC integration gives the following expression:

$$\mathbb{E}[y^* | \mathbf{x}^*, \mathcal{D}] = \frac{1}{n} \sum_{i=1}^n y_i, \quad (13)$$

where the samples y_i are drawn from the posterior distribution $p(y^* | \mathbf{x}^*, \mathcal{D})$.

Algorithm 2: Ensemble learning

framework Ensemble $(R, \{(\mathcal{S}_i, \mathcal{A}_i)\}_{i=1}^R, \mathcal{D}, \mathcal{E})$
Input : Number of models R , model architectures \mathcal{A}_i , sampling strategies \mathcal{S}_i , data set \mathcal{D} , aggregation strategy \mathcal{E}
Output: Ensemble predictor \hat{y}

```

1 for (  $i = 1$ ;  $i \leq R$ ;  $i = i + 1$  ) {
2   | Use strategy  $\mathcal{S}_i$  to obtain a data set  $\mathcal{D}_i$  from  $\mathcal{D}$ 
3   | Construct a model  $\hat{y}_i$  with architecture  $\mathcal{A}_i$ 
4   | Train model  $\hat{y}_i$  on  $\mathcal{D}_i$ 
5 }
6 Aggregate models  $\{\hat{y}_i\}_{i=1}^R$  into a new model  $\hat{y}$  using the strategy  $\mathcal{E}$ 
7 return  $\hat{y}$ 

```

3.2 Ensemble methods

Ensemble learning is a popular approach to enhance predictions by training multiple machine learning models and aggregating the individual predictions, for example by taking the mean [80]. In general one can consider ensemble methods as an intermediate step between Bayesian methods and pure point predictors. They yield an approximation to Bayesian inference, where the trained models represent a sample in the parameter space and the aggregation is a kind of MC integration. This also implies that they have a natural notion of uncertainty. However, this uncertainty estimate does not always easily admit a probabilistic interpretation, because the ensembles are in practice often of an ad hoc nature and only roughly represent approximate Bayesian inference. The general structure of an ensemble learning scheme is summarized in Algorithm 2.

Every ensemble allows for a naive construction of a prediction interval [35] when the aggregation strategy in Algorithm 2 is given by the arithmetic mean. By treating the predictions of the individual models in the ensemble as elements of a data sample, one can calculate the empirical mean and variance and use these as moment estimators for a normal distribution:

$$\hat{\mu}(\mathbf{x}) := \frac{1}{R} \sum_{i=1}^R \hat{y}_i(\mathbf{x}), \quad (14)$$

$$\hat{\sigma}^2(\mathbf{x}) := \frac{1}{R} \sum_{i=1}^R (\hat{y}_i(\mathbf{x}) - \hat{\mu}(\mathbf{x}))^2, \quad (15)$$

where R denotes the ensemble size. This gives rise to the following Gaussian prediction interval:

$$\Gamma_{\text{ensemble}}^\alpha(\mathbf{x}^*) := [\hat{\mu}(\mathbf{x}^*) - z^\alpha \hat{\sigma}(\mathbf{x}^*), \hat{\mu}(\mathbf{x}^*) + z^\alpha \hat{\sigma}(\mathbf{x}^*)], \quad (16)$$

where z^α is the two-tailed z -score corresponding to a significance level α for a standard normal distribution, e.g. $z^\alpha \approx 1.65$ for $\alpha = 0.1$.

Random forests (and other bagged ensembles)

Random forests, as introduced in [7] by Breiman, are an extension of the idea of decision trees. Their main feature is the use of *bagging* [6], where every decision tree

in the forest is constructed from a subset of the original training set (possibly with replacement). Random forests have been shown to be robust against overfitting and, due to their nonparametric nature, they can be reliably applied in a wide variety of situations.

The main disadvantage of random forests is the fact that they do not produce an uncertainty estimate by themselves aside from the crude construction (16) introduced above. There are two general methods for obtaining a more specialized prediction region from these predictors: Either one uses the *out-of-bag samples* to calculate an uncertainty estimate or one uses a stand-alone procedure to extract uncertainty estimates (the former could be considered as a specific case of the latter).

While the second approach will be discussed in Section 3.4 in more detail, the first one will be considered here. For a random forest, and in fact for any ensemble that implements the bagging procedure, it is possible to use *out-of-bag* (OOB) estimates to construct prediction regions. Because these ensembles do not use the full training set to train the individual models, one can for every training sample $\mathbf{z}_i \equiv (\mathbf{x}_i, y_i)$ use the subensemble $\hat{y}_{(i)}$ of models that were not trained on \mathbf{z}_i to get an estimate of the prediction accuracy. A well-established first step in this direction is the extension of the (infinitesimal) *jackknife* method for variance estimation [57] to bagged ensembles [15, 91]. This approach applies the general construction (16) to the ensemble consisting of the OOB estimators for all training samples. A related approach, first introduced in [38] and later generalized in [95], uses the OOB estimates to construct an empirical error distribution $D := \{y_i - \hat{y}_{(i)}(\mathbf{x}_i) : (\mathbf{x}_i, y_i) \in \mathcal{T}\}$. The prediction interval for a new instance \mathbf{x}^* is constructed as follows:

$$\Gamma_{\text{OOB}}^\alpha(\mathbf{x}^*) := [\hat{y}(\mathbf{x}^*) + D_{[\alpha/2]}, \hat{y}(\mathbf{x}^*) + D_{[1-\alpha/2]}], \quad (17)$$

where $D_{[\beta]}$ denotes the empirical β -quantile of D . In Section 3.4 this method will be seen to reduce to a *conformal predictor* in the case of symmetric error distributions as originally considered in [38]. A similar idea was used in [2], where a cross-validation-based framework, called the *jackknife+*, was introduced. Here, a prediction interval is obtained by using *k-fold cross validation* or *leave-one-out cross validation* to estimate the generalization error for arbitrary regression models. The above interval estimator for bagged ensembles is a specific instance of this approach. It should be noted that both the *jackknife+* method and the estimator (17), also called the *jackknife* method in [2], are nonparametric in contrast to the genuine jackknife estimator from [91].

Random forest quantile regression

Meinshausen [55] modified the random forest prediction method from the previous section to be able to directly estimate quantiles. Ordinary regression forests estimate the conditional mean of the response variable by taking a (weighted) average over the training instances in the leaves where the new instance belongs to. Quantile regression forests generalize this idea by using the property that the conditional distribution can be expressed as a conditional mean:

$$F(y^* | \mathbf{x}^*) = \mathbb{E}_{(\mathbf{x}, y) \sim P} [\mathbb{1}_{y \leq y^*} | \mathbf{x} = \mathbf{x}^*], \quad (18)$$

where $\mathbb{1}_{y \leq y^*}$ denotes the indicator function of the set $\{y \in \mathbb{R} : y \leq y^*\}$. The algorithm starts by building an ordinary random forest and then extracts the relevant weights to estimate the conditional cumulative distribution $\hat{F}(y^* | \mathbf{x}^*)$. The (conditional) quantiles are inferred through their defining equation

$$\hat{q}_\alpha(\mathbf{x}^*) := \inf\{y \in \mathbb{R} : \hat{F}(y | \mathbf{x}^*) \geq \alpha\}. \quad (19)$$

The main benefit of this approach over ordinary quantile regressors (see Section 3.3 below) is that one immediately obtains an approximation of the full conditional distribution and, accordingly, it is not necessary to train a separate model for each quantile. The downside of this method is its computational inefficiency. Because the full distribution needs to be modelled, the information about the leaf nodes of all trees needs to be stored and for every new data point \mathbf{z}^* both the conditional distribution (18) and the infimum in (19) need to be recalculated, whereas for ordinary regression forests only the mean of every leaf node has to be stored. Some numerical libraries have introduced approximations, such as random sampling in the leaf nodes (e.g. the `quantregForest` package in R), but these approximations have been reported to give considerably worse predictions [32].

Dropout networks

Dropout layers were initially introduced as a stochastic regularization technique by Hinton et al. [82]. During each forward pass at training time, a random subset of the network weights is set to zero with probability p . In essence this means that a Bernoulli prior with parameter p is placed over every weight. The idea behind this method was to reduce the “co-adaptation” behaviour of fully-connected networks, i.e., it makes it harder for the weights to work together to overfit the data. Gal and Ghahramani showed that dropout networks can be obtained as a variational approximation to deep Gaussian processes [23]. Therefore, it can be expected that, at least when normality assumptions are satisfied, this method can outperform more ad hoc approaches.

A second step in the development of dropout networks was the extension of the stochastic behaviour of the dropout layers to test time. By also randomly setting weights to zero during test time, an ensemble of different models can be obtained without having to retrain the model itself. Furthermore, because the Bernoulli priors induce a very sparse structure, MC integration can be expected to give a good approximation. A more in-depth study of stochastic regularization and Bayesian inference in dropout models is given in [22].

Kendall and Gal introduced a new method to estimate uncertainty using dropout networks [40]. Instead of having a neural network that produces a single point prediction, the network also estimates the predictive variance, thereby turning it into a *mean-variance estimator* [42, 60]. The loss function for this neural network contains a modified mean-squared error term and a penalty term for the variance:

$$\mathcal{L}_{\text{Gauss}}(\mathcal{T}) := \sum_{(\mathbf{x}, y) \in \mathcal{T}} \frac{|y - \hat{y}(\mathbf{x})|^2}{2\hat{\sigma}^2(\mathbf{x})} + \frac{1}{2} \log \hat{\sigma}^2(\mathbf{x}). \quad (20)$$

This loss function follows from the maximum likelihood principle with a Gaussian likelihood function $\mathcal{N}(\hat{y}(x), \hat{\sigma}^2(x))$ or, equivalently, by assuming that the error

$y - \hat{y}(x)$ is distributed according to $\mathcal{N}(0, \hat{\sigma}^2(x))$. Instead of maximum likelihood estimation, one could also use maximum a posteriori estimation with a Gaussian or Laplace prior, so that the common L^1 - or L^2 -penalties can be captured.

When making predictions, the conditional mean is again approximated by MC integration (12), i.e. one takes the average of multiple forward passes. The total predictive variance is given by the sum of the empirical variance of the ensemble and the variance predicted by the model itself:

$$\hat{\sigma}^2(\mathbf{x}) = \frac{1}{R} \sum_{i=1}^R \hat{y}_i(\mathbf{x})^2 - \left(\frac{1}{R} \sum_{i=1}^R \hat{y}_i(\mathbf{x}) \right)^2 + \frac{1}{R} \sum_{i=1}^R \hat{\sigma}_i^2(\mathbf{x}), \quad (21)$$

where R again denotes the ensemble size. With this approach, the predictive distribution is modelled as an approximation of a uniformly-weighted Gaussian mixture model [72], where the parameters of every component are estimated by a model in the ensemble, by a normal distribution with the same mean and variance as the mixture model.

The resulting prediction interval is virtually the same as (16):

$$\Gamma_{\text{MVE}}^\alpha(\mathbf{x}^*) := [\hat{\mu}(\mathbf{x}^*) - z^\alpha \hat{\sigma}(\mathbf{x}^*), \hat{\mu}(\mathbf{x}^*) + z^\alpha \hat{\sigma}(\mathbf{x}^*)]. \quad (22)$$

One can immediately expect that, analogous to general mean-variance estimators with a Gaussian prediction interval, this procedure does not give optimal intervals for data sets that do not follow a normal distribution. One of the consequences is that this model might suffer from the validity problems discussed in Section 3.1. This issue was already highlighted in the appendix to [23].

Deep ensembles

The idea behind deep ensembles [49] is the same as for any ensemble technique: training multiple models to obtain a better and more robust prediction. The loss functions of most (deep) models have multiple local minima and by aggregating multiple models one hopes to take into account all these minima. From this point of view the approach by Lakshminarayanan et al. [49] is very similar to that of [40]. However, the underlying philosophy is slightly different. First of all, although the same loss function is used, it is not obtained from a Bayesian framework, but rather chosen as a *proper scoring rule* [28], i.e. a loss function for distributional forecasts for which a model can never obtain a lower loss than the true distribution (it is said to be *strictly proper* if it has a unique minimum). By training a model with respect to a (strictly) proper scoring rule, the model is encouraged to approximate the true probability distribution. A disadvantage of these scoring rules is that they are only proper relative to a certain class of distributions, hence they still introduce distributional assumptions in the model. For example, the scoring rule (20) is only proper w.r.t. probability distributions with a finite second moment and strictly proper w.r.t. distributions that are determined by their first two moments. Secondly, instead of constructing an ensemble through MC sampling from the prior distribution, multiple models are independently trained and diversity is induced by using different initial parameters. This is motivated by a prior observation that random initialization often leads to superior performance when compared to other ensemble techniques [50].

A further modification introduced by Lakshminarayan et al. is the use of adversarial perturbations. Instead of using an ordinary gradient descent step

$$\theta \longrightarrow \theta - \varepsilon \nabla_{\theta} \mathcal{L}(\mathbf{X}, \mathbf{y}; \theta), \quad (23)$$

an additional contribution is added where the training instances \mathbf{x} are replaced by adversarial perturbations obtained using the Fast Gradient Sign Method [30] (this could be generalized by seeing the specific perturbation method as a hyperparameter of the model). For every instance (\mathbf{x}, y) , the perturbed version is obtained as follows:

$$\mathbf{x}' := \mathbf{x} + \boldsymbol{\eta} \odot \text{sgn}(\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, y; \theta)), \quad (24)$$

where \odot denotes elementwise multiplication. This allows the constant $\boldsymbol{\eta}$ to be a vector as to accompany features with different ranges. The modified update rule is obtained by replacing the loss function in the gradient descent step with the total loss

$$\mathcal{L}_{\text{tot}}(\mathbf{X}, \mathbf{y}; \theta) := \mathcal{L}(\mathbf{X}, \mathbf{y}; \theta) + \mathcal{L}(\mathbf{X}', \mathbf{y}; \theta), \quad (25)$$

where \mathbf{X}' denotes the design matrix obtained from the perturbed features.

Prediction intervals are constructed as in the previous section, i.e. a (conditionally) normal distribution is assumed and the intervals are given by Eq. (22). It was observed that this architecture shows improved modelling capabilities and robustness for uncertainty estimation. In [20] the improved performance is attributed to the multimodal behaviour in model space. Fort et al. argue that most networks tend to have the property that they only work in one specific subspace of the model space, while deep ensembles (due to random initialization) can model different modes.

Without the adversarial training, this model is similar to the one introduced by Khosravi et al. [44]. However, instead of training an ensemble of mean-variance estimators, an ensemble of point estimators is trained to predict y and in a second step a separate estimator $\hat{\sigma}$ for the data noise is trained using loss function (20), where the ensemble estimator is kept fixed.

3.3 Direct interval estimation methods

The class of direct interval estimators consists of all methods that are trained to directly output a prediction interval. Instead of modelling a distribution or extracting uncertainty from an ensemble, they are trained using a loss function that is specifically tailored to the construction of prediction intervals. The general structure of this approach is summarized in Algorithm 3. Because these methods are specifically made for estimating uncertainty, they can be expected to perform better than modified point estimators. However, this is also immediately their main disadvantage: they do not always produce a point estimate. Another disadvantage is that they are in general also specifically constructed for a predetermined confidence level α . Choosing a different value for α requires the model to be retrained.

Algorithm 3: Direct interval estimationframework *Direct* ($\mathcal{L}, \mathcal{A}, \mathcal{D}$)**Input** : Model architecture \mathcal{A} , loss function \mathcal{L} , data set \mathcal{D} **Output**: Interval estimator Γ

- 1 Construct a model Γ with architecture \mathcal{A}
- 2 Train the model Γ on \mathcal{D} w.r.t. the loss function \mathcal{L}
- 3 (Optional:) Post-process Γ as to obtain an interval estimator
- 4 **return** Γ

Quantile regression

Instead of estimating the conditional mean of a distribution, quantile regression models estimate the conditional quantiles. This has the benefit that one immediately obtains measures of the spread of the underlying distribution and, therefore, of the confidence of future predictions. A neural network quantile regressor \hat{q}_α for the α -quantile can easily be implemented by replacing the common mean-squared error loss by the following *pinball loss* or *quantile loss* [46]:

$$\mathcal{L}_{\text{pinball}}(\mathcal{T}) := \sum_{(\mathbf{x}, y) \in \mathcal{T}} \max \left((1 - \alpha)(\hat{q}_\alpha(\mathbf{x}) - y), \alpha(y - \hat{q}_\alpha(\mathbf{x})) \right). \quad (26)$$

This loss function tries to balance the number of data points below and above the (estimated) quantile. It can be obtained by reformulating the definition of quantiles as an optimization problem:

$$q_\alpha = \arg \min_{q \in \mathbb{R}} \left((1 - \alpha) \int_{-\infty}^q (q - y)p(y) \, dy + \alpha \int_q^{+\infty} (y - q)p(y) \, dy \right). \quad (27)$$

By differentiating the argument on the right-hand side with respect to q and equating it to 0, one obtains definition (19) of the α -quantile. The pinball loss (26) is then simply the loss function for the sample α -quantile, i.e. the α -quantile of the empirical distribution function. An intuition for this loss function can be gained from considering the example of the sample median ($\alpha = 0.5$). The median ξ is defined as the data point for which there are as many positive as negative residuals $r_i := y_i - \xi$. From an optimization point of view, this definition is equivalent to that of the minimizer of the average of all absolute residuals. The pinball loss estimates other quantiles by replacing this average with a weighted average.

To estimate multiple quantiles at the same time, one simply has to add multiple output nodes to the network and add terms of the form (26) for every quantile. Although quantile regression models do not produce a point estimate in the form of a conditional mean, they can give a prediction of the conditional median. These models do, however, suffer from a problem of underrepresented tails of the distribution. As for every neural network, the performance becomes worse when less data is available. Because the quantiles that are relevant for interval estimation are often located in the tails of the distribution, especially for very low significance levels α , the quantile regressors might not give optimal results.

High-Quality principle

As stated before, there are two quantities that are mainly used to evaluate the performance of interval estimators: the degree of coverage (1) and the average size of the prediction intervals (3). The idea that optimal prediction intervals should saturate inequality (2) and minimize the average size was dubbed the *High-Quality* (HQ) principle by Pearce et al. [67, 68]. The idea to construct a loss function based on the HQ principle was first proposed by Khosravi et al. in [43]. There, the Lower-Upper Bound Estimation (LUBE) network, predicting the lower and upper bounds of the prediction interval, was trained using the following loss function:

$$\mathcal{L}_{\text{LUBE}}(\mathcal{T}) := \frac{\text{MPIW}}{r} (1 + \exp(\lambda \max(0, (1 - \alpha) - \mathcal{C}))), \quad (28)$$

where \mathcal{C} denotes the coverage degree (1) estimated on the training set \mathcal{T} and, similarly, the Mean Prediction Interval Width (MPIW) is the average width (3) estimated on \mathcal{T} :

$$\text{MPIW} := \frac{1}{|\mathcal{T}|} \sum_{(\mathbf{x}, y) \in \mathcal{T}} |u(\mathbf{x}) - l(\mathbf{x})|. \quad (29)$$

The numbers r and λ are respectively the range of the response variable and a constant determining how much a deviation from optimal coverage should be penalized. The main idea is that the penalty should be proportional to the size of the intervals and that the penalty should be greater if the desired coverage is not obtained.

Although intuitively sensible, this method admits no true theoretical derivation. It is derived solely from heuristic arguments. Later, Pearce et al. introduced an alternative to the LUBE loss where some of the ad-hoc choices were formalised [68]. The most important modifications are the replacement of MPIW by *captured* MPIW, i.e. only the lengths of the intervals that contain the true value are taken into account to avoid artificially shrinking already invalid intervals, and replacing the exponential function by a term derived from a likelihood principle (the number of captured points can be modelled by a binomial distribution). The resulting loss function is

$$\mathcal{L}_{\text{QD}}(\mathcal{T}) := \text{MPIW}_{\text{capt}} + \lambda \frac{n}{\alpha(1 - \alpha)} \max(0, (1 - \alpha) - \mathcal{C})^2, \quad (30)$$

where n is the total number of data points and λ is a Lagrange multiplier determining the relative importance of the two contributions. Although this loss function is slightly more substantiated, it remains heuristic in nature and, hence, it does not admit any theoretical guarantees regarding its performance.

3.4 Conformal prediction

In this section a framework is considered that can both turn any point predictor into an interval estimator and turn any existing interval estimator into a valid estimator. As originally introduced [24, 74, 88], the *conformal prediction* (CP) framework allows for the construction of valid prediction regions given a data set and

a *nonconformity measure* that evaluates how different a new instance is from the given data set. The original implementation suffered from a high computational overhead because in general all calculations, including the training of any underlying models, had to be redone for every data point. A solution, called *inductive conformal prediction* (ICP) or *split-conformal prediction*, where the training phase of the underlying models was decoupled from the “conformalization” phase, was later introduced [63,64]. Although the theoretical guarantees for ICP are not as strong as for the original *transductive* framework, the computational speed-up outweighs this drawback and, therefore, only ICP will be considered in this study.

A short introduction to the inductive conformal prediction framework, along the lines of [89], is in order. It will be assumed that the following data has been provided:

- 1) a data set \mathcal{D} , and
- 2) a function $A : \mathcal{Z} \rightarrow \mathbb{R}$ (meeting some technical conditions).

The real-valued function A is called the *nonconformity measure* in the conformal prediction literature (note that it is not a measure in the mathematical sense). A common example is the following one:

$$A_{\text{point}}(\mathbf{x}, y) := |y - \hat{y}(\mathbf{x})|, \quad (31)$$

where \hat{y} can be any point predictor. Although it is the most frequently used type of nonconformity measure, not all nonconformity measures of relevance are of this form (see the next section).

Given a nonconformity measure A , possibly depending on a training set $\mathcal{T} \subset \mathcal{D}$, the framework proceeds by constructing a calibration “curve” using the complement $\mathcal{V} := \mathcal{D} \setminus \mathcal{T}$. For all $z_i \in \mathcal{V}$ the associated nonconformity score is defined as follows:

$$\alpha_i := A(\mathbf{z}_i). \quad (32)$$

The interval estimator Γ^α , associated with A and \mathcal{D} , is then defined as the function that constructs a region $\Gamma^\alpha(\mathbf{x}^*)$ for all $\mathbf{x}^* \in \mathcal{X}$ consisting of all $y^* \in \mathbb{R}$ for which the nonconformity score $A(\mathbf{z}^*)$ is smaller than the empirical $((1 - \alpha)(1 + \frac{1}{|\mathcal{V}|}))$ -quantile of $\{\alpha_i : \mathbf{z}_i \in \mathcal{V}\}$. The general procedure for constructing Γ^α is shown in Algorithm 4. The following theorem, which can be found in most of the conformal prediction literature, is the reason why this framework is so powerful.

Theorem (Marginal validity) *If the nonconformity scores $\{a_i : \mathbf{z}_i \in \mathcal{V}\} \cup \{\alpha^*\}$ are exchangeable for all validation sets \mathcal{V} and test points \mathbf{z}^* , the estimator Γ^α is (conservatively) valid:*

$$\text{Prob}\{y \in \Gamma^\alpha(\mathbf{x} \mid \mathcal{V})\} \geq 1 - \alpha, \quad (33)$$

where the probability is taken over both \mathbf{x} and \mathcal{V} . The estimator has been denoted by $\Gamma^\alpha(\cdot \mid \mathcal{V})$ to explicitly show the dependence of Algorithm 4 on the calibration set \mathcal{V} .

Algorithm 4: Inductive Conformal Prediction

function ICP ($\alpha, \mathcal{T}, \mathcal{V}, A$)**Input** : Significance level α , nonconformity measure A , training set \mathcal{T} and calibration set \mathcal{V} **Output:** Interval estimator Γ^α

```

1 (Optional) Train the underlying model of  $A$  on  $\mathcal{T}$ 
2 foreach  $(\mathbf{x}_i, y_i) \in \mathcal{V}$  do
3   | Apply (32):  $\alpha_i \leftarrow A(\mathbf{x}_i, y_i)$ 
4 end
5 Determine the critical value of  $A$ :  $\alpha^* \leftarrow ((1 - \alpha)(1 + \frac{1}{|\mathcal{V}|}))$ -quantile of  $\{\alpha_i : (\mathbf{x}_i, y_i) \in \mathcal{V}\}$ 
6 Construct an interval estimator as follows:

7   procedure  $\Gamma^\alpha(\mathbf{x} : \mathcal{X})$ 
8     | foreach  $y \in \mathbb{R}$  do
9       | Apply (32):  $\alpha_y \leftarrow A(\mathbf{x}, y)$ 
10    | end
11    | return  $\{y \in \mathbb{R} : \alpha_y \leq \alpha^*\}$ 

12 return  $\Gamma^\alpha$ 

```

In practice, one only has access to a single (calibration) data set and, consequently, one cannot take the expectation value over all possible data sets as is necessary to obtain marginal validity. However, for large calibration sets, the following theorem shows that one can expect approximately valid prediction intervals.

Theorem (Asymptotic validity) *If the nonconformity scores $\{a_i : \mathbf{z}_i \in \mathcal{V}\} \cup \{\alpha^*\}$ are exchangeable for all test points \mathbf{z}^* , the estimator Γ^α becomes (conservatively) valid when $|\mathcal{V}| \rightarrow \infty$.*

The exchangeability assumption [89] is the reason why the training and calibration steps in Algorithm 4 are performed on disjoint data sets. Note that this assumption is strictly weaker than being i.i.d. and will in most cases be, at least approximately, satisfied.

An important remark regarding this framework is that it is nonparametric and almost completely assumption-free. Except for the exchangeability assumption, this framework can work with any choice of data set, predictor or nonconformity measure. This not only makes it extremely versatile and applicable in almost any setting, it also assures that there can be no negative consequences coming from broken assumptions, something many other models do suffer from (see also further below in the section “Validity issues”). However, although virtually any function can act as a nonconformity measure, the specific choice can strongly influence the size of the resulting prediction regions.

Point predictors

The nonconformity measures defined in (31) are important due to the fact that they allow to turn a standard point predictor into a (valid) interval estimator. Because this choice of nonconformity measure only considers differences in the target

space, i.e. there is no \mathcal{X} -dependent scaling, the algorithm produces prediction intervals of constant size. This implies that the critical value from Algorithm 4 needs to be calculated only once and, therefore, that the prediction interval for a new instance \mathbf{x}^* can be constructed with constant time complexity as follows:

$$\Gamma_{\text{point}}^{\alpha}(\mathbf{x}^*) := [\hat{y}(\mathbf{x}^*) - \alpha^*, \hat{y}(\mathbf{x}^*) + \alpha^*], \quad (34)$$

where \hat{y} is the given point predictor.

The above procedure gives uniform (or homoscedastic) prediction intervals, which is in stark contrast with most bona fide interval estimators. Although computationally simple, it ought to be clear that this is not the generic situation. Different modifications to obtain heteroscedastic models have been proposed in the literature, the main one being to *normalize* [62] the nonconformity measure by a *dispersion* function $\sigma : \mathcal{X} \rightarrow \mathbb{R}$. This modification is often called *locally adaptive* or *normalized conformal prediction*. In the case of mean-variance estimators this could be the estimate of the predictive variance (or standard deviation), while for other models a popular choice is the empirical uncertainty of a data subset found using for example a k -nearest neighbour algorithm [65]. It should be clear that the latter choice, although interesting because it takes into account the number of similar data points and, hence, the data uncertainty, does introduce extra computational overhead.

Ensemble methods that make use of bagging, such as random forests, give rise to their own natural choice of both homoscedastic and heteroscedastic nonconformity measures. As mentioned in Section 3.2, with bagging, every individual subestimator only uses a subset of the training set and, therefore, one can use for every training instance the subensemble for which this instance was not used during training to produce an independent prediction. This allows the full data set to act as both a training and calibration set in stark contrast to the situation of ordinary CP methods, which are restricted by the exchangeability assumption. To this end, the following modified nonconformity measure was introduced by Johansson et al. [38]:

$$A_{\text{cal}}(\mathbf{x}_i, y_i) := |y_i - \hat{y}_{(i)}(\mathbf{x}_i)| \quad (35)$$

$$A_{\text{test}}(\mathbf{x}^*, y^*) := |y^* - \hat{y}(\mathbf{x}^*)|, \quad (36)$$

where $\hat{y}_{(i)}$ denotes the OOB estimator for the data point (\mathbf{x}_i, y_i) . Although it cannot be formally proven that this procedure satisfies the same validity properties as ordinary ICP, the fact that OOB estimates often overestimate the uncertainty make it plausible that the above measure will lead to (conservatively) valid prediction intervals. One should also note that this approach recovers the estimate (17) introduced in [95] whenever the residuals for the calibration set \mathcal{V} are distributed symmetrically around zero. To obtain formal guarantees, the authors later showed that this method can be modified [5]. Instead of using the asymmetric nonconformity measure above, they proposed the following scheme for every new observation \mathbf{x}^* :

1. Choose a random training instance \mathbf{z}_i .
2. Use the OOB predictor $\hat{y}_{(i)}$ to calculate the nonconformity score of a candidate \mathbf{z}^* (in the process also producing a point estimate).

3. Use the calibration scores of the complementary calibration set $\mathcal{T} \setminus \{\mathbf{z}_i\}$ to derive the prediction region for \mathbf{x}^* .

Although this method is theoretically valid, it does require the recalculation of the critical value for every new data point, thereby introducing some extra computational overhead. Furthermore, following from the fact that the point prediction is only made using a fraction of the total forest, the predictive performance can be expected to be worse (especially for small forests).

Similar to how the interval estimator (17) for bagged ensembles could be extended to a general cross-validation inspired framework [2], the above approach for conformal prediction can also be extended to a framework in which the full data set is used to calibrate the model through a k -fold or leave-one-out approach [87]. For this reason they are called *cross-conformal predictors*. As shown in [2], these methods always produce confidence regions that are contained in the intervals coming from jackknife+ estimators, with the added complexity that they are not necessarily connected, i.e. they might be disjoint unions of intervals.

Validity issues

There exists a wide variety of methods that are able to handle the task of generating prediction intervals, but most of them suffer from a major problem. In practice the coverage defined in (1) will not always satisfy the validity requirement (2), which implies that the predictions of these methods are not as meaningful as desired. They can still be used as a heuristic uncertainty measure, but they lack interpretability.

There are two main factors that contribute to this problem. First of all, as mentioned before, due to the complexity of present-day data sets, common assumptions can be violated. A simple example is the case of empirical averaging (16), where the standard deviation will only give rise to calibrated intervals if the residuals follow a Gaussian distribution. Secondly, it is possible that the interval estimators are only asymptotically valid and that the size of the data set is not sufficiently large. Since the validity of typical statistical techniques is only guaranteed to hold for infinitely large data sets, there is no guarantee that these will hold for $|\mathcal{D}| \ll \infty$.

In the context of classification problems, where especially the former issue plays a role [31], a wide variety of calibration methods is available: Platt scaling, temperature scaling, isotonic regression, etc. In general these methods take the output distribution of the trained predictor and modify it such that it becomes (approximately) well calibrated. However, these methods explicitly use the fact that the target space is finite in those cases. For regression problems, these methods can therefore not be applied in general. A convenient solution is given by the conformal prediction framework introduced above. It cannot only be used to turn a point predictor into a valid interval estimator, but it can also be used to calibrate an existing model.

Calibrating prediction intervals

In this section the models that predict the lower and upper bounds of prediction intervals are considered, for example the $\alpha/2$ - and $(1 - \alpha/2)$ -quantile estimates

for a given significance level α . For this class of estimators a reasonable choice of nonconformity measure is

$$A_{\text{PI}}(\mathbf{x}_i, y_i) := \max \left(\hat{l}(\mathbf{x}_i) - y_i, y_i - \hat{u}(\mathbf{x}_i) \right), \quad (37)$$

as originally introduced for quantile regression by Romano et al. in [73]. As was the case for point predictors and Eq. (31), this function only considers the distance between points in the target space and, therefore, the critical value of the ICP Algorithm 4 will be independent of the underlying features. This again allows one to compute the critical value only once and construct a modified prediction interval for every new instance \mathbf{x}^* as follows:

$$\Gamma_{\text{int}}^{\alpha}(\mathbf{x}^*) := \left[\hat{l}(\mathbf{x}^*) - \alpha^*, \hat{u}(\mathbf{x}^*) + \alpha^* \right]. \quad (38)$$

The idea behind this construction is very similar to that for point predictors. One estimates the amount by which the constructed intervals are too small (or wide) on average on the calibration set and corrects for these errors in the future. Two other normalized nonconformity measures were considered in [45, 76]. However, it was empirically observed that the inherent heteroscedasticity of quantile regression models is already sufficient and that the added normalization only increases the size of the prediction intervals without gaining much in terms of validity.

In the section on quantile regression it was noted that this approach tends to have a problem with modelling the tails of the distribution, with the added consequence that this can influence the validity at extreme significance levels. However, when combining such models with conformal prediction, validity is not an issue. Therefore, Romano et al. [73] introduced a “softening” factor w such that the underlying quantile regression model only has to estimate the $w\alpha/2$ - and $(1 - w\alpha/2)$ -quantiles. This avoids the problems with underrepresented tails, while preserving the validity of the model (potentially excessively broadening the intervals if the smoothing parameter is not correctly tuned).

3.5 Summary

In the preceding four sections, we introduced different classes of interval estimators, each having its own characteristics. In this section, we summarize the main properties for clarity and convenience. We identify four properties that are important for practical purposes. The first one is the main notion of this paper, namely validity, i.e. whether a model is guaranteed to produce (approximately) valid prediction intervals. Since none of the methods are valid for any finite-size data set, we only consider validity in the sense of the Marginal Validity Theorem of Section 3.4. The second property, which is becoming more and more important due to the increase in data size and complexity, is scalability. A third feature, which was not highlighted in the foregoing sections, but which might be important for certain applications is the possibility to include domain knowledge. Last but not least we indicate whether it is necessary to include a validation set, which is especially critical when available data is limited. The summary is shown in Table 1. For the four classes of methods, we also list the main references.

Method	Marginal validity	Scalability	Domain knowledge	Validation set
Bayesian methods [3, 4, 29, 34, 92, 93]	No (except for exact inference with correct priors)	Only scalable with approximate inference	Yes	No
Ensemble methods [35, 40, 49, 55, 60, 91]	No	Yes (when scalable models are used)	No	No
Direct interval estimation [46, 67]	No	Yes	No	Yes
Conformal prediction [2, 76, 87, 89]	Yes	Yes (for ICP)	No	Yes

Table 1: Summary of method characteristics and references.

4 Experimental setup

In this and the following section some of the models introduced above are experimentally investigated. They are evaluated and compared based on some general performance measures. Moreover, some general conclusions that can be used in future applications or research are derived.

4.1 Quality metrics

An optimal interval estimator should satisfy some conditions. To assess the quality of the models, the HQ principle from Section 3.3 is adopted. First of all a model ought to be valid (or calibrated) in the sense of Eq. (2). The more a model deviates from being well calibrated, the less reliable it becomes since the results cannot be trusted and relied upon. At the same time the results should also be interpretable, i.e. the end-user should be able to use the uncertainty information to make further decisions. The most obvious interpretation would be where small prediction intervals correspond to situations where the model is very confident. Therefore, the models are also compared based on the width of the prediction intervals. This is closely related to the fact that inequality (2) should not only be satisfied, but should even be saturated. To this end, the following two assessment measures are used:

1. Satisfy or even saturate the validity condition (2). This is measured by approximating the coverage probability (1) by a finite sum over the test set:

$$\mathcal{C}(\Gamma, P) \approx \frac{1}{|\mathcal{D}^*|} \sum_{(\mathbf{x}, y) \in \mathcal{D}^*} \mathbb{1}(y \in \Gamma(\mathbf{x})). \quad (39)$$

2. Minimize the average width of the intervals, as defined in Eq. (3):

$$\mathcal{W}(\Gamma, P) \approx \frac{1}{|\mathcal{D}^*|} \sum_{(\mathbf{x}, y) \in \mathcal{D}^*} |u(\mathbf{x}) - l(\mathbf{x})|. \quad (40)$$

Aside from the above quality measures, some other quantities might be of importance depending on the application. In the experiments only the predictive power is considered. The benefit of working with models that are built upon or include a point predictor is that one also gets a direct estimate of the response variable. Since this is important in many situations, the R^2 -coefficients are reported (as noted in the section on quantile regression, this method cannot estimate the conditional mean, instead the conditional median is estimated).

4.2 Data

Most of the data sets were obtained from the UCI repository [14]. Specific references are given in Table 2. This table also shows the number of data points and (used) features and the skewness and (Pearson) kurtosis of the response variable. All data sets were standardized (both features and target variables) before training. The data sets **blog** and **fb1** were also analysed after first taking a log transform of the response variable because these data sets are extremely skewed, which is reflected in the high skewness and kurtosis, as shown in the fourth column of Table 2, and are believed to follow a power law distribution. This strongly improved the R^2 -coefficient of the various models, but did not improve the prediction intervals, and therefore, these results are not included. The **crime** data set comes in two versions: the original data set consists of integer-valued data (count data), while the version used here was preprocessed using an unsupervised standardization algorithm [71]. Although standardized, the data set retains (some of) its count data properties. The **traffic** data set, aside of being very small, is also extremely sparse (on average 14 features are zero). It should be noted that all of the data sets used in this study were considered as ordinary (static) data sets. Even though some of them could be considered in a time series context, no autoregressive features were additionally extracted. The main reason to exclude autoregressive features is that most, if not all, methods considered in this study assume the data to be i.i.d. (or exchangeable), a property that is generically not valid for autoregressive data.

4.3 Architectures and training

All neural networks were constructed using the default implementations from **PyTorch** [66]. The general architecture for all neural-network-based models was fixed. The *Adam* optimizer was used for weight optimization with a fixed learning rate of 5×10^{-4} , in accordance with [73]. The number of epochs was limited to 100, unless stated otherwise. All neural networks contained only a single hidden layer with 64 neurons. The activation functions after the first layer and the hidden layer were of the ReLU type, while the activation function at the output node was simply a linear function.

Name	# samples	# features	Skewness / Kurtosis	Source
concrete	1030	8	0.42 / 2.68	[94]
naval	11934	14	0.0 / 1.8	[12]
turbine	9568	4	0.31 / 1.95	[39]
puma32H	8192	32	0.02 / 3.04	[13]
residential	372	105	1.26 / 5.15	[70]
crime2	1994	123	1.52 / 4.83	[71]
fb1	40949	54	14.29 / 301.44	[78]
blog	52397	280	12.69 / 235.3	[8]
traffic	135	18	1.07 / 3.74	[18]
star	2161	39	0.29 / 2.63	[1]

Table 2: Overview of the data sets.

For each of the four classes of interval estimators in Section 3, at least one example was chosen for a general comparison. Furthermore, to handle calibration issues, conformal prediction was chosen as a post-hoc method due to its nonparametric and versatile nature. Every model that produces a prediction interval (or an uncertainty estimate) is also re-examined after *conformalization*. Hyperparameter optimization, if applicable, was performed on a validation set containing 5% of the training samples, unless stated otherwise. The specific architectures are as follows:¹

1. Bayesian methods:

- (i) Gaussian Process: As kernel a standard radial basis function (RBF kernel) was used. This automatically incorporates out-of-distribution uncertainty. The default implementation from the `GPyTorch` library [25] was used. This library provides a multitude of different approximations and deep learning adaptations for Gaussian processes. A heteroscedastic likelihood function with trainable noise parameter was used and the model was trained for 50 epochs (following the library’s heuristics). Optimization of the noise parameter was handled internally, so no additional validation set was required.
- (ii) Approximate GP: Because in most present-day cases the data sets are too large to efficiently apply Gaussian processes, a variational approximation was included. The Stochastic Variational GP approximation [34], implemented in `GPyTorch`, was chosen. The same likelihood function and number of epochs were used as for the exact model.

2. Ensemble methods:

- (i) Dropout ensemble: An ensemble average of 50 MC samples was used. An early stopping criterion, based on the minimum of the loss function, was employed to avoid overfitting. The dropout probability was optimized over the interval $[0.05, 0.5]$ with steps of 0.05.

¹ All code is made publicly available at <https://github.com/nmdwolf/ValidPredictionIntervals>.

- (ii) Mean-variance estimator: The Gaussian dropout-based mean-variance estimator from [40] was chosen because this model extends the dropout ensemble by explicitly incorporating uncertainty. An ensemble average of 50 MC samples was used. The dropout probability was optimized over the interval $[0.05, 0.5]$ with steps of 0.05.
- (iii) Deep ensembles: The ensemble consisted of 5 estimators and the adversarial step size equalled 0.01 times the range of the corresponding dimension (cf. [49]). At training time L^2 -regularization with $\lambda = 10^{-6}$ was applied.

3. Direct interval estimation methods:

- (i) Neural network quantile regression: The same softening factor $w = 2$ as in [73, 76] was used. The early stopping criterion for neural networks was also modified to work with the average length and coverage degree of the prediction intervals instead of the loss function of the network. Dropout with $p = 0.1$ and L^2 -regularization with $\lambda = 10^{-6}$ were applied at training time.

4. Conformal prediction:

- (i) Neural network: A standard neural network point predictor was chosen as a baseline model. Early stopping as for the ensemble methods was used as a regularization method. Furthermore, dropout with $p = 0.1$ and L^2 -regularization with $\lambda = 10^{-6}$ were applied at training time.
- (ii) Random forest: The default implementation from `scikit-learn` [69] was used. The number of trees was fixed at 100, while all other hyperparameters were left at their default value.

In the remainder of the text, the selected models are abbreviated as follows:

Model	Abbreviation	Model	Abbreviation
Neural network	NN	Random forest	RF
Quantile regressor	QR	Deep ensemble	DE
Dropout ensemble	Drop	Gaussian dropout-MVE	MVE
Exact Gaussian Process	GP	Stochastic Variational GP	SVGP

All experimental results were obtained by evaluating the models on 50 different train/test-splits of the data sets in Table 2. The test set always contained 20% of the data. If a calibration set was needed for post-hoc calibration, the training set was further divided into two equal-sized sets. Because the models were tested both with and without post-hoc calibration, this approach also allows to see how the performance changes when halving the size of the training set. The confidence level was fixed at 0.9.

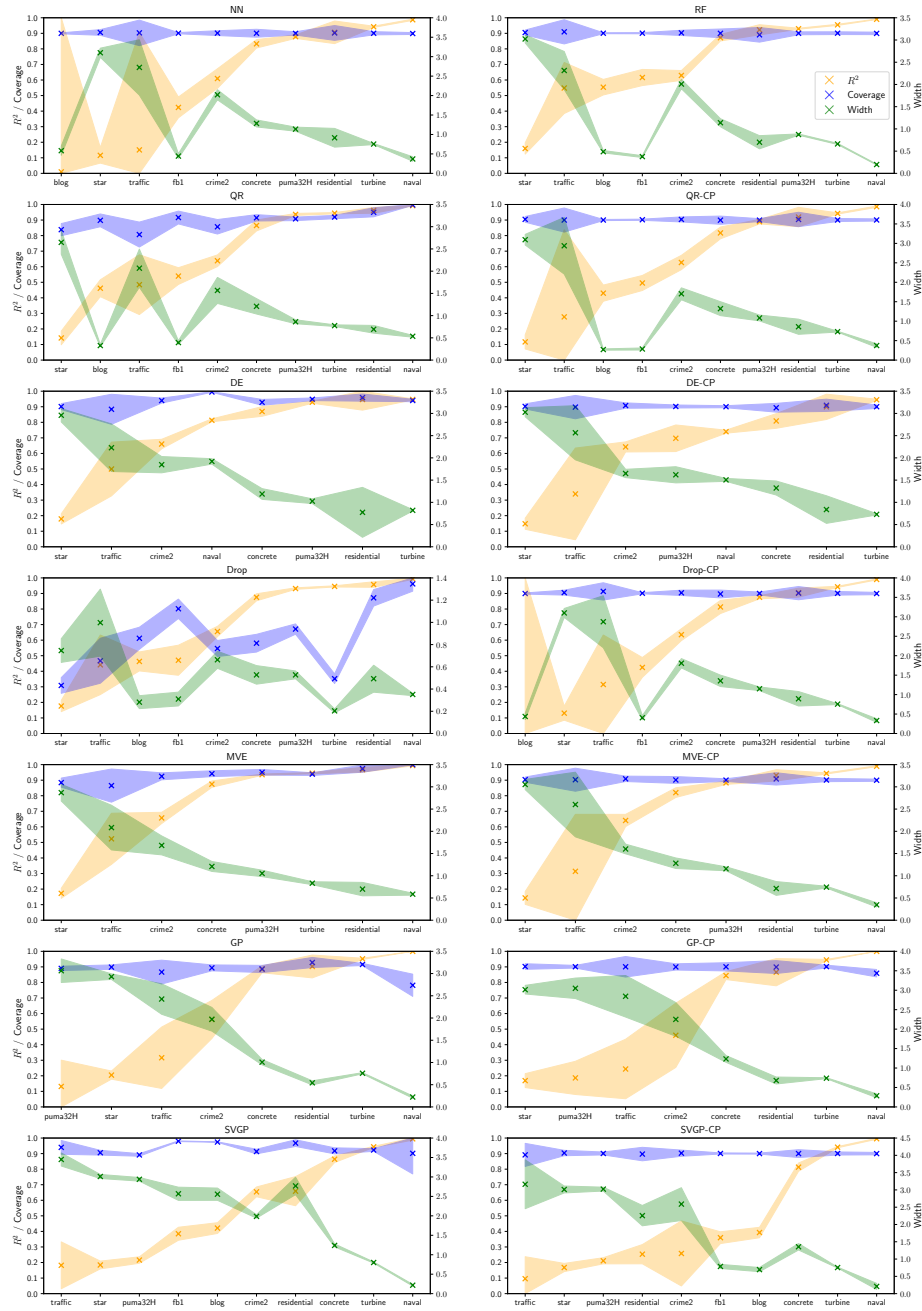


Fig. 1: Relation between confidence and predictive power. For each model, the R^2 -coefficient, coverage and average PI width are shown for all data sets in orange, blue and green, respectively.

Shaded areas represent the standard deviation. For each model, the data sets are sorted along the x-axis according to the R^2 -coefficient (in ascending order).

5 Results and discussion

In Fig. 1, both the coverage degree, average width and R^2 -coefficient are shown. For each model, the data sets are sorted according to increasing R^2 -coefficient (averaged over the different runs). For the exact Gaussian processes (GP), no results on the data sets **fb1** and **blog** are reported due to the large memory and time consumption. While these sets are not extraordinary compared to modern real-life data sets, they already require a considerable amount of memory due to the quadratic scaling. Approximate models, such as the variational approximation used in this study, will become imperative. The first row of Fig. 1 consists of the two conformalized point predictors (these could be considered baseline models). For the other rows, the left column shows the results of the models trained on the full data set, while the right column shows those of the calibrated (conformalized) models. When comparing the two columns, it is immediately clear that the coverage, indicated by the blue regions, is much more concentrated around the nominal value of 0.9 for the conformalized models, as is guaranteed by the Marginal Validity Theorem from Section 3.4. The shaded regions indicate the variability among the 50 different runs (standard deviation is used for Fig. 1). It is clear that a higher variability in the predictive power often corresponds to a higher variability in the interval quality (both coverage and average width). This is not surprising since a model in general performs worse when the uncertainty is higher. As most of the models explicitly use the predictions to build prediction intervals, this relation can be expected to be even stronger. For both the DE and MVE models the results for the **fb1** and **blog** data sets are missing because the average widths differed by about two orders of magnitude compared to the other data sets and models and were, therefore, deemed to be nonsensical. The R^2 -coefficient for the NN and Drop-CP models on the **blog** data set show a strong variability, while the average width remains small and almost constant. This can be explained by the strong skewness present in the data set. The R^2 -coefficient is sensitive to the extreme outliers, while the prediction intervals are not, as long as the outlier proportion is less than α . This also explains why almost all models give reasonably good intervals for both the **fb1** and **blog** data sets. The DE and MVE models form the exception, as mentioned above, since these methods inherently take into account the data uncertainty and cannot discard these outliers. A general summary of the results can be found in the tables in Appendix A.

In Fig. 2 the coverage degree and the average interval width are shown for a selection of data sets (**crime2**, **turbine**, **fb1** and **star**). Except for NN and RF, every model is reported in three variants. The first one is trained on half of the data set, the second one is trained on all of the data set and the third one is conformalized using a 50/50 split. As before, the average widths for MVE and DE are not shown for the data set **fb1** because these were too large to be sensible. The associated R^2 -coefficients are shown in Fig. 3. The results for the conformalized models are the same as for the those trained on half of the data set, since conformal prediction is a post-hoc method. Therefore, only the fully-trained and conformalized models are shown. Here again it is clear that the uncalibrated models do not approximately saturate the validity constraint. They either underestimate the uncertainty or produce overconservative prediction intervals. When comparing between the models trained on half of the data set and the full data set, it is generally the case that the

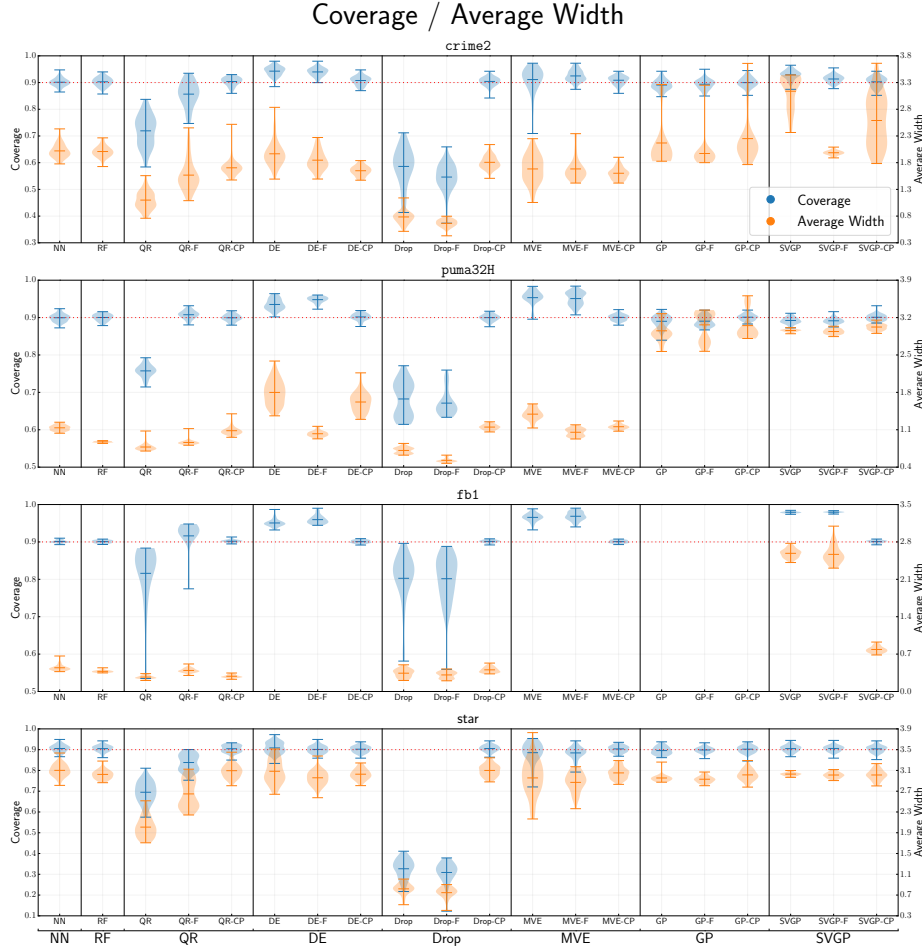


Fig. 2: Prediction interval properties for a selection of data sets. Blue and orange regions respectively show the coverage degree and the average interval width. For every model 3 variants are shown (except for NN and RF). Fully trained and conformalized models are indicated by F and CP, respectively. The red dashed line indicates the nominal coverage level.

fully-trained model achieves better results in terms of both coverage and average width. The models incorporating a probabilistic prior (DE, MVE and GP) come out as the most stable ones w.r.t. data size change.

For each of the selected models, Fig. 4 shows the best five models in terms of average width, excluding those that do not (approximately) satisfy the coverage constraint (2). This figure shows that there is quite some variation in the models. There is not a clear best choice. Because on most data sets the models produce uncalibrated prediction intervals, post-hoc calibrated models often produce the best intervals. However, when the models already produce (approximately) calibrated




Fig. 3: R^2 -coefficients for a selection of data sets. For every model 2 variants are shown (except for NN and RF): trained on the full data set and conformalized. The conformalized models are indicated by CP.

intervals from the start, they tend to give tighter intervals than their conformalized counterparts because there is more data available for training. It is also found that the models incorporating a Gaussian assumption (GP, SVGP, MVE and DE) fail to deliver optimal results in the case of non-Gaussian data (e.g. the **fb1** data). The desired confidence level can be attained, but there is a large trade-off with the average width. As noted before, the models (MVE and DE) using scoring rule (20) sometimes produce intervals where the average width is some orders of magnitude larger than for the other models. Leaving out the inherent uncertainty contribution to the variance (the third term in Eq. (21)) resolves this problem, but then the model loses its ability to approximately reach the required coverage level, similar to the ordinary dropout networks. However, as stated in the previous section, the data sets **fb1** and **blog** are strongly skewed. After reducing the skewness by discarding data points greater than a certain threshold, the average width also strongly decreases. This observation, combined with the fact that it are exactly the models explicitly incorporating data (or *aleatoric* [40]) uncertainty

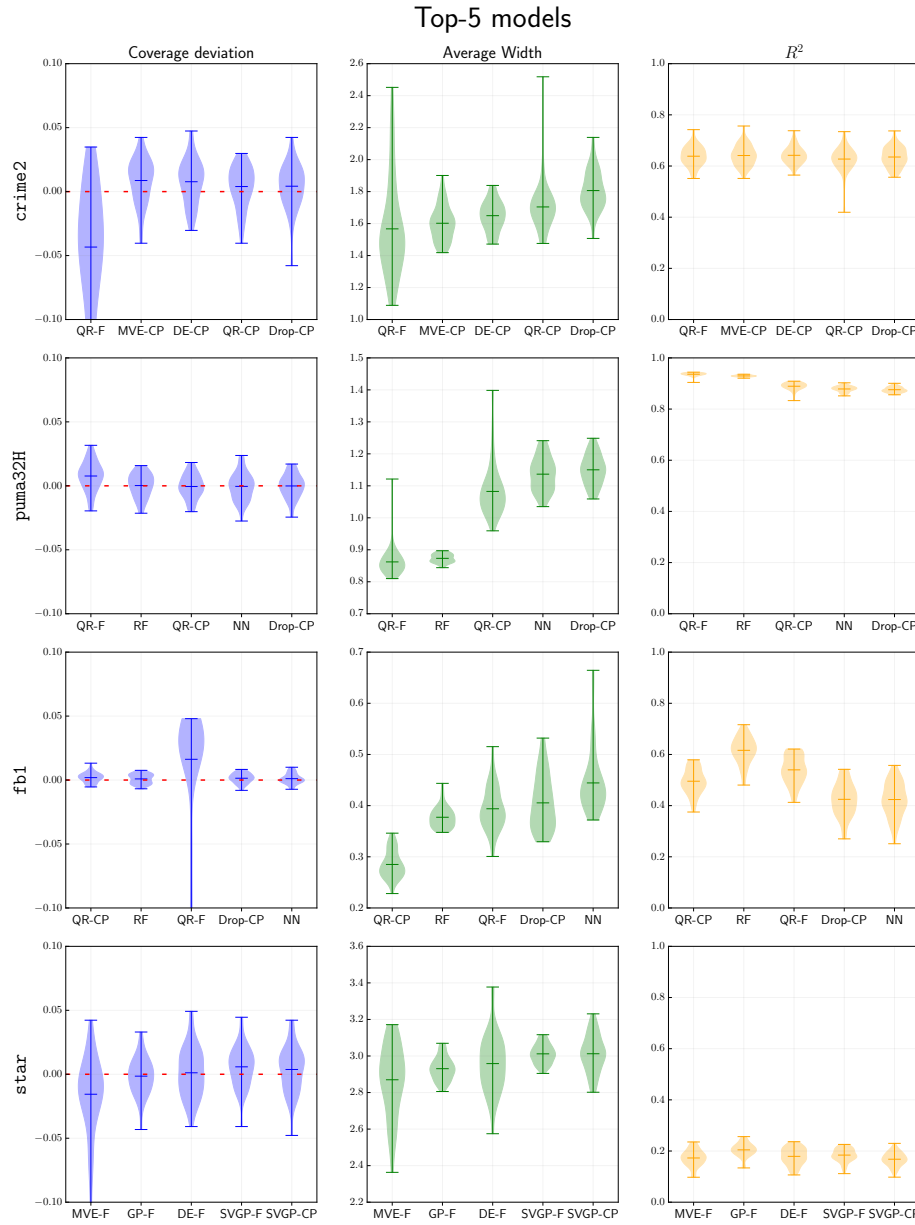


Fig. 4: Overview of the best models based on average width. The left column shows the deviation of the coverage from the desired level ($\alpha = 0.9$), the middle column shows the average width of the prediction intervals and the right column shows the R^2 -coefficient of the models. Error bars represent the standard deviation over train-test splits.

that suffer from extreme prediction intervals, might also point to a different interpretation. Namely, in the case of outliers, the model has almost no information of the conditional distribution, hence this corresponds to a situation of almost full uncertainty. From this point of view, an extremely wide prediction interval might be more sensible than a mild extrapolation of the lower uncertainty on the majority of the data points.

Because all data sets were standardized, the average widths of the prediction intervals can also be compared across data sets. From this perspective, the models seem to give significantly better results on the **fb1** data set. However, when making such a comparison, one should also look at the distribution of the data itself. The difference between the $(1 - \alpha/2)$ - and $\alpha/2$ -quantiles gives a rough estimate of the interval widths. The values of this quantity for the four selected data sets, averaged over the train-test splits, are given by (3.17, 3.49, 0.86, 3.24). Similar to the *relative absolute error*, one could define a “relative width” as the ratio of the average width as estimated by the given model and the quantile difference obtained from the unconditional data distribution, i.e. the interval width one would obtain when ignoring the underlying features in the data set. If this quantity is smaller than one, the model is able to improve on a simple featureless estimator. From this point of view, it is clear that the models do not perform as well as could be thought on first sight for the **fb1** data set. At the same time it is also clear that for the **star** data set, the models only barely improve on the naive estimate.

To see the influence of the training-calibration split on the resulting prediction intervals, two smaller experiments were performed where the training-calibration ratio was modified. In the first experiment the split ratio was changed from 50/50 to 75/25, i.e. more data was reserved for the training step. The average coverage was not significantly changed. However, the average width of the intervals decreased on average by 7%. Using more data to train the underlying models, thereby obtaining better predictions, will lead to tighter prediction intervals as long as the calibration set is not too small. This conclusion is in line with the observations from Figs. 1 and 2. This experiment was then repeated in an extreme fashion, where the models were trained on the full data set. Due to the lack of an independent calibration set, the post-hoc calibration was also performed on the same training set. This way the influence of violating the assumptions of the ICP Algorithm 4 and the associated validity theorem can also be investigated. It was found that for some models the coverage decreased sharply. Both of these observations should not come as a surprise. The amount by which the intervals are scaled can be interpreted as a hyperparameter of the model. In general it is better to use more data to train than to validate, as long as the validation data set is representative of the true population. Moreover, optimizing hyperparameters on the training set is known to lead to overfitting, which in this case corresponds to overly optimistic prediction intervals. The CP algorithm looks at how large the errors are on the calibration set, so as to be able to correct for them in the future. However, by using the training set to calibrate, future errors are underestimated and, therefore, the CP algorithm cannot fully correct for them.

6 Conclusions and future perspectives

In this study several types of prediction interval estimators for regression problems were reviewed and compared. Two main properties were taken into account: the coverage degree and the average width of the prediction intervals. It was found that without post-hoc calibration the methods derived from a probabilistic model attained the best coverage degree. However, after calibration, all methods attained the desired coverage and in certain cases the calibrated models even produced intervals with a smaller average width. It was also observed that the predictive power of the model and the quality of the prediction intervals are correlated. For the post-hoc calibration step the conformal prediction framework was used. To obtain the desired results, this method requires the data set to be split in a training and a calibration set. A small experiment confirmed that violating this condition leads to strongly uncalibrated intervals. Although this can be expected from an overfitting point of view, this also shows that the assumption does not merely serve a theoretical purpose. On the other hand, increasing the relative size of the training set can result in smaller prediction intervals without having a negative influence on the calibration.

Although a variety of methods was considered, it is not feasible to include all of them. The most important omission is a more detailed overview of Bayesian neural networks (although one can argue, as was done in the section on dropout networks, that some common neural networks are, at least partially, Bayesian by nature). The main reason for this omission is the large number of choices in terms of priors and approximations, both of which strongly depend on the problem at hand. On the level of calibration there are also some methods that were not included in this paper, mostly because they were either too specific or too complex for simple regression problems. For general regression models the literature on calibration methods is not as extensive as it is for classification models. Recently some advances were made in which β -calibration [48] was generalized to regression problems using a Gaussian process approach [81]. However, as mentioned before, a Gaussian process does not have a favorable scaling behaviour and also in this case certain approximations are necessary. Another technique that was recently introduced [47] calibrates the cumulative distribution function produced by a distribution predictor using isotonic regression. Although the technique itself is simple in spirit, it is only applicable to predictors that construct the full output distribution. In a similar vein [84], at the time of writing still under review, takes a distribution predictor and modifies the loss function such that the quantiles are calibrated without post-hoc calibration. The main benefit of this method is that it does not require a separate calibration set in stark contrast to conformal prediction, but it still requires the construction of the cumulative distribution. By dividing the target space in a finite number of bins, Keren et al. introduced an approach where the regression problem is approximated by a classification problem such that the usual tools for classifier calibration can be applied [41]. The main downside of this approach is that one loses the continuous nature of the initial problem. Another concept that was not covered is that of *predictive distributions* [75,77], where not only a single interval is considered, but a full distribution is estimated. This approach was combined with conformal prediction in [90] giving rise to *Conformal Predictive Systems*.

The choice of data sets in this comparative study was very broad and no specific properties were taken into account a priori. After comparing the results of the different models, it did become apparent that certain assumptions or properties can have a major influence on the performance of the models. The main examples were the normality assumption for mean-variance estimators or the proper scoring rule (20). Models using this scoring rule appeared to behave very badly when used for strongly skewed data sets. Since such data sets are gaining importance in the digital age, it would be interesting to both study methods tailored to these properties and how existing models behave on outliers. Another type of data that was not specifically considered in this study, but that is also becoming increasingly important, is time series data. The main issue with this kind of data, as was already mentioned in Section 4.2, is the autocorrelation, which (possibly) invalidates methods or theorems that make use of i.i.d. or exchangeability properties. This has an influence on both the construction of models and the structure of validation sets [11].

A further aspect that was not considered in this study is the conditional behaviour of the models. When constructing a model that optimizes the coverage probability (1), only the *marginal* coverage is controlled, i.e. the specific properties of an instance are not taken into account. In certain cases it might be relevant to not only attain global validity, but also guarantee the validity on a certain subset of the instance space. A general notion of conditional validity was considered in [51] and [86]. There it was also shown how to incorporate this notion in the conformal prediction framework.

References

1. Achilles, C., Bain, H.P., Bellott, F., Boyd-Zaharias, J., Finn, J., Folger, J., Johnston, J., Word, E.: Tennessee’s Student Teacher Achievement Ratio (STAR) project (2008). DOI 10.7910/DVN/SIWH9F
2. Barber, R.F., Candès, E.J., Ramdas, A., Tibshirani, R.J.: Predictive inference with the jackknife+. *The Annals of Statistics* **49**(1), 486–507 (2021)
3. Blei, D.M., Kucukelbir, A., McAuliffe, J.D.: Variational inference: A review for statisticians. *Journal of the American statistical Association* **112**(518), 859–877 (2017)
4. Blundell, C., Cornebise, J., Kavukcuoglu, K., Wierstra, D.: Weight uncertainty in neural network. In: *International Conference on Machine Learning*, pp. 1613–1622. PMLR (2015)
5. Boström, H., Linusson, H., Löfström, T., Johansson, U.: Accelerating difficulty estimation for conformal regression forests. *Annals of Mathematics and Artificial Intelligence* **81**(1), 125–144 (2017)
6. Breiman, L.: Bagging predictors. *Machine Learning* **24**(2), 123–140 (1996)
7. Breiman, L.: Random forests. *Machine Learning* **45**(1), 5–32 (2001)
8. Buza, K.: Feedback prediction for blogs. In: M. Spiliopoulou, L. Schmidt-Thieme, R. Janzing (eds.) *Data Analysis, Machine Learning and Knowledge Discovery*, pp. 145–152. Springer International Publishing, Cham (2014)
9. Cartagena, O., Parra, S., Muñoz-Carpintero, D., Marín, L.G., Sáez, D.: Review on fuzzy and neural prediction interval modelling for nonlinear dynamical systems. *IEEE Access* **9**, 23357–23384 (2021). DOI 10.1109/ACCESS.2021.3056003
10. Chen, Z., Huang, X.: End-to-end learning for lane keeping of self-driving cars. In: *2017 IEEE Intelligent Vehicles Symposium (IV)*, pp. 1856–1860. IEEE (2017)
11. Chernozhukov, V., Wüthrich, K., Yinchu, Z.: Exact and robust conformal inference methods for predictive machine learning with dependent data. In: *Conference On Learning Theory*, pp. 732–749. PMLR (2018)
12. Coraddu, A., Oneto, L., Ghio, A., Savio, S., Anguita, D., Figari, M.: Machine learning approaches for improving condition-based maintenance of naval propulsion plants. *Proceedings of the Institution of Mechanical Engineers, Part M: Journal of Engineering for the Maritime Environment* **230**(1), 136–153 (2016)

13. Corke, P.I.: A robotics toolbox for MATLAB. *IEEE Robotics & Automation Magazine* **3**(1), 24–32 (1996)
14. Dua, D., Graff, C.: UCI machine learning repository. <http://archive.ics.uci.edu/ml> (2017). Data set repository
15. Efron, B.: Jackknife-after-bootstrap standard errors and influence functions. *Journal of the Royal Statistical Society: Series B (Methodological)* **54**(1), 83–111 (1992)
16. Fang, Y., Xu, P., Yang, J., Qin, Y.: A quantile regression forest based method to predict drug response and assess prediction reliability. *PLOS ONE* **13**(10) (2018)
17. Faulkenberry, D.G.: A method of obtaining prediction intervals. *Journal of the American Statistical Association* **68**(342), 433–435 (1973)
18. Ferreira, R.P., Affonso, C., Sassi, R.J.: Combination of artificial intelligence techniques for prediction the behavior of urban vehicular traffic in the city of sao paulo. In 10th Brazilian Congress on Computational Intelligence (CBIC) pp. 1–7 (2011)
19. Fink, D.: A compendium of conjugate priors. Tech. rep., Montana State Univeristy (1997)
20. Fort, S., Hu, H., Lakshminarayanan, B.: Deep ensembles: A loss landscape perspective (2019). URL [url{https://arxiv.org/abs/1912.02757}](https://arxiv.org/abs/1912.02757)
21. Fraser, D.A., Guttman, I., et al.: Tolerance regions. *Annals of Mathematical Statistics* **27**(1), 162–179 (1956)
22. Gal, Y.: Uncertainty in deep learning. *University of Cambridge* **1**(3) (2016)
23. Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: M.F. Balcan, K.Q. Weinberger (eds.) *Proceedings of The 33rd International Conference on Machine Learning, Proceedings of Machine Learning Research*, vol. 48, pp. 1050–1059. PMLR, New York, New York, USA (2016)
24. Gammernan, A., Vovk, V., Vapnik, V.: Learning by transduction. In: *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence, UAI’98*, pp. 148—155. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1998)
25. Gardner, J.R., Pleiss, G., Bindel, D., Weinberger, K.Q., Wilson, A.G.: Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. In: *Advances in Neural Information Processing Systems* (2018)
26. Geng, C., Huang, S.J., Chen, S.: Recent advances in open set recognition: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* pp. 1–1 (2020). DOI 10.1109/TPAMI.2020.2981604
27. Gentle, J.E.: Monte Carlo methods for statistical inference. In: *Computational Statistics*, pp. 417–433. Springer (2009)
28. Gneiting, T., Raftery, A.E.: Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association* **102**(477), 359–378 (2007)
29. Goan, E., Fookes, C.: Bayesian neural networks: An introduction and survey. In: *Case Studies in Applied Bayesian Data Science*, pp. 45–87. Springer (2020)
30. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: Y. Bengio, Y. LeCun (eds.) *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings* (2015)
31. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: *International Conference on Machine Learning*, pp. 1321–1330. PMLR (2017)
32. Havlicek, O.: GitHub Issue: Very slow QRF prediction. <https://github.com/scikit-garden/scikit-garden/issues/74>
33. Hein, M., Andriushchenko, M., Bitterwolf, J.: Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 41–50 (2019)
34. Hensman, J., Matthews, A., Ghahramani, Z.: Scalable variational gaussian process classification. In: *Artificial Intelligence and Statistics*, pp. 351–360. PMLR (2015)
35. Heskes, T.: Practical confidence and prediction intervals. In: *Proceedings of the 9th International Conference on Neural Information Processing Systems, NIPS’96*, pp. 176—182. MIT Press (1996)
36. Hinton, G.E., van Camp, D.: Keeping the neural networks simple by minimizing the description length of the weights. In: *Proceedings of the Sixth Annual Conference on Computational Learning Theory, COLT ’93*, pp. 5–13. Association for Computing Machinery, New York, NY, USA (1993). DOI 10.1145/168304.168306
37. Jiang, X., Osl, M., Kim, J., Ohno-Machado, L.: Calibrating predictive model estimates to support personalized medicine. *Journal of the American Medical Informatics Association* **19**(2), 263–274 (2012)

38. Johansson, U., Boström, H., Löfström, T., Linusson, H.: Regression conformal prediction with random forests. *Machine Learning* **97**(1-2), 155–176 (2014)
39. Kaya, H., Tüfekci, P., Gürgen, F.S.: Local and global learning methods for predicting power of a combined gas & steam turbine. In: *Proceedings of the International Conference on Emerging Trends in Computer and Electronics Engineering (ICETCEE)*, pp. 13–18 (2012)
40. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? In: I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (eds.) *Advances in Neural Information Processing Systems*, vol. 30, pp. 5580–5590. Curran Associates, Inc. (2017)
41. Keren, G., Cummins, N., Schuller, B.: Calibrated prediction intervals for neural network regressors. *IEEE Access* **6**, 54033–54041 (2018)
42. Khosravi, A., Nahavandi, S., Creighton, D., Atiya, A.F.: Comprehensive review of neural network-based prediction intervals and new advances. *IEEE Transactions on Neural Networks* **22**(9), 1341–1356 (2011). DOI 10.1109/TNN.2011.2162110
43. Khosravi, A., Nahavandi, S., Creighton, D., Atiya, A.F.: Lower upper bound estimation method for construction of neural network-based prediction intervals. *IEEE Transactions on Neural Networks* **22**(3), 337–346 (2011). DOI 10.1109/TNN.2010.2096824
44. Khosravi, A., Nahavandi, S., Srinivasan, D., Khosravi, R.: Constructing optimal prediction intervals by using neural networks and bootstrap method. *IEEE Transactions on Neural Networks and Learning Systems* **26**(8), 1810–1815 (2015). DOI 10.1109/TNNLS.2014.2354418
45. Kivaranovic, D., Johnson, K.D., Leeb, H.: Adaptive, distribution-free prediction intervals for deep networks. In: *International Conference on Artificial Intelligence and Statistics*, pp. 4346–4356. PMLR (2020)
46. Koenker, R., Hallock, K.F.: Quantile regression. *Journal of Economic Perspectives* **15**(4), 143–156 (2001). DOI 10.1257/jep.15.4.143
47. Kuleshov, V., Fenner, N., Ermon, S.: Accurate uncertainties for deep learning using calibrated regression. In: J. Dy, A. Krause (eds.) *Proceedings of the 35th International Conference on Machine Learning, Proceedings of Machine Learning Research*, vol. 80, pp. 2796–2804. PMLR, Stockholmsmässan, Stockholm Sweden (2018)
48. Kull, M., Filho, T.S., Flach, P.: Beta calibration: a well-founded and easily implemented improvement on logistic calibration for binary classifiers. In: A. Singh, J. Zhu (eds.) *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, Proceedings of Machine Learning Research*, vol. 54, pp. 623–631. PMLR, Fort Lauderdale, FL, USA (2017)
49. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, pp. 6405–6416. Curran Associates Inc., Red Hook, NY, USA (2017)
50. Lee, S., Purushwalkam, S., Cogswell, M., Crandall, D.J., Batra, D.: Why M heads are better than one: Training a diverse ensemble of deep networks. *CoRR* **abs/1511.06314** (2015). URL <http://arxiv.org/abs/1511.06314>
51. Lei, J., Wasserman, L.: Distribution-free prediction bands for non-parametric regression. *Journal of the Royal Statistical Society: Series B: Statistical Methodology* pp. 71–96 (2014)
52. Liu, J., Paisley, J., Kioumourtoglou, M.A., Coull, B.: Accurate uncertainty estimation and decomposition in ensemble learning. In: *Advances in Neural Information Processing Systems*, vol. 32. Curran Associates, Inc. (2019)
53. MacKay, D.J.: A practical bayesian framework for backpropagation networks. *Neural Computation* **4**(3), 448–472 (1992)
54. Malinin, A., Gales, M.: Predictive uncertainty estimation via prior networks. In: *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, pp. 7047–7058. Curran Associates Inc. (2018)
55. Meinshausen, N.: Quantile regression forests. *Journal of Machine Learning Research* **7**(Jun), 983–999 (2006)
56. Michelmore, R., Wicker, M., Laurenti, L., Cardelli, L., Gal, Y., Kwiatkowska, M.: Uncertainty quantification with statistical guarantees in end-to-end autonomous driving control. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 7344–7350 (2020). DOI 10.1109/ICRA40945.2020.9196844
57. Miller, R.G.: The jackknife-a review. *Biometrika* **61**(1), 1–15 (1974)

58. Naeini, M.P., Cooper, G., Hauskrecht, M.: Obtaining well calibrated probabilities using bayesian binning. In: Proceedings of the AAAI Conference on Artificial Intelligence, *AAAI'15*, vol. 29, pp. 2901–2907. AAAI Press (2015)
59. Neal, R.M.: Bayesian Learning for Neural Networks. Springer-Verlag, Berlin, Heidelberg (1996)
60. Nix, D.A., Weigend, A.S.: Estimating the mean and variance of the target probability distribution. In: Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94), vol. 1, pp. 55–60. IEEE (1994)
61. Ott, M., Auli, M., Grangier, D., Ranzato, M.: Analyzing uncertainty in neural machine translation. In: International Conference on Machine Learning (ICML), pp. 3956–3965. PMLR (2018)
62. Papadopoulos, H., Gammernan, A., Vovk, V.: Normalized nonconformity measures for regression conformal prediction. In: Proceedings of the IASTED International Conference on Artificial Intelligence and Applications (AIA 2008), pp. 64–69 (2008)
63. Papadopoulos, H., Proedrou, K., Vovk, V., Gammernan, A.: Inductive confidence machines for regression. In: European Conference on Machine Learning, pp. 345–356. Springer (2002)
64. Papadopoulos, H., Vovk, V., Gammernan, A.: Conformal prediction with neural networks. In: 19th IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2007), vol. 2, pp. 388–395 (2007)
65. Papadopoulos, H., Vovk, V., Gammernan, A.: Regression conformal prediction with nearest neighbours. *Journal of Artificial Intelligence Research* **40**, 815–840 (2011)
66. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (eds.) *Advances in Neural Information Processing Systems* 32, pp. 8024–8035. Curran Associates, Inc. (2019)
67. Pearce, T.: Uncertainty in neural networks; bayesian ensembles, priors & prediction intervals. Ph.D. Thesis, University of Cambridge (2020)
68. Pearce, T., Brintrup, A., Zaki, M., Neely, A.: High-quality prediction intervals for deep learning: A distribution-free, ensembled approach. In: International Conference on Machine Learning, pp. 4075–4084. PMLR (2018)
69. 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.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)
70. Raffei, M.H., Adeli, H.: A novel machine learning model for estimation of sale prices of real estate units. *Journal of Construction Engineering and Management* **142**(2) (2016)
71. Redmond, M., Baveja, A.: A data-driven software tool for enabling cooperative information sharing among police departments. *European Journal of Operational Research* **141**(3), 660–678 (2002)
72. Reynolds, D.A.: Gaussian mixture models. *Encyclopedia of Biometrics* **741**, 659–663 (2009)
73. Romano, Y., Patterson, E., Candes, E.: Conformalized quantile regression. In: *Advances in Neural Information Processing Systems*, pp. 3543–3553 (2019)
74. Saunders, C., Gammernan, A., Vovk, V.: Transduction with confidence and credibility. In: *Proceedings of the International Joint Conference on Artificial Intelligence*, pp. 722–726 (1999)
75. Schweder, T., Hjort, N.L.: Confidence, likelihood, probability, vol. 41. Cambridge University Press (2016)
76. Sesia, M., Candès, E.J.: A comparison of some conformal quantile regression methods. *Stat* **9**(1) (2020)
77. Shen, J., Liu, R.Y., Xie, M.g.: Prediction with confidence — A general framework for predictive inference. *Journal of Statistical Planning and Inference* **195**, 126–140 (2018)
78. Singh, K., Sandhu, R.K., Kumar, D.: Comment volume prediction using neural networks and decision trees. In: *IEEE UKSim-AMSS 17th International Conference on Computer Modelling and Simulation (UKSim2015)* (2015)
79. Singh, S.P., Kumar, A., Darbari, H., Singh, L., Rastogi, A., Jain, S.: Machine translation using deep learning: An overview. In: *2017 International Conference on Computer, Communications and Electronics (Comptelix)*, pp. 162–167. IEEE (2017)

80. Sollich, P., Krogh, A.: Learning with ensembles: How over-fitting can be useful. In: Proceedings of the 8th International Conference on Neural Information Processing Systems, NIPS'95, pp. 190–196. MIT Press, Cambridge, MA, USA (1995)
81. Song, H., Diethe, T., Kull, M., Flach, P.: Distribution calibration for regression. In: K. Chaudhuri, R. Salakhutdinov (eds.) Proceedings of the 36th International Conference on Machine Learning, *Proceedings of Machine Learning Research*, vol. 97, pp. 5897–5906. PMLR, Long Beach, California, USA (2019)
82. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* **15**(56), 1929–1958 (2014)
83. Teye, M., Azizpour, H., Smith, K.: Bayesian uncertainty estimation for batch normalized deep networks. In: International Conference on Machine Learning, pp. 4907–4916. PMLR (2018)
84. Utpala, S., Rai, P.: Quantile regularization: Towards implicit calibration of regression models. arXiv preprint arXiv:2002.12860 (2020)
85. Van Amersfoort, J., Smith, L., Teh, Y.W., Gal, Y.: Uncertainty estimation using a single deep deterministic neural network. In: International Conference on Machine Learning, pp. 9690–9700. PMLR (2020)
86. Vovk, V.: Conditional validity of inductive conformal predictors. In: S.C.H. Hoi, W. Buntine (eds.) Proceedings of the Asian Conference on Machine Learning, *Proceedings of Machine Learning Research*, vol. 25, pp. 475–490. PMLR, Singapore Management University, Singapore (2012)
87. Vovk, V.: Cross-conformal predictors. *Annals of Mathematics and Artificial Intelligence* **74**(1), 9–28 (2015)
88. Vovk, V., Gammerman, A., Saunders, C.: Machine-learning applications of algorithmic randomness. In: Proceedings of the Sixteenth International Conference on Machine Learning, ICML '99, pp. 444–453. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1999)
89. Vovk, V., Gammerman, A., Shafer, G.: Algorithmic Learning in a Random World. Springer Science & Business Media (2005)
90. Vovk, V., Shen, J., Manokhin, V., Xie, M.g.: Nonparametric predictive distributions based on conformal prediction. In: Conformal and Probabilistic Prediction and Applications, pp. 82–102. PMLR (2017)
91. Wager, S., Hastie, T., Efron, B.: Confidence intervals for random forests: The jackknife and the infinitesimal jackknife. *The Journal of Machine Learning Research* **15**(1), 1625–1651 (2014)
92. Williams, C., Rasmussen, C.: Gaussian processes for regression. In: Advances in Neural Information Processing Systems, pp. 514–520 (1996)
93. Wilson, A., Nickisch, H.: Kernel interpolation for scalable structured gaussian processes (KISS-GP). In: International Conference on Machine Learning, pp. 1775–1784. PMLR (2015)
94. Yeh, I.C.: Modeling of strength of high-performance concrete using artificial neural networks. *Cement and Concrete Research* **28**(12), 1797–1808 (1998). DOI [https://doi.org/10.1016/S0008-8846\(98\)00165-3](https://doi.org/10.1016/S0008-8846(98)00165-3)
95. Zhang, H., Zimmerman, J., Nettleton, D., Nordman, D.J.: Random forest prediction intervals. *The American Statistician* **74**(4), 392–406 (2020)

A Experimental results

In this section all experimental results are summarized. For every metric (R^2 -coefficient, coverage and average width) the results are split into two tables, with five data sets each, for clarity. The rows represent the various models and the columns correspond to the different data sets. For every combination the average over 50 random train-test splits is reported. The standard deviation is shown between parentheses. As mentioned in Section 5, there are no results available for exact Gaussian processes on the data sets **fb1** and **blog** due to exceeding the runtime limit; this is indicated by “OoT” (out of time). Moreover, certain values

were deemed out-of-range, e.g. R^2 -coefficients below -1 or cases where the standard deviation was unreasonably large; this is indicated by “OoR” (out of range).

	concrete	naval	turbine	puma32H	residential
NN-CP	0.833 (0.027)	0.987 (0.004)	0.942 (0.004)	0.879 (0.011)	0.907 (0.073)
RF-CP	0.872 (0.021)	0.990 (0.002)	0.955 (0.004)	0.930 (0.003)	0.923 (0.033)
QR	0.865 (0.035)	0.992 (0.002)	0.944 (0.004)	0.936 (0.006)	0.962 (0.018)
QR-CP	0.817 (0.038)	0.986 (0.004)	0.942 (0.004)	0.889 (0.014)	0.919 (0.062)
DE	0.870 (0.031)	0.813 (0.009)	0.947 (0.004)	0.929 (0.007)	0.946 (0.068)
DE-CP	0.809 (0.047)	0.740 (0.010)	0.945 (0.004)	0.697 (0.085)	0.899 (0.081)
Drop	0.875 (0.023)	0.994 (0.003)	0.945 (0.004)	0.931 (0.005)	0.957 (0.018)
Drop-CP	0.813 (0.042)	0.990 (0.003)	0.943 (0.004)	0.876 (0.010)	0.905 (0.029)
MVE	0.874 (0.020)	0.994 (0.002)	0.945 (0.004)	0.935 (0.004)	0.967 (0.016)
MVE-CP	0.820 (0.031)	0.990 (0.003)	0.943 (0.004)	0.883 (0.010)	0.931 (0.036)
GP	0.884 (0.022)	1.000 (0.000)	0.952 (0.004)	0.132 (0.169)	0.903 (0.073)
GP-CP	0.844 (0.025)	1.000 (0.000)	0.945 (0.004)	0.187 (0.107)	0.866 (0.087)
SVGP	0.864 (0.021)	0.995 (0.008)	0.943 (0.004)	0.216 (0.021)	0.659 (0.094)
SVGP-CP	0.813 (0.027)	0.995 (0.003)	0.941 (0.004)	0.211 (0.020)	0.254 (0.062)

Table 3: Summary of R^2 -coefficients (part I). For the data sets in each column, the results are shown for all models (standard deviation between parentheses). For clarity, models that come in both a default and a conformalized version have been grouped together.

	crime2	fb1	blog	traffic	star
NN-CP	0.609 (0.063)	0.424 (0.067)	0.010 (1.390)	0.151 (0.755)	0.116 (0.050)
RF-CP	0.629 (0.032)	0.616 (0.052)	0.554 (0.050)	0.549 (0.164)	0.160 (0.036)
QR	0.638 (0.037)	0.540 (0.054)	0.462 (0.054)	0.485 (0.192)	0.142 (0.044)
QR-CP	0.627 (0.046)	0.495 (0.048)	0.430 (0.052)	0.277 (0.560)	0.117 (0.045)
DE	0.660 (0.031)	0.371 (0.047)	0.289 (0.034)	0.500 (0.173)	0.179 (0.033)
DE-CP	0.642 (0.032)	0.353 (0.048)	0.287 (0.031)	0.340 (0.295)	0.148 (0.035)
Drop	0.656 (0.033)	0.471 (0.097)	0.463 (0.061)	0.441 (0.191)	0.177 (0.036)
Drop-CP	0.635 (0.038)	0.425 (0.064)	OoR	0.315 (0.318)	0.131 (0.045)
MVE	0.658 (0.036)	0.411 (0.065)	OoR	0.523 (0.164)	0.173 (0.032)
MVE-CP	0.641 (0.039)	0.379 (0.074)	OoR	0.314 (0.366)	0.143 (0.041)
GP	0.562 (0.122)	OoT	OoT	0.316 (0.197)	0.205 (0.026)
GP-CP	0.460 (0.206)	OoT	OoT	0.244 (0.191)	0.169 (0.045)
SVGP	0.654 (0.031)	0.385 (0.042)	0.421 (0.035)	0.182 (0.151)	0.184 (0.025)
SVGP-CP	0.258 (0.209)	0.360 (0.038)	0.393 (0.034)	0.096 (0.141)	0.168 (0.026)

Table 4: Summary of R^2 -coefficients (part II). For the data sets in each column, the results are shown for all models (standard deviation between parentheses). For clarity, models that come in both a default and a conformalized version have been grouped together. “OoT” stands for “out of time”, for these combinations of data sets and models, the maximum runtime was exceeded. “OoR” stands for “out of range”, these values exceeded reasonable ranges for the R^2 -coefficient.

	concrete	naval	turbine	puma32H	residential
NN-CP	0.901 (0.025)	0.899 (0.006)	0.900 (0.010)	0.900 (0.011)	0.902 (0.048)
RF-CP	0.900 (0.028)	0.900 (0.008)	0.901 (0.009)	0.900 (0.009)	0.890 (0.047)
QR	0.914 (0.020)	0.997 (0.003)	0.919 (0.011)	0.908 (0.011)	0.950 (0.029)
QR-CP	0.899 (0.026)	0.900 (0.007)	0.900 (0.010)	0.899 (0.009)	0.904 (0.046)
DE	0.928 (0.019)	0.995 (0.002)	0.940 (0.008)	0.948 (0.008)	0.959 (0.023)
DE-CP	0.893 (0.028)	0.900 (0.007)	0.900 (0.011)	0.902 (0.010)	0.909 (0.040)
Drop	0.580 (0.057)	0.962 (0.047)	0.352 (0.031)	0.671 (0.033)	0.872 (0.053)
Drop-CP	0.897 (0.025)	0.900 (0.007)	0.901 (0.012)	0.900 (0.009)	0.902 (0.043)
MVE	0.942 (0.019)	1.000 (0.001)	0.939 (0.009)	0.951 (0.018)	0.974 (0.023)
MVE-CP	0.901 (0.022)	0.900 (0.007)	0.901 (0.010)	0.901 (0.008)	0.908 (0.040)
GP	0.887 (0.023)	0.782 (0.071)	0.916 (0.007)	0.890 (0.014)	0.928 (0.032)
GP-CP	0.900 (0.026)	0.859 (0.024)	0.902 (0.009)	0.901 (0.009)	0.899 (0.041)
SVGP	0.918 (0.020)	0.902 (0.132)	0.923 (0.009)	0.891 (0.008)	0.967 (0.022)
SVGP-CP	0.900 (0.025)	0.900 (0.008)	0.901 (0.010)	0.901 (0.009)	0.897 (0.043)

Table 5: Summary of coverage degrees (part I). For the data sets in each column, the results are shown for all models (standard deviation between parentheses). For clarity, models that come in both a default and a conformalized version have been grouped together.

	crime2	fb1	blog	traffic	star
NN-CP	0.901 (0.015)	0.901 (0.004)	0.900 (0.003)	0.903 (0.082)	0.906 (0.018)
RF-CP	0.903 (0.017)	0.901 (0.004)	0.901 (0.003)	0.910 (0.078)	0.905 (0.016)
QR	0.857 (0.047)	0.916 (0.041)	0.898 (0.041)	0.807 (0.080)	0.838 (0.039)
QR-CP	0.904 (0.016)	0.902 (0.004)	0.900 (0.004)	0.900 (0.077)	0.904 (0.018)
DE	0.940 (0.016)	0.960 (0.011)	0.959 (0.010)	0.884 (0.097)	0.901 (0.019)
DE-CP	0.908 (0.017)	0.901 (0.004)	0.900 (0.003)	0.898 (0.075)	0.902 (0.016)
Drop	0.546 (0.051)	0.802 (0.062)	0.612 (0.071)	0.467 (0.145)	0.309 (0.050)
Drop-CP	0.904 (0.017)	0.901 (0.004)	0.900 (0.004)	0.913 (0.056)	0.905 (0.016)
MVE	0.925 (0.022)	0.969 (0.012)	0.965 (0.012)	0.865 (0.107)	0.884 (0.029)
MVE-CP	0.909 (0.017)	0.901 (0.004)	0.900 (0.003)	0.903 (0.074)	0.904 (0.017)
GP	0.894 (0.019)	OoT	OoT	0.867 (0.077)	0.899 (0.014)
GP-CP	0.900 (0.020)	OoT	OoT	0.901 (0.066)	0.902 (0.018)
SVGP	0.913 (0.015)	0.979 (0.002)	0.974 (0.002)	0.940 (0.044)	0.906 (0.015)
SVGP-CP	0.903 (0.020)	0.901 (0.004)	0.900 (0.003)	0.892 (0.074)	0.904 (0.017)

Table 6: Summary of coverage degrees (part II). For the data sets in each column, the results are shown for all models (standard deviation between parentheses). For clarity, models that come in both a default and a conformalized version have been grouped together. “OoT” stands for “out of time”, for these combinations of data sets and models, the maximum runtime was exceeded.

	concrete	naval	turbine	puma32H	residential
NN-CP	1.284 (0.090)	0.370 (0.057)	0.754 (0.014)	1.136 (0.052)	0.915 (0.241)
RF-CP	1.139 (0.090)	0.199 (0.010)	0.663 (0.013)	0.873 (0.014)	0.700 (0.148)
QR	1.209 (0.173)	0.535 (0.021)	0.775 (0.018)	0.862 (0.044)	0.693 (0.087)
QR-CP	1.322 (0.179)	0.374 (0.046)	0.731 (0.011)	1.083 (0.075)	0.857 (0.195)
DE	1.188 (0.121)	1.921 (0.061)	0.820 (0.015)	1.026 (0.053)	0.775 (0.562)
DE-CP	1.322 (0.158)	1.505 (0.033)	0.730 (0.015)	1.620 (0.182)	0.838 (0.311)
Drop	0.527 (0.084)	0.351 (0.010)	0.204 (0.017)	0.527 (0.035)	0.493 (0.122)
Drop-CP	1.354 (0.147)	0.334 (0.047)	0.755 (0.012)	1.150 (0.047)	0.893 (0.186)
MVE	1.209 (0.112)	0.586 (0.022)	0.832 (0.033)	1.053 (0.072)	0.700 (0.151)
MVE-CP	1.282 (0.118)	0.346 (0.047)	0.744 (0.020)	1.157 (0.044)	0.714 (0.155)
GP	1.006 (0.060)	0.222 (0.025)	0.757 (0.013)	3.064 (0.265)	0.542 (0.046)
GP-CP	1.234 (0.087)	0.287 (0.049)	0.739 (0.012)	3.050 (0.263)	0.679 (0.090)
SVGP	1.240 (0.038)	0.217 (0.021)	0.800 (0.015)	2.939 (0.047)	2.769 (0.224)
SVGP-CP	1.356 (0.094)	0.210 (0.087)	0.754 (0.013)	3.023 (0.056)	2.256 (0.294)

Table 7: Summary of average widths (part I). For the data sets in each column, the results are shown for all models (standard deviation between parentheses). For clarity, models that come in both a default and a conformalized version have been grouped together.

	crime2	fb1	blog	traffic	star
NN-CP	2.021 (0.130)	0.444 (0.059)	0.583 (0.118)	2.724 (0.711)	3.101 (0.117)
RF-CP	2.009 (0.106)	0.377 (0.020)	0.487 (0.027)	2.314 (0.428)	3.023 (0.090)
QR	1.567 (0.293)	0.394 (0.041)	0.327 (0.027)	2.067 (0.430)	2.648 (0.267)
QR-CP	1.704 (0.153)	0.285 (0.026)	0.271 (0.016)	2.939 (0.729)	3.095 (0.136)
DE	1.847 (0.183)	OoR	OoR	2.229 (0.535)	2.958 (0.150)
DE-CP	1.649 (0.093)	OoR	OoR	2.565 (0.609)	3.027 (0.102)
Drop	0.663 (0.077)	0.310 (0.062)	0.282 (0.058)	0.997 (0.301)	0.746 (0.106)
Drop-CP	1.807 (0.120)	0.405 (0.053)	0.438 (0.081)	2.873 (0.668)	3.100 (0.115)
MVE	1.683 (0.213)	OoR	OoR	2.083 (0.503)	2.870 (0.189)
MVE-CP	1.602 (0.109)	OoR	OoR	2.602 (0.727)	3.052 (0.114)
GP	1.972 (0.266)	OoT	OoT	2.427 (0.343)	2.931 (0.063)
GP-CP	2.251 (0.433)	OoT	OoT	2.845 (0.534)	3.014 (0.111)
SVGP	1.986 (0.045)	2.567 (0.170)	2.556 (0.158)	3.449 (0.168)	3.012 (0.052)
SVGP-CP	2.590 (0.475)	0.787 (0.059)	0.695 (0.061)	3.164 (0.707)	3.012 (0.100)

Table 8: Summary of average widths (part II). For the data sets in each column, the results are shown for all models (standard deviation between parentheses). For clarity, models that come in both a default and a conformalized version have been grouped together. “OoT” stands for “out of time”, for these combinations of data sets and models, the maximum runtime was exceeded. “OoR” stands for “out of range”, these values exceeded reasonable ranges for the average interval width.