COPULA-BASED CONFORMAL PREDICTION FOR MULTI-TARGET REGRESSION

A PREPRINT

Soundouss Messoudi
HEUDIASYC - UMR CNRS 7253
Université de Technologie de Compiègne
60203 COMPIEGNE - FRANCE
soundouss.messoudi@hds.utc.fr

Sébastien Destercke
HEUDIASYC - UMR CNRS 7253
Université de Technologie de Compiègne
60203 COMPIEGNE - FRANCE
sebastien.destercke@hds.utc.fr

Sylvain Rousseau
HEUDIASYC - UMR CNRS 7253
Université de Technologie de Compiègne
60203 COMPIEGNE - FRANCE
sylvain.rousseau@hds.utc.fr

January 29, 2021

ABSTRACT

There are relatively few works dealing with conformal prediction for multi-task learning issues, and this is particularly true for multi-target regression. This paper focuses on the problem of providing valid (i.e., frequency calibrated) multi-variate predictions. To do so, we propose to use copula functions applied to deep neural networks for inductive conformal prediction. We show that the proposed method ensures efficiency and validity for multi-target regression problems on various data sets.

Keywords Inductive conformal prediction · Copula functions · Multi-target regression · Deep neural networks.

1 Introduction

The most common supervised task in machine learning is to learn a single-task, single-output prediction model. However, such a setting can be ill-adapted to some problems and applications.

On the one hand, producing a single output can be undesirable when data is scarce and when producing reliable, possibly set-valued predictions is important (for instance in the medical domain where examples are very hard to collect for specific targets, and where predictions are used for critical decisions). Such an issue can be solved by using conformal prediction approaches [1]. It was initially proposed as a transductive online learning approach to provide set predictions (in the classification case) or interval predictions (in the case of regression) with a statistical guarantee depending on the probability of error tolerated by the user, but was then extended to handle inductive processes [2]. On the other hand, there are many situations where there are multiple, possibly correlated output variables to predict at once, and it is then natural to try to leverage such correlations to improve predictions. Such learning tasks are commonly called Multi-task in the literature [3].

Most research work on conformal prediction for multi-task learning focuses on the problem of multi-label prediction [4, 5], where each task is a binary classification one. Conformal prediction for multi-target regression has been less explored, with only a few studies dealing with it: Kuleshov *et al.* [6] provide a theoretical framework to use conformal predictors within manifold (e.g., to provide a mono-dimensional embedding of the multi-variate output), while Neeven and Smirnov [7] use a straightforward multi-target extension of a conformal single-output k-nearest neighbor regressor [8] to provide weather forecasts. However, this latter essentially verifies validity (i.e., having well-calibrated

outputs) for each individual target. Recently, we proposed a simple method to have an approximate validity for the multi-variate prediction [9], that generally provided overly conservative results.

In this paper, we propose a new conformal prediction method fitted to multi-target regression, that makes use of copulas [10] (a common tool to model dependence between multi-variate random variables) to provide valid multi-variate predictions. The interest of such a framework is that it remains very easy to apply while linking multi-variate conformal predictions to the theoretically sound framework that are copulas. Experiments also show that it works quite well, and allows to improve upon previous heuristics [9].

Section 2 provides a general overview of our problem: a brief introduction to conformal prediction and multi-target regression will be presented in Sections 2.1 and 2.2, before raising the problematic of applying conformal prediction to the multi-target regression setting in Section 2.3. We will then present our setting in Section 3: we will first recall the needed basic principles and theorems of copulas in Section 3.1, before detailing our conformal multi-target approach in Section 3.2. The experiments and their results are described in Section 4.

2 Inductive conformal prediction (ICP) for Multi-Target Regression

This section recalls the basics of inductive conformal regression and multi-target regression, before introducing the issues we will tackle in this paper.

2.1 Inductive conformal regression

In regression tasks, conformal prediction is a method that provides a statistical guarantee to the predictions by giving an interval prediction instead of a point prediction in the regression case. By statistical guarantee, it is meant that the set-valued predictions cover the true value with a given frequency, i.e., they are calibrated. It was first introduced as a transductive online learning approach [11] and then adapted to the inductive framework [2] where one uses a model induced from training examples to get conformal predictions for the new instances. The two desirable features in conformal regressors are (a) *validity*, i.e. the error rate does not exceed ϵ for each chosen confidence level $1 - \epsilon$, and (b) *efficiency*, meaning prediction intervals are as small as possible.

Let $\{z_1 = (x_1, y_1), z_2 = (x_2, y_2), \ldots, z_n = (x_n, y_n)\}$ be the successive pairs of an object $x_i \in X$ and its real-valued label $y_i \in \mathbb{R}$, which constitute the observed examples. Assuming that the underlying random variables are exchangeable (a weaker condition than i.i.d.), we can predict $y_{n+1} \in \mathbb{R}$ for any new object $x_{n+1} \in X$ by following the inductive conformal framework.

The first step consists of splitting the original data set $Z = \{z_1, \ldots, z_n\}$ into a *training set* $Z^{tr} = \{z_1, \ldots, z_l\}$ and a *calibration set* $Z^{cal} = \{z_{l+1}, \ldots, z_n\}$, with $|Z^{cal}| = n - l$. Then, an *underlying algorithm* is trained on Z^{tr} to obtain the *non-conformity measure* A_l, a measure that evaluates the strangeness of an example compared to other examples of a bag, called the non-conformity score. Hence, we can calculate the non-conformity score α_k for an example z_k compared to the other examples in the bag $\{z_1, \ldots, z_l\}$ with $\alpha_k = A_l(\{z_1, \ldots, z_l\}, z_k)$.

By computing the non-conformity score α_i for each example z_i of Z^{cal} using this equation, we get the sequence $\alpha_{l+1}, \ldots, \alpha_n$. When making a prediction for a new example x_{n+1}, we use the underlying algorithm to associate to any possible prediction $\hat{y}$ its non-conformity score $\alpha_{n+1}^{\hat{y}}$, and calculate its *p-value* which indicates the proportion of less conforming examples than z_{n+1}, with:

$$p(\hat{y}_{n+1}) = \frac{|\{i = l + 1, \ldots, n, n + 1 : \alpha_i \geq \alpha_{n+1}^{\hat{y}}\}|}{n - l + 1}. \tag{1}$$

The final step before producing the conformal prediction consists of choosing the *significance level* $\epsilon \in (0, 1)$ to get a prediction set with a *confidence level* of $1 - \epsilon$, which is the statistical guarantee of coverage of the true value y_{n+1} by the interval prediction $\hat{\mathbf{y}}_{n+1}$ such that

$$\hat{\mathbf{y}}_{n+1} = \{\hat{y}_{n+1} \in \mathbb{R} : p(\hat{y}_{n+1}) > \epsilon\}.$$

The most basic non-conformity measure in a regression setting is the absolute difference between the actual value y_i and the predicted value $\hat{y}_i$ by the underlying algorithm. The non-conformity score is then calculated as follows:

$$\alpha_i = |y_i - \hat{y}_i|. \tag{2}$$

The sequence of non-conformity scores $\alpha_{l+1}, \ldots, \alpha_n$ for all examples in Z^{cal} are obtained and sorted in descending order. Then, we compute the index of the $(1 - \epsilon)$-percentile non-conformity score α_s, based on the chosen significance

level ϵ, such as:

$$\mathbb{P}(|y_i - \hat{y}_i| \leq \alpha_s) \geq 1 - \epsilon. \tag{3}$$

Finally, the prediction interval for each new example x_{n+1}, which covers the true output y_{n+1} with probability $1 - \epsilon$ is calculated as:

$$\hat{\mathbf{y}}_{n+1} = [\hat{y}_{n+1} - \alpha_s, \hat{y}_{n+1} + \alpha_s]. \tag{4}$$

The drawback of this standard non-conformity measure is that all prediction intervals are equally sized ($2\alpha_s$) for a given confidence level. Adopting a *normalized* non-conformity measure instead provides personalized individual bounds for each new example by scaling the standard non-conformity measure with σ_i, a term that estimates the difficulty of predicting y_i. This means that using a *normalized* non-conformity measure gives a smaller prediction interval for "easy" examples, and a bigger one for "hard" examples. Thus, two distinct examples with the same α_s calculated by (2) will have two different interval predictions depending on their difficulty. In this case, the normalized non-conformity score is as follows:

$$\alpha_i = \frac{|y_i - \hat{y}_i|}{\sigma_i}. \tag{5}$$

Thus, we have:

$$\mathbb{P}\left(\frac{|y_i - \hat{y}_i|}{\sigma_i} \leq \alpha_s\right) \geq 1 - \epsilon, \tag{6}$$

which becomes an equality if the method is perfectly calibrated. For a new example x_{n+1}, the prediction interval becomes :

$$\hat{\mathbf{y}}_{n+1} = [\hat{y}_{n+1} - \alpha_s \sigma_{n+1}, \hat{y}_{n+1} + \alpha_s \sigma_{n+1}]. \tag{7}$$

The value σ_i can be defined in various ways. A popular approach proposed by Papadopoulos and Haralambous [12] consists of training a small neural network to estimate the error of the underlying algorithm by predicting the value $\mu_i = \ln(|y_i - \hat{y}_i|)$. In this case, the non-conformity score is defined as:

$$\alpha_i = \frac{|y_i - \hat{y}_i|}{\exp(\mu_i) + \beta}, \tag{8}$$

where $\beta \geq 0$ is a sensitivity parameter. With the significance level ϵ, we have:

$$\mathbb{P}\left(\frac{|y_i - \hat{y}_i|}{\exp(\mu_i) + \beta} \leq \alpha_s\right) \geq 1 - \epsilon. \tag{9}$$

For a new example x_{n+1}, the prediction interval is:

$$\hat{\mathbf{y}}_{n+1} = [\hat{y}_{n+1} - \alpha_s(\exp(\mu_{n+1}) + \beta), \hat{y}_{n+1} + \alpha_s(\exp(\mu_{n+1}) + \beta)]. \tag{10}$$

Other approaches use different algorithms to normalize the non-conformity scores, such as regression trees [13] and k-nearest neighbors [8]. Before introducing the problem of multi-target regression, let us first note that, assuming that our method is well-calibrated and that $|y_i - \hat{y}_i|/\sigma_i$ is associated to a random variable Q, (6) can be rewritten as

$$\mathbb{P}(Q \leq \alpha_s) = 1 - \epsilon := F_Q(\alpha_s), \tag{11}$$

which will be instrumental when dealing with copulas and multi-variate outputs later on. Also note that this means that specifying a confidence ϵ uniquely defines a value α_s.

2.2 Multi-target regression (MTR)

In multi-target regression, the feature space X is the same as in standard regression, but the target space $Y \subset \mathbb{R}^m$ is made of m real-valued targets. This means that observations are i.i.d pairs (x_i, y_i) drawn from a probability distribution on $X \times Y$, where each instance $x_i \in X$ is associated to an m dimensional real-valued target $y_i = (y_i^1, \ldots, y_i^m) \in Y$. The usual objective of multi-target regression is then to learn a predictor $h : X \rightarrow Y$, i.e. to predict multiple outputs based on the input features characterizing the data set, which generalizes standard regression. There are two distinct approaches to treat MTR called *algorithm adaptation* and *problem transformation* methods.

For *algorithm adaptation* approaches, standard single-output regression algorithms are extended to the multi-target regression problem. Many models were adapted to the MTR problem, such as Support Vector Regressors [14], regression trees [15], kernel methods [16] and rule ensembles [17].

In *problem transformation*, one usually decomposes the initial multi-variate problems into several simpler problems, thus allowing the use of standard classification methods without the need for an adaptation that can be tricky or

computationally costly. A prototypical example of such a transformation is the chaining method [18], where one predicts each target sequentially, using the output and predictions of previous targets as inputs for the next one, thus capturing some correlations between the targets.

As our goal here is not to produce a new MTR method, but rather to propose a flexible means to make their predictions reliable through conformal prediction, we will not make a more detailed review of those methods. The reader interested in different methods can consult for instance [18]. We will now detail how conformal prediction and MTR can be combined. Let us just mention that exploiting the possible relationships allow in general to improve performances of the methods [19, 20].

2.3 Inductive conformal prediction for Multi-Target Regression

As said before, previous studies about conformal MTR focused on providing valid and efficient inferences target-wise [7], thus potentially neglecting the potential advantages of exploiting target relations. Our main goal in this paper is to provide an easy conformal MTR method allowing to do so.

Within the MTR setting, we have a multi-dimensional output $\{Y^1, \ldots, Y^m\}$ (we will use superscripts to denote the dimensions, and subscripts to denote sample indices) with $Y^j \in \mathbb{R}, j \in \{1, \ldots, m\}$ the different individual real-valued m targets. Let $\underline{\hat{y}}^j_{n+1}, \overline{\hat{y}}^j_{n+1}$ be respectively the lower and upper bounds of the interval predictions given by the non-conformity measure for each target Y^j given a new instance x_{n+1}. We define the hyper-rectangle $[\hat{\mathbf{y}}_{n+1}]$ as the following Cartesian product:

$$[\hat{\mathbf{y}}_{n+1}] = \times_{j=0}^m [\underline{\hat{y}}^j_{n+1}, \overline{\hat{y}}^j_{n+1}]. \tag{12}$$

This hyper-rectangle forms the volume $\prod_{j=0}^m (\overline{\hat{y}}^j_{n+1} - \underline{\hat{y}}^j_{n+1})$ to which a global prediction y_{n+1} of a new example x_{n+1} should belong in order to be valid, i.e. each single prediction y^j_{n+1} for each individual target Y^j should be between the bounds $\underline{\hat{y}}^j_{n+1}, \overline{\hat{y}}^j_{n+1}$ of its interval prediction. With this view, the objective of the conformal prediction framework for MTR in the normalized setting is to satisfy a global significance level ϵ_g required by the user such that:

$$\mathbb{P}(y_{n+1} \in [\hat{\mathbf{y}}_{n+1}]) \geq 1 - \epsilon_g. \tag{13}$$

This probability can also be written as follows:

$$\mathbb{P}(y^1_{n+1} \in [\underline{y^1_{n+1}}, \overline{y^1_{n+1}}], \ldots, y^m_{n+1} \in [\underline{y^m_{n+1}}, \overline{y^m_{n+1}}])$$
$$= \mathbb{P}\left(\frac{|y^1_{n+1} - \hat{y}^1_{n+1}|}{\sigma^1_{n+1}} \leq \alpha^1_s, \ldots, \frac{|y^m_{n+1} - \hat{y}^m_{n+1}|}{\sigma^m_{n+1}} \leq \alpha^m_s \right) \geq 1 - \epsilon_g. \tag{14}$$

Thus, we need to find the individual non-conformity scores $\alpha^1_s, \ldots, \alpha^m_s$, defined for instance by target-wise confidence levels ϵ_j, such that we ensure a global confidence level $1 - \epsilon_g$. Extending (11) and considering the random variables $Q^j = |y^j - \hat{y}^j|/\sigma^j, j \in \{1, \ldots, m\}$, we get:

$$\mathbb{P}(Q^1 \leq \alpha^1_s, \ldots, Q^m \leq \alpha^m_s) \geq 1 - \epsilon_g. \tag{15}$$

Should we know the joint distribution in (15), and therefore the dependence relations between target predictions, it would be relatively easy to get the individual significance levels[1] ϵ_j associated to the individual non-conformity scores α^j_s such that we satisfy the chosen confidence level $1 - \epsilon_g$. Yet, such a joint distribution is usually unknown. The next section proposes a simple and efficient method to do so, leveraging the connection between (15) and copulas. Before doing that, note again that under the assumption that we are well calibrated, we can transform (15) into

$$F(\alpha^1_s, \ldots, \alpha^m_s) = 1 - \epsilon_g, \tag{16}$$

where F denotes here the joint cumulative distribution induced by $\mathbb{P}$.

3 Copula-based conformal Multi-Target Regression

This section introduces our approach to obtain valid or better conformal prediction in the multi-variate regression setting. We first recall some basics of copulas and refer to Nelsen [10] for a full introduction, before detailing how we apply them to conformal approaches.

---

[1]Note that there may be multiple choices for such individual levels. Here we will fix them to be equal for simplicity.

3.1 Overview on copulas

A copula is a mathematical function that can describe the dependence between multiple random variables. The term "copula" was first introduced by Sklar [21] in his famous theorem, which is one of the fundamentals of copula theory, now known as Sklar's theorem. However, these tools have already been used before, as for instance in Fréchet's paper [22] and Höffding's work [23, 24] (reprinted as [25]). Copulas are popular in the statistical and financial fields [26], but they are nowadays more and more used in other domains as well, such as hydrology [27], medicine [28], and machine learning [29].

Let $\mathbf{Q} = (Q^1, \ldots, Q^m)$ be an m-dimensional random vector composed of the random variables $Q^1, \ldots, Q^m$. Let its cumulative distribution function (c.d.f.) be $F = F_Q : \mathbb{R}^m \to [0,1]$. This c.d.f. carries two important pieces of information:

- The c.d.f. of each random variable Q^j s.t. $F_j(q^j) = \mathbb{P}(Q^j \leq q^j)$, for all $j \in \{1, \ldots m\}$.
- The dependence structure between them.

The objective of copulas is to isolate the dependence structure from the marginals Q^j by transforming them into uniformly distributed random variables U^j and then expressing the dependence structure between the U^j's. In other words, an m-dimensional copula $C : [0,1]^m \to [0,1]$ is a c.d.f. with standard uniform marginals. It is characterized by the following properties:

1. C is grounded, i.e. if $u^j = 0$ for at least one $j \in \{1, \ldots, m\}$, then $C(u^1, \ldots, u^m) = 0$.
2. If all components of C are equal to 1 except u^j for all $u^j \in [0,1]$ and $j \in \{1, \ldots, m\}$, then $C(1, \ldots, 1, u^j, 1, \ldots, 1) = u^j$.
3. C is m-increasing, i.e., for all $\mathbf{a}, \mathbf{b} \in [0,1]^m$ with $\mathbf{a} \leq \mathbf{b}$:

$$\Delta_{(\mathbf{a}, \mathbf{b}]} C = \sum_{j \in \{0,1\}^m} (-1)^{\sum_{k=1}^m j_k} C(a_1^{j_1} b_1^{1-j_1}, \ldots, a_m^{j_m} b_m^{1-j_m}) \geq 0.$$

The last inequality simply ensures that the copula is a well-defined c.d.f. inducing non-negative probability for every event. The idea of copulas is based on probability and quantile transformations [30]. Using these latter, we can see that all multivariate distribution functions include copulas and that we can use a mixture of univariate marginal distributions and a suitable copula to produce a multivariate distribution function. This is described in Sklar's theorem [21] as follows:

Theorem 3.1 (Sklar's theorem) *For any m-dimensional cumulative distribution function (c.d.f.) F with marginal distributions $F_1, \ldots, F_m$, there exists a copula $C : [0,1]^m \to [0,1]$ such that:*

$$F(\mathbf{q}) = F(q^1, \ldots, q^m) = C(F_1(q^1), \ldots, F_m(q^m)), \quad \mathbf{q} \in \mathbb{R}^m. \tag{17}$$

If F_j is continuous for all $j \in \{1, \ldots, m\}$, then C is unique.

Denoting the pseudo inverse of F_j as $F_j^{\leftarrow}$ [30], we can get from (17) that

$$C(\mathbf{u}) = C(u^1, \ldots, u^m) = F(F_1^{\leftarrow}(u^1), \ldots, F_m^{\leftarrow}(u^m)). \tag{18}$$

There are a few noticeable copulas, among which are:

- the product copula: $\Pi(\mathbf{u}) = \prod_{j=1}^m u^j$;
- the Fréchet-Höffding upper bound copula [2]: $M(\mathbf{u}) = \min_{1 \leq j \leq m}\{u^j\}$;
- the Fréchet-Höffding lower bound copula [3]: $W(\mathbf{u}) = \max\{\sum_{j=1}^m u^j - m + 1, 0\}$.

While the product copula corresponds to classical stochastic independence, the Fréchet-Höffding bound copulas play an important role as they correspond to extreme cases of dependence [31]. Indeed, any m-dimensional copula C is such that $W(\mathbf{u}) \leq C(\mathbf{u}) \leq M(\mathbf{u}), \mathbf{u} \in [0,1]^m$.

Another important class of copulas are so-called Archimedean copulas, which are based on generator functions ϕ of specific kinds. More precisely, a continuous, strictly decreasing, convex function $\phi : [0,1] \to [0, \infty]$ satisfying

---

[2] M is a copula for all $m \geq 2$.

[3] W is a copula if and only if $m = 2$.

$\phi(1) = 0$ is known as an Archimedean copula generator. It is known as a strict generator if $\phi(0) = \infty$. The generated copula is then given by

$$C(u^1, \ldots, u^m) = \phi^{[-1]}(\phi(u^1) + \ldots + \phi(u^m)). \tag{19}$$

Table 1 provides examples and details of three one parameter Archimedean copula families [30], which are particularly convenient in estimation problems (being based on a single parameter).

Family	Generator $\phi(t)$	θ range	Strict	Lower	Upper
Gumbel [32]	$(-\ln t)^\theta$	$\theta \geq 1$	Yes	Π	M
Clayton [33]	$\frac{1}{\theta}(t^{-\theta} - 1)$	$\theta \geq -1$	$\theta \geq 0$	W	M
Frank [34]	$-\ln\left(\frac{e^{-\theta t} - 1}{e^{-\theta} - 1}\right)$	$\theta \in \mathbb{R}$	Yes	W	M

Table 1: Archimedean copula families.

3.2 Copula-based conformal Multi-Target Regression

Let us now revisit our previous problem of finding the significance levels ϵ_j for each target so that the hyper-rectangle prediction $[\hat{y}]$ covers the true value with confidence $1 - \epsilon_g$. Let us first consider (16). Following Sklar's theorem, we have

$$\begin{aligned}
F(\alpha_s^1, \ldots, \alpha_s^m) &= C(F_1(\alpha_s^1), \ldots, F_m(\alpha_s^m)) \\
&= C(1 - \epsilon^1, \ldots, 1 - \epsilon^m) \\
&= 1 - \epsilon_g
\end{aligned}$$

where the second line is obtained from (6). Clearly, if we knew the copula C, then we could search for values ϵ_j providing the desired global confidence.

A major issue is then to obtain or estimate the copula modelling the dependence structure between the targets and their confidence levels. As copulas are classically estimated from multi-variate observations, a simple means that we will use here is to estimate them from the non-conformity scores generated from the calibration set Z^{cal}. Namely, if α_i^j is the non-conformity score corresponding to the j^{th} target of the z_i example of Z^{cal} for $i \in \{l+1, \ldots, n\}$, we simply propose to estimate a copula C from the matrix

$$A = \begin{bmatrix} \alpha_{l+1}^1 & \alpha_{l+1}^2 & \cdots \\ \vdots & \ddots & \\ \alpha_n^1 & & \alpha_n^m \end{bmatrix}. \tag{20}$$

3.3 On three specific copulas

We will now provide some detail about the copulas we performed experiments on. They have been chosen to go from the one requiring the most assumptions to the one requiring the least assumptions.

3.3.1 The Independent copula

The Independent copula means that the m targets are considered as being independent, with no relationship between them. It is a strong assumption, but it does not require any estimation of the copula. In this case, (15) becomes:

$$\Pi(F_1(\alpha_s^1), \ldots, F_m(\alpha_s^m)) = \prod_{j=1}^m F_j(\alpha_s^j) = \prod_{j=1}^m \mathbb{P}(Q^j \leq \alpha_s^j)$$

$$\geq \prod_{j=1}^m (1 - \epsilon^j) = 1 - \epsilon_g,$$

If we assume that all $\epsilon^1, \ldots, \epsilon^m$ equal the same value ϵ_t, then:

$$\prod_{j=1}^m (1 - \epsilon^j) = (1 - \epsilon_t)^m = 1 - \epsilon_g.$$

6

Thus, we simply obtain

$$\epsilon_t = 1 - \sqrt[m]{1 - \epsilon_g}. \tag{21}$$

This individual significance level ϵ_t is then used to calculate the different non-conformity scores α_s^j for each target in the multi-target regression problem for the Independent copula.

3.3.2 The Gumbel copula

The Gumbel copula is a member of the Archimedean copula family which depends on only one parameter, and in this sense is a good representative of parametric copulas. It comes down to applying the generator function $\phi(F_j(\alpha_s^j)) = (-\ln F_j(\alpha_s^j))^\theta$ and its inverse $\phi^{[-1]}(F_j(\alpha_s^j)) = \exp -(F_j(\alpha_s^j))^{1/\theta}$ to (19), resulting in the expression

$$C_G^\theta(F_1(\alpha_s^1), \ldots, F_m(\alpha_s^m)) = \exp - \left(\sum_{j=1}^m \left(-\ln F_j(\alpha_s^j) \right)^\theta \right)^{1/\theta}. \tag{22}$$

In this case, we need to estimate the parameter θ. Since the marginals $F_j(\alpha^j)$ are unknown, we also need to estimate them. In our case, we will simply use the empirical c.d.f. induced by the non-conformity scores α_i^j of matrix A. An alternative would be to also assume a parametric form of the F_j, but this seems in contradiction with the very spirit of non-conformity scores. In particular, we will denote by $\hat{F}_j$ the empirical cumulative distribution such that

$$\hat{F}_j(\beta) = \frac{|\{\alpha_i^j : \alpha_i^j \leq \beta, i \in \{l+1, \ldots, n\}\}|}{n - l}, \quad \beta \in \mathbb{R}.$$

The parameter θ can then be estimated from matrix A using the Maximum Pseudo-Likelihood Estimator [35] with a numerical optimization, for instance by using the Python library "copulae"[4]. Once this is obtained, we then get for a particular choice of ϵ_j that

$$C_G^{\hat{\theta}} = \exp - \left(\sum_{j=1}^m \left(-\ln(1 - \epsilon_j) \right)^{\hat{\theta}} \right)^{1/\hat{\theta}} \tag{23}$$

$$= \exp - \left(\sum_{j=1}^m \left(-\ln F_j(\alpha_s^j) \right)^{\hat{\theta}} \right)^{1/\hat{\theta}} \tag{24}$$

And we can search for values ϵ_j that will make this equation equal to $1 - \epsilon_g$, using the estimations $\hat{F}_j$. The solution is especially easy to obtain analytically if we consider that $\epsilon^1 = \ldots = \epsilon^m = \epsilon_t$, as we then have that

$$\epsilon_t = 1 - (1 - \epsilon_g)^{1/\sqrt[\hat{\theta}]{m}},$$

and one can then obtain the corresponding non-conformity scores $\alpha_s^1, \ldots, \alpha_s^m$ by replacing F_j by $\hat{F}_j$.

We chose this particular family of Archimedean copulas because its lower bound is the Independent copula (as seen in Table 1). We can easily verify this by taking $\hat{\theta} = 1$. Thus, we can capture independence if it is verified, and otherwise search in the direction of positive dependence. One reason for such a choice is that previous experiments [9] indicate that the product copula gives overly conservative results.

3.3.3 The Empirical copula

Parametric copulas, as all parametric models, have the advantage of requiring less data to be well estimated, while having the possibly important disadvantage that they induce some bias in the estimation, that is likely to grow as the number of target increases. The Empirical copula presents a non-parametric way of estimating the marginals directly from the observations [36, 37]. It is defined as follows [35]:

$$C_E(\mathbf{u}) = \frac{1}{n - l} \sum_{i=l+1}^n \mathbb{1}_{\mathbf{u}_i \leq \mathbf{u}} = \frac{1}{n - l} \sum_{i=l+1}^n \prod_{j=1}^m \mathbb{1}_{u_i^j \leq u^j}, \quad \mathbf{u} \in [0, 1]^m, \tag{25}$$

---

[4]https://pypi.org/project/copulae/

where $\mathbb{1}_A$ is the indicator function of event A, and the inequalities $\mathbf{u}_i \leq \mathbf{u}$ for $i \in \{l+1, \ldots, n\}$ need to be understood component-wise. $\mathbf{u}_i$ are the pseudo-observations that replace the unknown marginal distributions, which are defined as:

$$\mathbf{u}_i = (u_i^1, \ldots, u_i^m) = (\hat{F}_1(\alpha_i^1), \ldots, \hat{F}_m(\alpha_i^m)), \quad i \in \{l+1, \ldots, n\}, \tag{26}$$

where distributions $\hat{F}_j$ are defined as before. Simply put, the Empirical copula corresponds to consider as our joint probability the Empirical joint cumulative distribution. We then have that

$$C_E(F_1(\alpha_s^1), \ldots, F_m(\alpha_s^m)) = \frac{1}{n-l} \sum_{i=l+1}^{n} \prod_{j=1}^{m} \mathbb{1}_{u_i^j \leq F_j(\alpha_s^j)}. \tag{27}$$

Using that $F_j(\alpha_s^j) = 1 - \epsilon_j$, we can then search for values of ϵ_j, $j = 1, \ldots, m$ that will make (27) equal to $1 - \epsilon_g$. Note that in this case, even assuming that $\epsilon^1 = \ldots = \epsilon^m = \epsilon_t$ will require an algorithmic search, which is however easy as C_E is an increasing function, meaning that we can use a simple dichotomic search.

4 Evaluation

In this section, we describe the experimental setting (underlying algorithm, data sets and performance metrics) and the results of our study.

4.1 Experimental setting

We choose to work with a deep neural network as the underlying algorithm. We keep the same underlying algorithm for all non-conformity measures, since our focus is to compare between the three copula functions chosen to get the different non-conformity scores.

To compute the non-conformity scores over the calibration set, we use the normalized non-conformity score given by (8) as described in [12], and predict $\mu_i = \ln(|y_i - \hat{y}_i|)$ simultaneously for all targets by a single multivariate multi-layer perceptron. In this case, μ_i represents the estimation of the underlying algorithm's error. As mentioned before, the approach can be adapted to any conformal regression approach.

Experiments are conducted on normalized data with a mean of 0 and a standard deviation of 1 to simplify the deep neural network optimization, with a 10-fold cross validation to avoid the impact of biased results, and with a calibration set equal to 10% of the training examples for all data sets. We take the value $\beta = 0.1$ for the sensitivity parameter and do not optimize it when calculating the normalizing coefficient μ_i. After getting the proper training data (X^{tr}, Y^{tr}), calibration data (X^{cal}, Y^{cal}) and test data (X^{ts}, Y^{ts}) for each fold, we follow the steps described below:

1. Train the underlying algorithm (a deep neural network) on the proper training data (X^{tr}, Y^{tr}). Its architecture is composed of a first dense layer applied to the input with "selu" activation (scaled exponential linear units [38]), three hidden dense layers with dropouts and "selu" activation, and a final dense layer with m outputs and a linear activation.

2. Predict $\hat{Y}^{cal}$ and $\hat{Y}^{ts}$ for calibration and test data respectively using the underlying algorithm.

3. Train the normalizing multi-layer perceptron on the proper training data $(X^{tr}, \mu_{tr} = \ln(|Y^{tr} - \hat{Y}^{tr}|))$, corresponding to the error estimation of the underlying algorithm. The normalizing MLP consists of three hidden dense layers with "selu" activation and dropouts and a final dense layer with m outputs for predicting all targets simultaneously.

4. Predict μ_{cal} and μ_{ts} for calibration and test data respectively using the normalizing MLP.

5. If needed, get an estimation[5] of the copula C from the matrix A of calibration non-conformity scores.

6. For each global significance level ϵ_g:
   - Get the individual significance level $\epsilon_j = \epsilon_t$ for $j \in \{1, \ldots, m\}$ and calculate $\alpha_s = \{\alpha_s^1, \ldots, \alpha_s^m\}$ for all targets using calibration data, according to the methods mentioned in Section 3.3.
   - Get the interval predictions for the test data with:

$$\left[\hat{Y}^{ts} - \alpha_s(\exp(\mu_{ts}) + \beta), \hat{Y}^{ts} + \alpha_s(\exp(\mu_{ts}) + \beta)\right]. \tag{28}$$

---

[5]In the case of the Gumbel copula, we use a Maximum Pseudo-Likelihood Estimator with a numerical optimization using the BFGS algorithm

Names	Examples	Features	Targets
music origin [39]	1059	68	2
indoor loc [40]	21049	520	3
scpf [41]	1137	23	3
sgemm [42]	241600	14	4
rf1 [41]	9125	64	8
rf2 [41]	9125	576	8
scm1d [41]	9803	280	16
scm20d [41]	8966	61	16

Table 2: Information on the used multi-target regression data sets.

Remark 4.1 *We choose $\epsilon_j = \epsilon_t$ for $j \in \{1, \ldots, m\}$ as we have no indication that individual targets should be treated with different degree of cautiousness. However, since copulas are functions from $[0,1]^m$ to $[0,1]$, there is in principle no problem in considering different confidence degrees for different tasks, if an application calls for it. How to determine and elicit such degrees is however, to our knowledge, an open question.*

The implementation was done using Python and Tensorflow. The copula part of our experiments was based on the book [35] and the Python library "copulae".

We use eight data sets with different numbers of targets and varying sizes. They are summarized in Table 2.

4.2 Results

This section presents the results of our experiments, investigating in particular the validity and efficiency of the proposed approaches. Figures 1 and 2 detail these results for "music origin" and "sgemm". The figures for all other data sets can be found in A.

To verify the validity of each non-conformity measure, we calculate the accuracy of each one and compare it with the calibration line. This line represents the case where the error rate is exactly equal to ϵ_g for a confidence level $1 - \epsilon_g$, which is the desired outcome of using conformal prediction. In multi-target regression, the accuracy is computed based on whether the observation y belongs to the hyper-rectangle $[\hat{y}]$ or not depending on the significance level ϵ_g. Thus, a correctly predicted example must verify that all its individual predictions y_i for each individual target Y_i is in its corresponding individual interval predictions. Concretely, for each considered confidence level ϵ_g and test example $x \in X^{ts}$, we obtain a prediction $[\hat{y}]_{\epsilon_g}$. From this, we can compute the empirical validity as the percentage of times that $[\hat{y}]_{\epsilon_g}$ contains the true observed value, i.e.,

$$\frac{\sum_{(x,y) \in Z^{ts}} \mathbb{1}_{y \in [\hat{y}]_{\epsilon_g}}}{|Z^{ts}|}.$$

Doing it for several values of ϵ_g, we obtain a calibration curve that should be as close as possible to the identity function.

The results of the error rate or accuracy curves are shown in sub-figure a of each figure for the Independent, Gumbel and Empirical multivariate non-conformity measures. The outcomes clearly show that the best performance is obtained by using the Empirical copula, where the model is well calibrated. For most of the studied data sets, the Empirical copula accuracy curve is almost perfectly aligned with the calibration line, and thus almost exactly valid. This is due to the fact that Empirical copula functions non-parametrically estimate the marginals based on the observations, which enables the model to better adapt to the dependence structure of each data set. This dependence structure is neglected when using an Independent copula-based non-conformity measure, as the m targets are treated as if they were independent, and so the link between them is not exploited when computing ϵ_t. This also means that the difference between the Empirical and the Independent copula-based non-conformity measures is bigger when there is a strong dependence between the non-conformity scores, and is an indication of the strength of this dependence. For instance, we can deduce that the targets are strongly related for "sgemm" by the big gap between the Independent and Empirical accuracy curves (sub-figure 2a). For the Gumbel copula, the accuracy curve is generally closer to the calibration line than the one for the Independent copula. This supports the existence of a dependence structure between the targets, since the lower bound of the Gumbel copula is the Independent copula, which means that if the targets were in fact independent, the two curves would perfectly match. This can be seen in sub-figure 1a for "music origin", where the accuracy curves almost overlap all the time, meaning that the targets are likely to be independent.
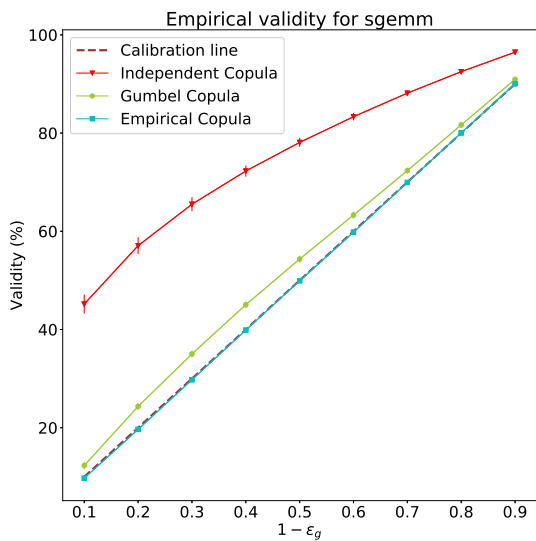
9
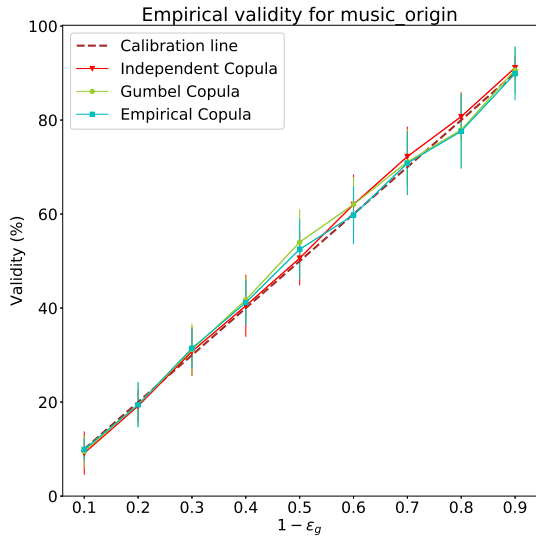
(a) Empirical validity
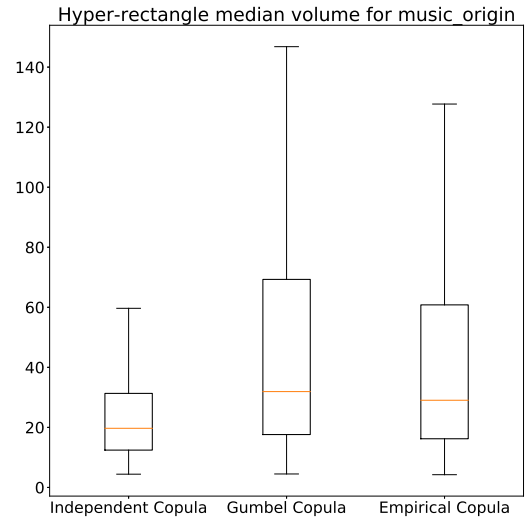
(b) Hyper-rectangle median volume

Figure 1: Results for music origin.

(a) Empirical validity

(b) Hyper-rectangle median volume

Figure 2: Results for sgemm.

Data sets	Independent	Gumbel	Empirical
music origin	$7.06 \times 10^1 \pm 5.12$	$8.48 \times 10^1 \pm 5.72$	$\mathbf{2.90 \times 10^1 \pm 5.48}$
indoor loc	2.99 ± 1.17	2.00 ± 1.28	$\mathbf{3.24 \times 10^{-1} \pm 1.28}$
scpf	9.04 ± 5.07	2.73 ± 5.64	$\mathbf{-1.42 \pm 4.16}$
sgemm	$2.54 \times 10^1 \pm 1.00$	$3.26 \pm 6.53 \times 10^{-1}$	$\mathbf{-1.35 \times 10^1 \pm 3.00 \times 10^{-1}}$
rf1	5.60 ± 1.59	3.46 ± 1.56	$\mathbf{-9.35 \times 10^{-3} \pm 1.51}$
rf2	6.09 ± 1.86	2.19 ± 2.27	$\mathbf{-3.61 \times 10^{-1} \pm 2.14}$
scm1d	$1.44 \times 10^1 \pm 1.82$	$1.03 \times 10^1 \pm 2.98$	$\mathbf{-7.03 \times 10^{-1} \pm 2.32}$
scm20d	$1.68 \times 10^1 \pm 1.43$	$1.02 \times 10^1 \pm 2.35$	$\mathbf{-1.34 \pm 2.25}$

Table 3: Validity (average gap between the empirical validity curve and the calibration line in percentage) summarized results for all data sets.

Data sets	Independent	Gumbel	Empirical
music origin	$\mathbf{1.97 \times 10^1 \pm 2.99}$	$3.19 \times 10^1 \pm 1.73 \times 10^1$	$2.90 \times 10^1 \pm 1.39 \times 10^1$
indoor loc	$1.70 \times 10^{-1} \pm 5.12 \times 10^{-2}$	$9.54 \times 10^{-2} \pm 2.04 \times 10^{-2}$	$\mathbf{8.69 \times 10^{-2} \pm 1.86 \times 10^{-2}}$
scpf	5.10 ± 5.31	3.06 ± 3.7	$\mathbf{2.39 \pm 3.67}$
sgemm	$1.17 \times 10^{-3} \pm 5.69 \times 10^{-4}$	$2.56 \times 10^{-4} \pm 1.95 \times 10^{-4}$	$\mathbf{2.20 \times 10^{-4} \pm 1.60 \times 10^{-4}}$
rf1	$1.18 \times 10^{-2} \pm 1.52 \times 10^{-2}$	$5.52 \times 10^{-3} \pm 1.05 \times 10^{-2}$	$\mathbf{3.61 \times 10^{-3} \pm 6.00 \times 10^{-3}}$
rf2	$2.56 \times 10^{-3} \pm 1.87 \times 10^{-3}$	$7.48 \times 10^{-4} \pm 8.44 \times 10^{-4}$	$\mathbf{7.00 \times 10^{-4} \pm 8.48 \times 10^{-4}}$
scm1d	$3.49 \times 10^4 \pm 2.89 \times 10^4$	$1.28 \times 10^4 \pm 1.20 \times 10^4$	$\mathbf{1.15 \times 10^3 \pm 1.22 \times 10^3}$
scm20d	$5.43 \times 10^6 \pm 4.43 \times 10^6$	$1.80 \times 10^5 \pm 2.15 \times 10^5$	$\mathbf{4.14 \times 10^4 \pm 6.66 \times 10^4}$

Table 4: Efficiency (hyper-rectangle median volume for $\epsilon_g = 0.1$) summarized results for all data sets.

From the empirical validity results, we also noticed that the Empirical copula non-conformity measure can be slightly invalid sometimes (sub-figure 4a for "scpf"). We explain this by the fewer number of examples, in which case one could use a more regularized form than the Empirical copula. However, when a lot of examples are available (for instance, more than 20000 observations for "sgemm"), the validity curve of the Empirical copula non-conformity measure is perfectly aligned with the calibration line, meaning that this measure is exactly valid (sub-figure 2a).

In single-output regression, efficiency is measured by the size of the intervals, and a method is all the more efficient as predicted intervals are small. To assess efficiency in multi-target regression, we can simply compute the volume of the obtained predictions $[\hat{\mathbf{y}}]_{\epsilon_g}$, after (12). For each experiment, we then compute the median value of those hyper-rectangle volumes (for the estimation to be robust against very large hyper-rectangles).

Efficiency results are shown in sub-figure b for all data sets for $\epsilon_g = 0.1$. They show that, in general, the Independent copula has a bigger median hyper-rectangle volume compared to the Gumbel and Empirical copulas, especially in those cases where the existence of a dependence structure is confirmed by the calibration curves. This is due to the fact that using an Independent copula ignores the dependence between the non-conformity scores, which leads to an over-estimation of the global hyper-rectangle error. This impact is avoided when using the Empirical copula because it takes advantage of the dependence structure to construct better interval predictions. Another remark concerning efficiency is that the box plots for Empirical copula are tighter than the other two, which shows that the values are homogeneous on all folds compared to the Independent copula for instance, where the variation is much more visible.

The empirical validity and hyper-rectangle median volume results are summarized in Tables 3 and 4. The validity simply provides the average difference between a perfect calibration (the identity function) and the observed curve for each copula. This means, in particular, that a negative value indicates that the observed frequency is in average below the specified confidence degree.

The numbers confirm our previous observations on the graphs, as the average gap is systematically higher for the Independent copula and lower for the Empirical one, with Gumbel in-between. We can however notice that while the Empirical copula provides the best results, it is also often a bit under the calibration line, indicating that if conservativeness is to be sought, one should maybe prefer the Gumbel copula. About the same conclusions can be given regarding efficiency, with the Empirical copula giving the best results and the Independent one the worst.

5   Conclusion and discussion

In this paper, we provided a quite easy and flexible way to obtain valid conformal predictions in a multi-variate regression setting. We did so by exploiting a link between non-conformity scores and copulas, a commonly used tool to model multi-variate distribution.

Experiments on various data sets for a small choice of representative copulas show that the method indeed allows to improve upon the naive independence assumption. Those first results indicate in particular that while parametric, simple copulas may provide valid results for some data sets, more complex copulas may be needed in general to obtain well calibrated predictions, with the cost that good estimations of such copulas require a lot of calibration data.

As future lines of work, we would like to explore further the flexibility of our framework, for instance by exploring the possibility of using vines [43] to model complex dependencies, or by proposing protocols allowing to obtain ϵ_g from different individual, user-defined confidence degrees, taking up on our Remark 4.1.

Finally, while we mostly focused on multi-variate regression in the present paper, it would be interesting to try to extend the current approach to other multi-task settings, such as multi-label problems. A possibility could be to make such problems continuous, as proposed for instance by Liu [29].

6   Acknowledgments

References

[1] Glenn Shafer and Vladimir Vovk. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(Mar):371–421, 2008.

[2] Harris Papadopoulos, Kostas Proedrou, Volodya Vovk, and Alex Gammerman. Inductive confidence machines for regression. In *European Conference on Machine Learning*, pages 345–356. Springer, 2002.

[3] Rich Caruana. A dozen tricks with multitask learning. In *Neural networks: tricks of the trade*, pages 165–191. Springer, 1998.

[4] Huazhen Wang, Xin Liu, Ilia Nouretdinov, and Zhiyuan Luo. A comparison of three implementations of multi-label conformal prediction. In *International Symposium on Statistical Learning and Data Sciences*, pages 241–250. Springer, 2015.

[5] Ran Wang, Sam Kwong, Xu Wang, and Yuheng Jia. Active k-labelsets ensemble for multi-label classification. *Pattern Recognition*, page 107583, 2020.

[6] Alexander Kuleshov, Alexander Bernstein, and Evgeny Burnaev. Conformal prediction in manifold learning. In *Conformal and Probabilistic Prediction and Applications*, pages 234–253, 2018.

[7] Jelmer Neeven and Evgueni Smirnov. Conformal stacked weather forecasting. In *Conformal and Probabilistic Prediction and Applications*, pages 220–233, 2018.

[8] Harris Papadopoulos, Vladimir Vovk, and Alexander Gammerman. Regression conformal prediction with nearest neighbours. *Journal of Artificial Intelligence Research*, 40:815–840, 2011.

[9] Soundouss Messoudi, Sébastien Destercke, and Sylvain Rousseau. Conformal multi-target regression using neural networks. In *Conformal and Probabilistic Prediction and Applications*, pages 65–83. PMLR, 2020.

[10] Roger B Nelsen. An introduction to copulas, volume 139 of. *Lecture Notes in Statistics*, 1999.

[11] Alex Gammerman, Volodya Vovk, and Vladimir Vapnik. Learning by transduction. *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, page 148–155, 1998.

[12] Harris Papadopoulos and Haris Haralambous. Reliable prediction intervals with regression neural networks. *Neural Networks*, 24(8):842–851, 2011.

[13] Ulf Johansson, Henrik Linusson, Tuve Löfström, and Henrik Boström. Interpretable regression trees using conformal prediction. *Expert systems with applications*, 97:394–404, 2018.

[14] Matilde Sánchez-Fernández, Mario de Prado-Cumplido, Jerónimo Arenas-García, and Fernando Pérez-Cruz. Svm multiregression for nonlinear channel estimation in multiple-input multiple-output systems. *IEEE transactions on signal processing*, 52(8):2298–2307, 2004.

[15] Glenn De'Ath. Multivariate regression trees: a new technique for modeling species–environment relationships. *Ecology*, 83(4):1105–1117, 2002.

[16] Luca Baldassarre, Lorenzo Rosasco, Annalisa Barla, and Alessandro Verri. Multi-output learning via spectral filtering. *Machine learning*, 87(3):259–301, 2012.

[17] Timo Aho, Bernard Ženko, and Sašo Džeroski. Rule ensembles for multi-target regression. In *2009 Ninth IEEE International Conference on Data Mining*, pages 21–30. IEEE, 2009.

[18] Eleftherios Spyromitros-Xioufis, Grigorios Tsoumakas, William Groves, and Ioannis Vlahavas. Multi-target regression via input space expansion: treating targets as inputs. *Machine Learning*, 104(1):55–98, 2016.

[19] Sebastian Ruder. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*, 2017.

[20] Rich Caruana. Multitask learning: A knowledge-based source of inductive bias icml. *Google Scholar Google Scholar Digital Library Digital Library*, 1993.

[21] M Sklar. Fonctions de repartition an dimensions et leurs marges. *Publ. inst. statist. univ. Paris*, 8:229–231, 1959.

[22] Maurice Fréchet. Sur les tableaux de corrélation dont les marges sont données. *Ann. Univ. Lyon, 3ˆ e serie, Sciences, Sect. A*, 14:53–77, 1951.

[23] Wassilij Höffding. Masstabinvariante korrelationstheorie. *Schriften des Mathematischen Instituts und Instituts fur Angewandte Mathematik der Universitat Berlin*, 5:181–233, 1940.

[24] Wassily Höffding. Masstabinvariante korrelationsmasse für diskontinuierliche verteilungen. *Archiv für mathematische Wirtschafts-und Sozialforschung*, 7:49–70, 1941.

[25] Wassily Höffding. Scale—invariant correlation theory. In *The collected works of Wassily Höffding*, pages 57–107. Springer, 1994.

[26] Paul Embrechts, Alexander McNeil, and Daniel Straumann. Correlation and dependence in risk management: properties and pitfalls. *Risk management: value at risk and beyond*, 1:176–223, 2002.

[27] Anne-Catherine Favre, Salaheddine El Adlouni, Luc Perreault, Nathalie Thiémonge, and Bernard Bobée. Multivariate hydrological frequency analysis using copulas. *Water resources research*, 40(1), 2004.

[28] Aristidis K Nikoloulopoulos and Dimitris Karlis. Multivariate logit copula model with an application to dental data. *Statistics in Medicine*, 27(30):6393–6406, 2008.

[29] Weiwei Liu. Copula multi-label learning. In *Advances in Neural Information Processing Systems*, pages 6337–6346, 2019.

[30] Alexander J McNeil, Rüdiger Frey, and Paul Embrechts. *Quantitative risk management: concepts, techniques and tools-revised edition*. Princeton university press, 2015.

[31] Thorsten Schmidt. Coping with copulas. *Copulas-From theory to application in finance*, pages 3–34, 2007.

[32] Emil Julius Gumbel. Distributions des valeurs extremes en plusiers dimensions. *Publ. Inst. Statist. Univ. Paris*, 9:171–173, 1960.

[33] Christian Genest and Louis-Paul Rivest. Statistical inference procedures for bivariate archimedean copulas. *Journal of the American statistical Association*, 88(423):1034–1043, 1993.

[34] Maurice J Frank. On the simultaneous associativity off (x, y) andx+y- f (x, y). *Aequationes mathematicae*, 19(1):194–226, 1979.

[35] Marius Hofert, Ivan Kojadinovic, Martin Mächler, and Jun Yan. *Elements of copula modeling with R*. Springer, 2019.

[36] Ludger Ruschendorf. Asymptotic distributions of multivariate rank order statistics. *The Annals of Statistics*, pages 912–923, 1976.

[37] Frederik Hendrik Ruymgaart. Asymptotic theory of rank tests for independence. *MC Tracts*, 1978.

[38] Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. Self-normalizing neural networks. In *Advances in neural information processing systems*, pages 971–980, 2017.

[39] Fang Zhou, Q Claire, and Ross D King. Predicting the geographical origin of music. In *2014 IEEE International Conference on Data Mining*, pages 1115–1120. IEEE, 2014.

[40] Joaquín Torres-Sospedra, Raúl Montoliu, Adolfo Martínez-Usó, Joan P Avariento, Tomás J Arnau, Mauri Benedito-Bordonau, and Joaquín Huerta. Ujiindoorloc: A new multi-building and multi-floor database for wlan fingerprint-based indoor localization problems. In *2014 international conference on indoor positioning and indoor navigation (IPIN)*, pages 261–270. IEEE, 2014.

[41] Grigorios Tsoumakas, Eleftherios Spyromitros-Xioufis, Jozef Vilcek, and Ioannis Vlahavas. Mulan: A java library for multi-label learning. *Journal of Machine Learning Research*, 12(71):2411–2414, 2011.

[42] Cedric Nugteren and Valeriu Codreanu. Cltune: A generic auto-tuner for opencl kernels. In *2015 IEEE 9th International Symposium on Embedded Multicore/Many-core Systems-on-Chip*, pages 195–202. IEEE, 2015.

[43] Harry Joe and Dorota Kurowicka. *Dependence modeling: vine copula handbook*. World Scientific, 2011.

A Validity and efficiency figures

This appendix contains the figures for empirical validity and hyper-rectangle median volume for all remaining data sets.

(a) Empirical validity

(b) Hyper-rectangle median volume
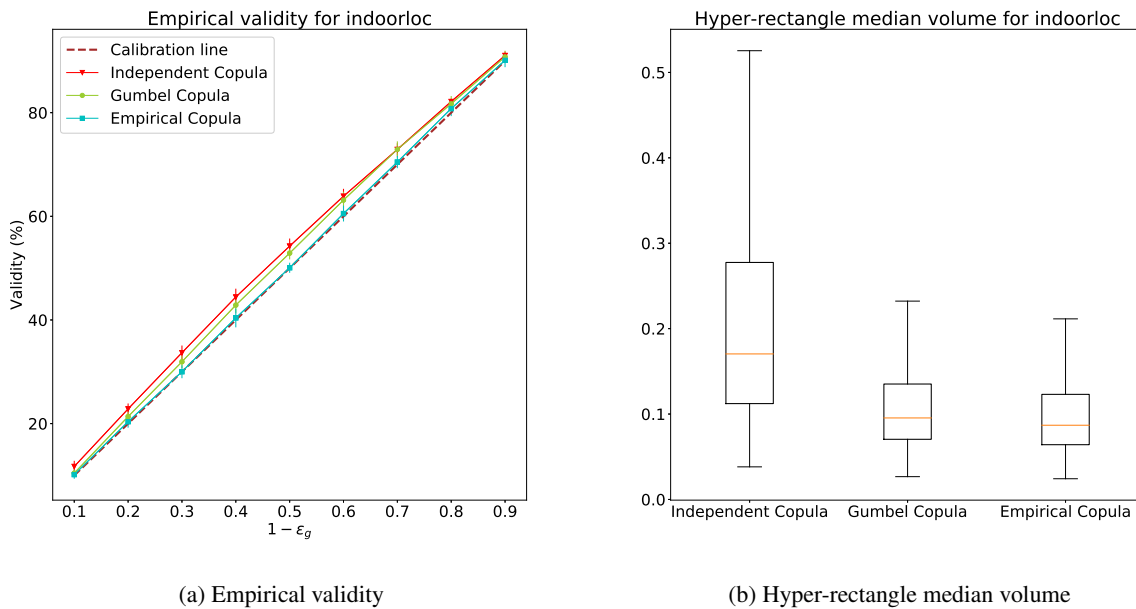
Figure 3: Results for indoor loc.

(a) Empirical validity
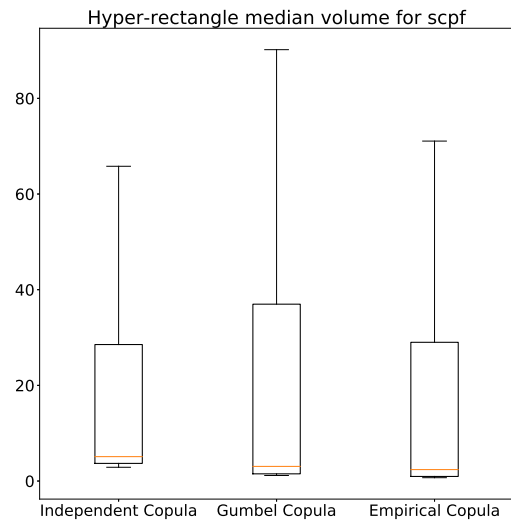
(b) Hyper-rectangle median volume

Figure 4: Results for scpf.
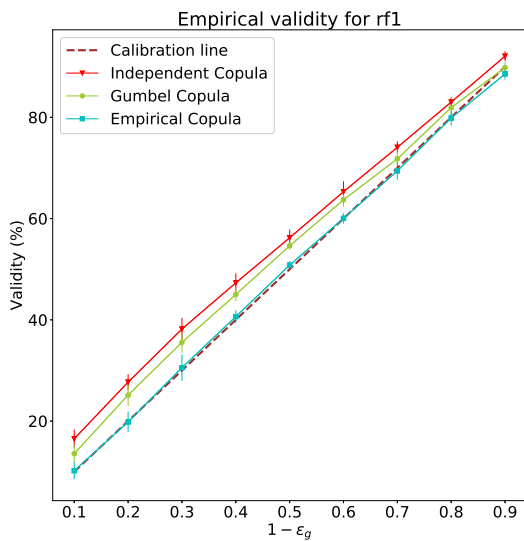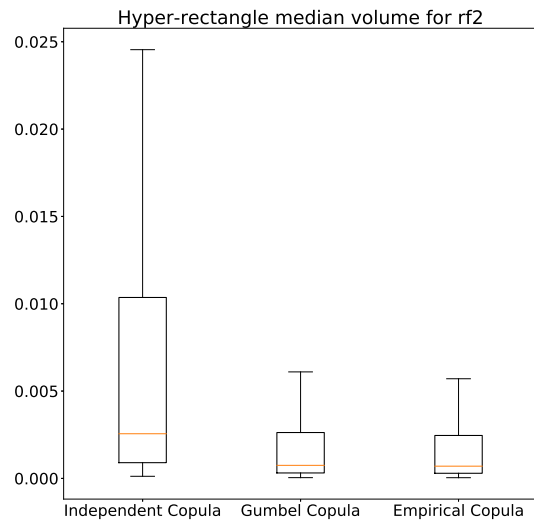

(a) Empirical validity
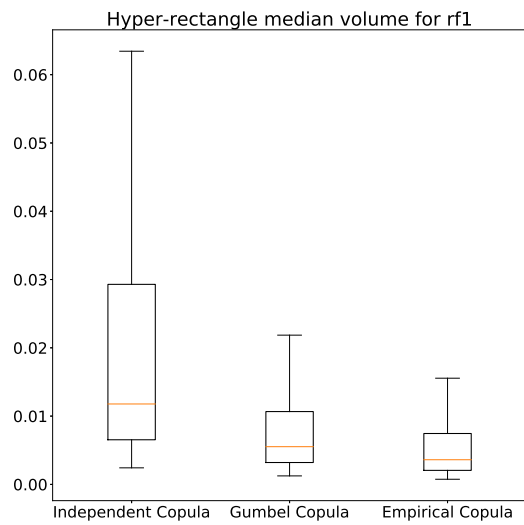
(b) Hyper-rectangle median volume
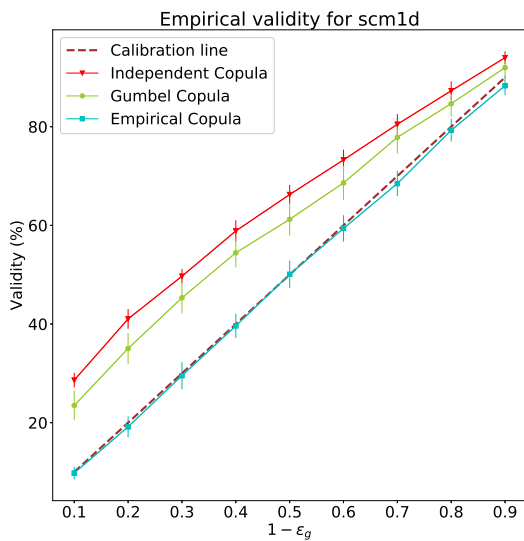
Figure 5: Results for rf1.
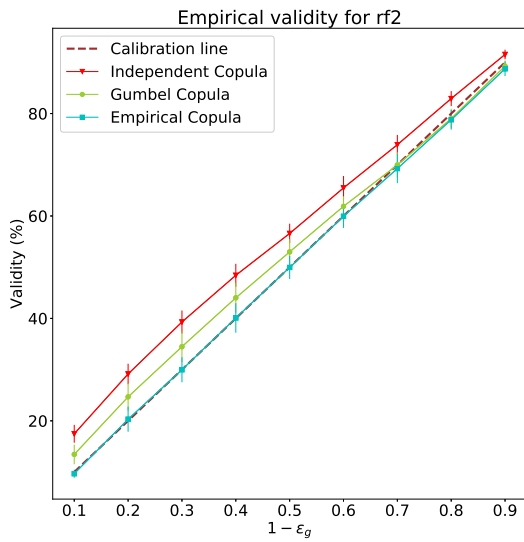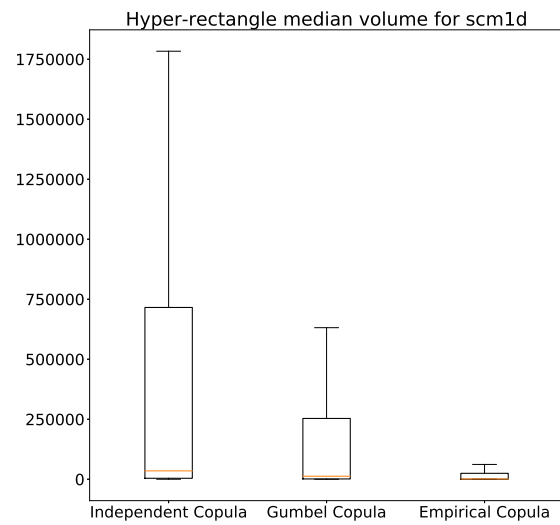
(a) Empirical validity

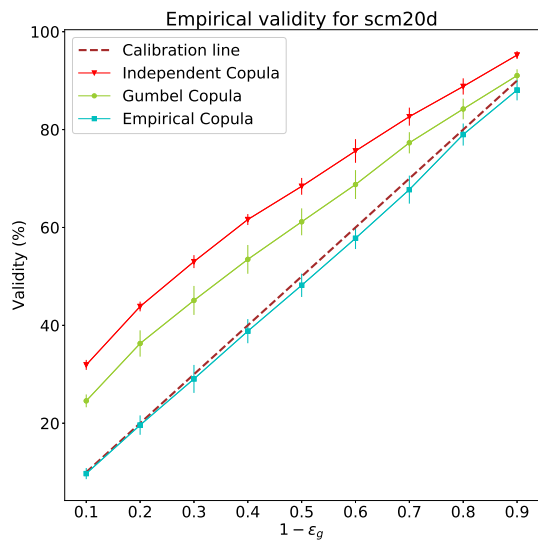(b) Hyper-rectangle median volume

Figure 6: Results for rf2.
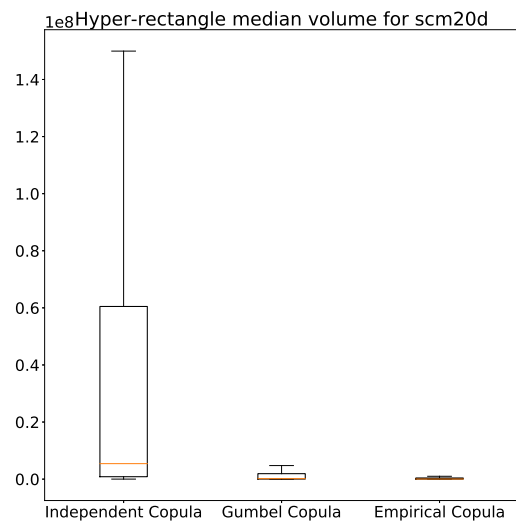

(a) Empirical validity

(b) Hyper-rectangle median volume

Figure 7: Results for scm1d.

(a) Empirical validity

(b) Hyper-rectangle median volume

Figure 8: Results for scm20d.