

Uncertain Trees: Dealing with Uncertain Inputs in Regression Trees

Myriam Tami¹, Marianne Clausel², Emilie Devijver¹, Adrien Dulac¹, Eric Gaussier¹,
Stefan Janaqi³ & Meriam Chebre⁴

¹ *University Grenoble Alpes, CNRS, Grenoble INP*, LIG, F-38000 Grenoble France*

² *University of Lorraine, Nancy · IEC - Institut Elie Cartan, 54052 Nancy France*

³ *Ecole des Mines d'Alès, 6 avenue de Clavières, F-30319 Ales Cedex France*

⁴ *TOTAL S.A., 24 cours Michelet - Défense 10, 92069 Paris La Défense Cedex France*

Abstract

Tree-based ensemble methods, as Random Forests and Gradient Boosted Trees, have been successfully used for regression in many applications and research studies. Furthermore, these methods have been extended in order to deal with uncertainty in the output variable, using for example a quantile loss in Random Forests [16]. To the best of our knowledge, no extension has been provided yet for dealing with uncertainties in the input variables, even though such uncertainties are common in practical situations. We propose here such an extension by showing how standard regression trees optimizing a quadratic loss can be adapted and learned while taking into account the uncertainties in the inputs. By doing so, one no longer assumes that an observation lies into a single region of the regression tree, but rather that it belongs to each region with a certain probability. Experiments conducted on several data sets illustrate the good behavior of the proposed extension.

1 Introduction

Most measures carried out in the real world, *e.g.*, by sensors embedded in different machines or by analyses of samples, are uncertain if not erroneous in some cases. This uncertainty may be due to the generating process of the samples being measured or from the intrinsic limitations of any measurement process. Considering such measures, that constitute many of the data sets used in data science applications in both industry and academy, as certain is thus in the best case a naive position. Our experiments illustrate this point inasmuch as the method we propose here to handle uncertainty outperforms standard approaches on several, real benchmark data sets.

However, if several methods have been developed to obtain uncertainty estimates from data sets, very few studies have been devoted to designing data science methods that can deal with such uncertainties. We

address this problem here in the context of regression trees, a popular machine learning method at the basis of widely used ensemble methods as Random Forests.

In this context, recent studies have focused on providing conditional quantiles, as opposed to conditional means, so as to better represent the output variable and avoid the uncertainty inherent to point estimates. The work on quantile regression forests developed by Meinshausen [16] is a good illustration of this. We take here a different approach and directly model the uncertainty of the input variables in the regression trees we consider. This leads to a regression tree in which observations no longer belong to a single leaf. Instead, each observation has a non-null probability of being assigned to any leaf of the tree. The construction process associated to such trees is slightly more complex than the ones of standard trees (it involves in particular the inversion of a $K \times K$ matrix, where K is the number of leaves of the tree), but the improvements obtained on the prediction fully justify this additional complexity, as shown in Section 4 on both benchmark and modified (with an additional noise) data sets.

The idea of including information on the uncertainty of the input data in a regression method is not new and is related to uncertainty propagation and sensitivity analysis. Several authors have indeed proposed to integrate uncertainties on the inputs in different regression methods, as multivariate linear models [17] or neural networks [11]. To the best of our knowledge, our approach is the first one to address this problem in the context of regression trees and ensemble methods based on such trees, as Random Forests.

The remainder of the paper is organized as follows: Section 2 presents the general model we rely on and its main properties; Section 3 then describes the algorithms for constructing the regression tree and the associated prediction rule, while Section 4 presents the experiments conducted and the numerical results obtained. Finally, Section 5 positions our work wrt previous studies while

*Institute of Engineering Univ. Grenoble Alpes

Section 6 concludes the paper.

2 Regression trees with uncertain input

Let Y be an output random variable and $\mathbf{X} = (X^1, \dots, X^p)$ a p -dimensional input random vector. A classical question is to find some relationship between \mathbf{X} and Y , estimating the so called link function f involved in the model $Y = f(\mathbf{X}) + \epsilon$.

Tree-based ensemble methods, as Random Forests or Gradient Boosted Trees, are popular machine learning methods, developed to address the above regression problem [13]. In these methods, a set of regressors are constructed and aggregated in a convenient way. The building blocks are decision regression trees [3], which are defined from a partition of the space \mathcal{X} of input variables into K regions $(R_k)_{1 \leq k \leq K}$ obtained by dyadic splits minimizing a risk function. A weight γ_k is associated to each region leading to a piece-wise predictor, for a new input $\mathbf{x} \in \mathcal{X}$, of the form:

$$(2.1) \quad T(\mathbf{x}; \Theta) = \sum_{k=1}^K \gamma_k \mathbb{1}_{\{\mathbf{x} \in R_k\}}$$

where $\Theta = (R_k, \gamma_k)_{1 \leq k \leq K}$ is the set of parameters, learned from a training data set, defining the tree. Both categorical and quantitative inputs can in theory be considered. For the sake of simplicity, we however focus in this study on quantitative inputs, thus considering that $\mathcal{X} = \mathbb{R}^p$.

To deal with uncertainty in the inputs, we introduce an auxiliary latent random vector \mathbf{U} representing the true value of the data and consider the general regression function that relates $\mathbf{U}|\mathbf{X}$ to Y through $Y = f(\mathbf{U}|\mathbf{X}) + \epsilon$. The standard regression model is obtained from this general model by considering that $U^j|X^j$ is a Dirac at $X^j = x^j$ (or equivalently is Gaussian with mean x^j and variance 0). We further assume here that the variables are independent of each other and that the true measure given the observation, on each variable, is Gaussian, leading to the following complete model:

$$(2.2) \quad \begin{cases} Y = f(\mathbf{U}|\mathbf{X}) + \epsilon, \\ U^j|X^j = x^j \sim \mathcal{N}(x^j, \sigma_{U_j}^2) \text{ for } 1 \leq j \leq p. \end{cases}$$

The Gaussian distribution is widespread, can be used to approximate several other distributions and is easy to manipulate, hence its use here. Other distributions can nevertheless be considered for the measurement error, but this is beyond the scope of this study. In the remainder, we will denote $\boldsymbol{\sigma}_{\mathbf{U}}^2 = (\sigma_{U_1}^2, \dots, \sigma_{U_p}^2)$ the vector of variances of the Gaussian distributions. In practical situations, these variances may be given (for

example when the data corresponds to measurements from machines for which the uncertainty is known) or may be directly learned from the data.

When one is dealing with uncertain inputs, and we want to stress again that this is the general situation, one can no longer assume that observations are hard assigned to regions. Instead, each observation has a probability of being associated to each region, leading to the following prediction rule that directly generalizes Eq. (2.1):

$$(2.3) \quad T_{\text{un}}(\mathbf{x}; \Theta) = \sum_{k=1}^K \gamma_k \mathbb{P}(\mathbf{U} \in R_k | \mathbf{X} = \mathbf{x}, \boldsymbol{\sigma}_{\mathbf{U}}^2).$$

Note that the set of parameters Θ now includes the variances $\boldsymbol{\sigma}_{\mathbf{U}}^2$. In addition, because of the independence assumption at the basis of the model retained, one has, for $1 \leq k \leq K$:

$$\mathbb{P}(\mathbf{U} \in R_k | \mathbf{X} = \mathbf{x}, \boldsymbol{\sigma}_{\mathbf{U}}^2) = \prod_{j=1}^p \mathbb{P}(U^j \in R_k^j | X^j = x^j, \sigma_{U_j}^2),$$

where the interval R_k^j corresponds to the region R_k projected onto the j th variable. We now turn to the procedure for learning the parameters of the model.

2.1 Estimation procedure The estimation procedure of the parameters is based, in this study, on the minimization of the empirical quadratic risk, which is the standard risk considered in regression trees. More precisely, for the learning set defined by:

$$(2.4) \quad \mathcal{L}_n = \left\{ (x_i^1, \dots, x_i^p, y_i)_{1 \leq i \leq n} \right\},$$

with $(\mathbf{x}_i, y_i)_{1 \leq i \leq n}$ the observed sample, we define the empirical risk:

$$(2.5) \quad R_{\text{emp}}(T_{\text{un}}(\cdot; \hat{\Theta}), \mathcal{L}_n) = \frac{1}{n} \sum_{i=1}^n (y_i - T_{\text{un}}(\mathbf{x}_i; \hat{\Theta}))^2$$

where T_{un} has been introduced in Eq. (2.3). This criterion has to be minimized on the training set wrt the parameters denoted by $\Theta = \{(R_k, \gamma_k)_{1 \leq k \leq K}, \boldsymbol{\sigma}_{\mathbf{U}}^2\}$: regions of the tree, associated weights, and the variances of the vector $\mathbf{U}|\mathbf{X}$. To do so, we introduce the matrix $P \in M_{n,K}(\mathbb{R})^1$ defined, for $1 \leq i \leq n$ and for $1 \leq k \leq K$, by:

$$\begin{aligned} P_{i,k} &= \mathbb{P}(\mathbf{U}_i \in R_k | \mathbf{X}_i = \mathbf{x}_i, \boldsymbol{\sigma}_{\mathbf{U}}^2) \\ &= \prod_{j=1}^p \mathbb{P}(U_i^j \in R_k^j | X_i^j = x_i^j, \sigma_{U_j}^2). \end{aligned}$$

¹ P depends on the number of regions considered. It is thus a dynamic matrix that evolves during the construction of the tree. For clarity sake, we do not explicit this dependence in our notation but the reader has to keep this in mind.

Estimating γ – It is easy to see that when fixing the regions $(R_k)_{1 \leq k \leq K}$ and the vector of variances σ_U^2 , minimizing Eq. (2.5) with respect to γ corresponds to a weighted average of y_1, \dots, y_n , in a way similar to the linear regression model if $P^T P$ is not singular:

$$(2.6) \quad \hat{\gamma} = \underset{\gamma \in \mathbb{R}^K}{\operatorname{argmin}} \left\{ R_{emp}(T_{un}(\cdot; \hat{\Theta}), \mathcal{L}_n) \right\} \\ = (P^T P)^{-1} P^T \mathbf{y},$$

where \mathbf{y} denotes the vector of n univariate outputs y_1, \dots, y_n . Indeed, by definition:

$$R_{emp}(T_{un}(\cdot; \Theta), \mathcal{L}_n) = \frac{1}{n} \sum_{i=1}^n \left(y_i - \sum_{k=1}^K \gamma_k P_{i,k} \right)^2.$$

Differentiating wrt γ_k , for all $1 \leq k \leq K$, one gets:

$$\frac{\partial}{\partial \gamma_k} R_{emp}(T_{un}(\cdot; \Theta), \mathcal{L}_n) = 0 \\ \Leftrightarrow 2 \sum_{i=1}^n \left[P_{i,k} \cdot \left(y_i - \sum_{k'=1}^K P_{i,k'} \gamma_{k'} \right) \right] = 0 \\ \Leftrightarrow \sum_{i=1}^n P_{i,k} \cdot y_i = \sum_{i=1}^n \sum_{k'=1}^K P_{i,k} P_{i,k'} \gamma_{k'}.$$

So that, if $P^T P$ is not singular:

$$P^T \mathbf{y} = P^T P \gamma \Leftrightarrow \gamma = (P^T P)^{-1} P^T \mathbf{y},$$

which is a minimum. In Section 2.3, we derive assumptions under which $P^T P$ is indeed not singular. In practice, one can always use the pseudo-inverse of $P^T P$, which we will do in our experiments. Lastly, note that $\hat{\gamma}$ depends on the regions through the matrix P .

Estimating $(R_k)_{1 \leq k \leq K}$ – The regions are constructed in a way similar to the construction process of standard trees with the difference that here, during the construction process of the regression tree, the number of regions K is not fixed and increases step by step, which implies that the size of the matrix P is also changing during the iterative process. Let us assume that K regions have already been identified, meaning that the current tree has K leaves, each leaf corresponding to a region (*i.e.*, hyper-rectangle). As in standard regression trees, each *current* region R_k , $1 \leq k \leq K$, can be decomposed into two sub-regions wrt a variable X^j and a splitting point s_k^j , for $1 \leq j \leq p$:

$$\begin{cases} R_{k,L}^j = \{ \mathbf{x} \in R_k \mid X^j < s_k^j \} \\ R_{k,R}^j = \{ \mathbf{x} \in R_k \mid X^j > s_k^j \}. \end{cases}$$

This decomposition adds a new region for which the associated elements, $P \in M_{n, K+1}(\mathbb{R})$ and $\gamma \in \mathbb{R}^{K+1}$

can readily be computed. P has now $K+1$ regions corresponding to the current regions $R_{k'}$ (with $k' \neq k$) and the two new regions $R_{k,L}^j$ and $R_{k,R}^j$. For each current region R_k , one is looking for the best split, *i.e.*, the best variable X^j and the best splitting point s_k^j that minimize:

$$\sum_{i=1}^n \left(y_i - \sum_{l=1}^{K+1} \gamma_l P_{i,l} \right)^2.$$

The sum includes all possible observations as each observation has a non null probability to belong to any region. Using Eq. (2.6), the above problem can be reformulated, for each current region R_k , as:

$$(2.7) \quad \underset{1 \leq j \leq p, s \in \mathcal{S}_k^j}{\operatorname{argmin}} \left\{ \sum_{i=1}^n \left(y_i - \sum_{l=1}^{K+1} \left((P^T P)^{-1} P^T \mathbf{y} \right)_l P_{i,l} \right)^2 \right\},$$

where \mathcal{S}_k^j denotes the set of splitting points corresponding to the middle points of the j^{th} coordinates of the observations in R_k sorted according to X^j .

The decomposition corresponding to the best split is then used to build the child nodes of R_k , which is replaced, in the set of current regions, by its two children. In this process, that is repeated till a stopping criterion is met², the number of regions, the matrix P and the weights γ are gradually updated. Section 3 provides the algorithm corresponding to this construction.

Estimating σ_U^2 – Lastly, the vector of variances, σ_U^2 , can either be set according to some high level principles, or can be learned through a grid search on a validation set guaranteeing. The latter is more demanding wrt computational resources, but is likely to lead to better results. However, in our experiments, we use the former strategy, with the aim of showing that our approach is robust in the sense that it yields good results even when σ_U is set *a priori*.

2.2 Final prediction For a new observation $\mathbf{x} \in \mathbb{R}^p$, one first computes its distribution over all the obtained regions:

$$\forall k, 1 \leq k \leq K, P_k^{\mathbf{x}} = \mathbb{P}(U \in R_k \mid \mathbf{X} = \mathbf{x}, \sigma_U^2).$$

The prediction is then a direct application of (2.3):

$$(2.8) \quad T_{un}(\mathbf{x}; \hat{\Theta}) = \sum_{k=1}^K \hat{\gamma}_k P_k^{\mathbf{x}},$$

where γ , $(R_k)_{1 \leq k \leq K}$ and σ_U^2 are estimated as described above.

²Any standard stopping criterion can be used here.

2.3 A sufficient condition on the invertibility of $P^T P$ The formula used for the construction of the tree relies on the inverse on the matrix $P^T P \in M_{K,K}(\mathbb{R})$. Even if numerically the Moore-Penrose pseudo-inverse might be used to approximate this inverse, we provide here a sufficient condition on the invertibility of P . Without loss of generality, the regions involved in the definition of the regression tree are of the form $R_k = \prod_{j=1}^p [a_k^j, b_k^j]$ for $1 \leq k \leq K$, where $(a_k^j, b_k^j) \in (-\infty, +\infty)^2$. As usual, we denote q_α the α quantile of the standard Gaussian distribution. The general invertibility result is stated in Theorem 2.1.

THEOREM 2.1. *With the model defined in (2.2), if the following assumption is satisfied:*

$$\forall j, 1 \leq j \leq p, \sigma_{U_j} < \frac{\min_{1 \leq k \leq K} (b_k^j - a_k^j)}{2q_{\frac{1+0.5}{2}^{1/p}}},$$

then the matrix $P^T P$ is invertible.

Roughly speaking, the matrix $P^T P$ is invertible provided that the standard deviation σ_{U_j} is sufficiently small for all $1 \leq j \leq p$. The smaller the regions and/or the larger the number p of input variables, the lower the uncertainty on the measurement has to be to ensure the invertibility of the matrix $P^T P$.

To prove Theorem 2.1, we prove that P is of full rank. To do so, we first prove the following result:

PROPOSITION 2.1. *Let us fix $1 \leq k \leq K$ and consider $1 \leq i \leq n$ such that $a_k^j < x_i^j < b_k^j$. Assume that*

$$(2.9) \quad \forall 1 \leq j \leq p, \sigma_{U_j} < \frac{b_k^j - a_k^j}{2q_{\frac{1+0.5}{2}^{1/p}}}.$$

Then, $P_{i,k} > 0.5$.

Proof. Observe that a sufficient condition to have $P_{i,k} > 0.5$ is

$$\forall j, \frac{1}{\sigma_{U_j} \sqrt{2\pi}} \int_{R_k^j} e^{-\frac{(u-x_i^j)^2}{2\sigma_{U_j}^2}} du > 0.5^{1/p}.$$

We now search a sufficient condition to have the inequality just above. To get this inequality the following condition is sufficient: for all $1 \leq j \leq p$,

$$(2.10) \quad \mathbb{P} \left(\frac{U_i^j - x_i^j}{\sigma_{U_j}} < \frac{b_k^j - x_i^j}{\sigma_{U_j}} \mid X_i = x_i \right) - \mathbb{P} \left(\frac{U_i^j - x_i^j}{\sigma_{U_j}} < \frac{a_k^j - x_i^j}{\sigma_{U_j}} \mid X_i = x_i \right) > 0.5^{1/p}.$$

Note that this last condition is satisfied if we have, for all $1 \leq j \leq p$,

$$\begin{cases} \mathbb{P} \left(\frac{U_i^j - x_i^j}{\sigma_{U_j}} < \frac{b_k^j - x_i^j}{\sigma_{U_j}} \mid X_i = x_i \right) > \frac{1+0.5^{1/p}}{2}, \\ \mathbb{P} \left(\frac{U_i^j - x_i^j}{\sigma_{U_j}} < \frac{a_k^j - x_i^j}{\sigma_{U_j}} \mid X_i = x_i \right) < \frac{1-0.5^{1/p}}{2}, \end{cases}$$

which is equivalent to

$$\frac{a_k^j - x_i^j}{\sigma_{U_j}} < q_{\frac{1-0.5^{1/p}}{2}} \quad \text{and} \quad \frac{b_k^j - x_i^j}{\sigma_{U_j}} > q_{\frac{1+0.5^{1/p}}{2}}.$$

Note that $q_{\frac{1-0.5^{1/p}}{2}} = -q_{\frac{1+0.5^{1/p}}{2}} < 0$. A sufficient condition is then, for all $1 \leq j \leq p$,

$$\sigma_{U_j} < \min_{1 \leq k \leq K} \left\{ \frac{\min(x_i^j - a_k^j, b_k^j - x_i^j)}{q_{\frac{1+0.5^{1/p}}{2}}} \right\}.$$

Since $a_k^j < x_i^j < b_k^j$, a condition even more conservative is the following:

$$\forall j, \sigma_{U_j} < \frac{b_k^j - a_k^j}{2q_{\frac{1+0.5^{1/p}}{2}}}.$$

This concludes the proof. ■

We assume in the following that Assumption (2.9) is satisfied. Then, the set $I_k = \{1 \leq i \leq n, P_{i,k} > 0.5\}$ is non-empty. Let us consider, for all $1 \leq k \leq K$, i_k a representative of this set and let us introduce the matrix $Q \in M_{K,K}(\mathbb{R})$ defined, for all $1 \leq l, k \leq K$, by:

$$Q_{l,k} = P_{i_l,k}.$$

PROPOSITION 2.2. *Assume that Assumption (2.9) holds. Then the $K \times K$ matrix Q is invertible.*

Proof. We first show that Q is a strictly dominant diagonal matrix, i.e.:

$$\forall k, Q_{kk} > \sum_{k' \neq k} Q_{kk'}.$$

Indeed, by Proposition 2.1, we know that $P_{i_k,k} > 0.5$ and k is the only region where it is true:

$$\sum_{l \neq k} Q_{l,k} \leq \sum_{l \neq k} P_{i_l,k} = 1 - P_{i_k,k} < 0.5 < P_{i_k,k} = Q_{k,k}.$$

According to Hadamard's Lemma, we know that a strictly dominant diagonal matrix is invertible, which concludes the proof. ■

As Q is invertible, P is of full rank, leading to the fact that $P^T P$ is invertible which concludes the proof of Theorem 2.1.

2.4 Extension to Random Forests It is straightforward to use the uncertain regression trees introduced above in Random Forests, leading to *Uncertain Random Forests*. Indeed, each tree of the forest is now an uncertain regression tree that can be constructed as outlined before. Assuming a forest of τ uncertain trees and denoting $\hat{\gamma}^t$, $1 \leq t \leq \tau$, the weights estimated for each tree and $P_{k_t}^{\mathbf{x}, t}$ the probability distribution of a new observation $\mathbf{x} \in \mathbb{R}^p$ over the K_t regions ($1 \leq k_t \leq K_t$) of the t^{th} tree ($1 \leq t \leq \tau$) of the forest, the prediction rule of the uncertain random forest takes the form:

$$(2.11) \quad RF_{\text{un}}(\mathbf{x}; \hat{\Theta}_{t, 1 \leq t \leq \tau}) = \frac{1}{\tau} \sum_{t=1}^{\tau} \sum_{k_t=1}^{K_t} \hat{\gamma}_{k_t}^t P_{k_t}^{\mathbf{x}, t}.$$

As one can note, the above prediction rule is a direct extension of Eq. (2.8).

3 Associated algorithms

Algorithm 3.1 describes the construction of uncertain regression trees. This construction parallels the one of standard regression trees except that we consider a matrix P encoding the probability distribution of training data points over regions, which is built dynamically (as the regions and the weights) by adding a new column when a given region is split into two sub-regions.

Each time a region R , corresponding to a current leaf of the tree being constructed, is considered (through the pop function in Algorithm 3.1), one identifies the best split (j^*, s^*) that minimizes the empirical risk in Eq. (2.7) among all possible splitting points in $\mathcal{S}_j(R)$ of each variable j . The set $\mathcal{S}_j(R)$ is defined by $\mathcal{S}_j(R) = \{(x_{i_{l+1}}^j + x_{i_l}^j)/2 \mid 1 \leq l \leq r-1\}$, with $(\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_r})$ corresponding to the r observations of the learning set \mathcal{L}_n belonging to the region R , sorted such that $x_{i_1}^j \leq x_{i_2}^j \leq \dots \leq x_{i_r}^j$.

The k -th column of P corresponding to the current region is finally replaced by the probability distribution of the training data points over its *left sub-region* (denoted R_L) and an additional column is added to P for the *right sub-region* (denoted R_R). The weights $\hat{\gamma}$ are easily deduced from P at each step through (2.6).

The algorithm finally outputs the set of regions $\mathcal{R}(S)$ corresponding to the leaves of the tree and the weights $\hat{\gamma}$.

ALGORITHM 3.1. *Uncertain regression trees*

Input: \mathcal{L}_n
Initialization: $K = 1$, $S \leftarrow (1, \mathcal{X})$, $P = \mathbf{1}_n$
while stopping criterion not met
 $(k, R) = S.\text{pop}()$
 for $j = 1$ to p
 Construct $\mathcal{S}_j(R)$

for $s \in \mathcal{S}_j(R)$

$R_L = \{\mathbf{x} \in R, x_j \leq s\}$
 $R_R = \{\mathbf{x} \in R, x_j \geq s\}$
 $Q_L = (\mathbb{P}(U_i \in R_L | X_i = x_i))_{1 \leq i \leq n}$
 $Q_R = (\mathbb{P}(U_i \in R_R | X_i = x_i))_{1 \leq i \leq n}$
 $P^{[j, s]} \leftarrow \text{Merge} \{P[1 : k-1], Q_L,$
 $P[k : k+1 : K], Q_R\}$

end for

end for
 $(j^*, s^*) = \underset{1 \leq j \leq p, s \in \mathcal{S}_j(R)}{\text{argmin}} \text{risk}(P^{[j, s]})$ #defined in (2.7)

$P = P^{[j^*, s^]}$

Update $\hat{\gamma}$ acc. to Eq. 2.6

$K = K+1$

$S.\text{append}((k, R_L), (K, R_R))$

end while

Output: $(\mathcal{R}(S), \hat{\gamma})$

Note that in this version the tree is constructed in a depth-first manner, and that the usual stopping criteria for regression trees can be used (as imposing a minimum number of observations in each leaf or a maximal depth for the tree).

Algorithm 3.2 provides the pseudo-code for the prediction rule given in (2.8) for a new covariate \mathbf{x}^{pred} .

ALGORITHM 3.2. *Prediction*

Input: $(R_k)_{1 \leq k \leq K}$, $\hat{\gamma}$, \mathbf{x}^{pred} , σ_U^2

$\hat{y}^{\text{pred}} = 0$

for $k = 1$ to K

$\hat{y}^{\text{pred}+} = \mathbb{P}(U \in R_k | \mathbf{X} = \mathbf{x}^{\text{pred}}, \sigma_U^2) \hat{\gamma}_k$

end for

Output: \hat{y}^{pred}

4 Experimental validation

The goal of our experiments is to assess how uncertain regression trees:

- compare to standard regression trees,
- behave in *Uncertain Random Forests*.

By standard regression trees we mean here regression trees based on the quadratic risk and the prediction rule given in Eq. (2.1). In addition, we consider a trade-off between standard and uncertain regression trees based on standard trees (and thus avoiding the complex construction process outlined before) but using the prediction rule given in (2.8). The matrix P and the weights γ are thus computed only once, when the standard trees have been built. The rationale for this method is to rely on the strengths of the two approaches: a simple construction tree process and a robust prediction rule. As one can conjecture, this method yields results in between the two approaches.

We make use in our experiments of four benchmark data sets commonly used in regression tasks and described in Section 4.1. We first use these data sets without any modification, to illustrate the fact that *real* data sets contain uncertain inputs. The uncertain regression trees proposed here indeed outperform standard trees and Random Forests on these data sets, as shown in Section 4.2. We then modify these data sets by adding a uniform noise bounded by a quantity proportional to the empirical variance of the data. This additional perturbation aims at assessing the robustness of the different methods (standard and uncertain trees) in situations where the input data is highly uncertain. Once again, the results obtained show the good behavior of the uncertain trees and Random Forests (Section 4.3). In all our experiments, the results are evaluated with the Root Mean Squared Errors (RMSE), which is a standard evaluation measure in regressions tasks. To ensure that all the available data is used for both training and testing, we further rely on 5-fold cross-validation and report the mean RMSE and its standard deviation over the 5 folds.

The stopping criterion for the trees, both standard and uncertain, is based on the fact that all leaves should contain at least ten percent of the training data. For Random Forests, three features are randomly selected for constructing a tree, which roughly corresponds to one third of the features on the data sets considered, a standard proportion in Random Forests.

In this study, the vector of variances is fixed according to some high level principle. In particular, when no additional noise is introduced on the data, the variance for a particular variable $U^j|X^j$ is set to the empirical variance of X^j (in other words, we assume that the variance of the *true* values is of the same order as the variance of the observed values). When some noise is added to the data, the variance of $U^j|X^j$ is set to one half of the variance of the observed, noisy data (in this case, the variance of the true values should be lower than the empirical variance of the observed values; we arbitrarily chose one half in this study).

Lastly, our algorithms have been implemented from `scikit-learn` [21], using Cython [1].

4.1 Data Sets Experiments are conducted on 4 publicly available data sets, popular in the regression tree literature. The main characteristics of these data sets are summarized in Table 1. As we focus in this paper on quantitative variables, we simply deleted the categorical variables from the data sets. Several applications are considered, among which environment (concentration in ozone over a day for the data set *Ozone* introduced in [5]), health (data set *Diabete*, introduced in [21] and used in [8]), economy (data set *BigMac* about

price of sandwiches, available in R package `alr3` and used in [16]) or biology (data set *Abalone* introduced in [18] and used in [16] among others).

Data set	sample size n	number of variables p
BigMac	69	9
Ozone	112	9
Diabete	442	10
Abalone	500	7

Table 1: Characteristics of data sets used in our experiments, ordered by sample size n .

4.2 Results on benchmark data sets As mentioned in the introduction, data sets are by nature uncertain, so we illustrate our method on the four data sets introduced in Section 4.1. The noise reflects the uncertainty, so we propose to use the empirical standard deviation vector as the input parameter σ_U . Table 2 displays the results obtained for each data set in terms of mean and standard deviation of RMSE using 5-fold cross validation. Standard trees, standard Random Forests (with $\tau = 100$ trees) as well as uncertain trees and standard trees with uncertain prediction are compared.

For all data sets, the best performances are achieved by uncertain trees. As expected, the standard Random Forest are performing better than considering only one tree. However, as one can see in Table 2, uncertain trees yield better results than standard Random Forests. Performances of standard trees with uncertain prediction are better than the ones of standard trees, however not always better than the ones of standard Random Forests.

Finally, one can note that the standard deviation of uncertain trees is smaller than the standard deviation of other methods, meaning that uncertain trees yield more stable results.

4.3 Results on artificial uncertain data sets To assess the robustness of uncertain trees to and uncertain Random Forests to uncertainty in the input variable, we add artificial noise, which could correspond to some measure error, to the input observations. By doing so, one can consider that the variability of the data is coming from two sources, on the one hand the variance in the latent variables (related to the variance of \mathbf{U} with the notations of Section 2) and on the other hand the variance of the uncertainty (related to the variance of $\mathbf{X}|\mathbf{U}$), both leading to the variance of the observations. Then, we assume here that σ_U is smaller than the standard deviation of the observations σ_X , which can be estimated through the data set. Basically it means that

Data sets	BigMac	Ozone	Diabete	Abalone
Standard tree	25.09 (12.3)	17.82 (4.3)	60.29 (4.3)	2.70 (0.3)
Standard RF, $\tau = 100$	19.78 (11.0)	15.86 (4.1)	58.18 (4.3)	2.65 (0.4)
Standard tree with uncertain prediction	21.49 (8.7)	16.79 (4.1)	57.05 (3.7)	2.41 (0.2)
Uncertain tree	18.74 (8.9)	15.39 (3.7)	56.56 (3.3)	2.33 (0.3)

Table 2: Average RMSE based on 5-fold cross-validation for the 4 benchmark data sets. Standard deviations are given in parentheses. For each uncertain tree based method, σ_U are set to the empirical standard deviations of the observed input variables.

the main part of the variance is due to the uncertainty.

To construct artificial uncertain data sets satisfying this condition, we consider in this section *BigMac*, *Ozone*, *Diabete* and *Abalone* data sets introduced in Section 4.1. A noise is added to each observation using the following process. For each observation from an input variable X^j , $1 \leq j \leq p$, we add a noise generated from the product of a Rademacher variable and a uniform variable coming from the interval $[\hat{\sigma}_{X_j}/10, \hat{\sigma}_{X_j}/4]$.

We consider here standard trees, standard Random Forests with $\tau = 500$ trees, standard trees with uncertain prediction, uncertain trees and uncertain Random Forests with $\tau = 15$ trees. The results obtained are displayed in Table 3.

As one can note, uncertain trees outperforms here again standard trees. RMSE scores are a bit higher than in Table 2, which is not surprising given the noise added to the data sets.

Similarly, uncertain Random Forests outperform all the other methods, including standard Random Forests even though the number of trees is one order of magnitude less.

5 Discussion

Regression trees have been introduced in the 1980s through the popular CART algorithm [3], notably allowing one to deal with both categorical and numerical input variables. They constitute the basic building block of state-of-the-art ensemble methods based on the combination of random trees, notably Random Forests introduced by Breiman in [4] to circumvent the instability of Regression trees [12]. Since Random Forests are particularly well suited for Big Data analysis (see [7]), many applications have been addressed in various domains with Random Forests, for example in ecology or genomics. [22] provides a review of the use of Random Forests for data mining purposes. An interesting feature of Random Forests is the possibility to quantify the importance of input variables in the model obtained (see [6] for more details on that point). Another interesting feature, which was empirically established,

is their robustness to noise. They are thus, to a certain extent, robust to uncertain inputs (even though no mechanism was specifically designed for that). This said, explicitly modelling the uncertainty as done in the uncertain regression trees proposed here allows one to outperform the standard version of Random Forests, as illustrated in our experiments.

Several adaptations of ensemble methods for quantile regression have been proposed, as quantile Random Forests or quantile boosting trees [9, 14, 15, 16, 23]. These studies however focus on the uncertainty in the output variable (by producing conditional quantiles rather than a conditional mean) and not on the uncertainties in the inputs, as done here. It is of course possible to combine both approaches, which we intend to do in the future.

Lastly, the idea of including information on the uncertainty of the input data in a regression method is not new and is related to uncertainty propagation and sensitivity analysis. Several authors have indeed proposed to integrate uncertainties on the inputs in different regression methods, as multivariate linear models [17] or neural networks [11]. In each case, the methods have been improved, showing the benefits of explicitly modelling uncertainties in the input data. To the best of our knowledge, our approach is the first one to address this problem in the context of regression trees (and ensemble methods based on such trees). Our conclusion on the benefits of this approach parallels the ones of the above-mentioned studies.

6 Conclusion

We have introduced in this study uncertain regression trees which can deal with uncertain inputs. In such trees, observations no longer belong to a single region, but rather have a non-null probability to be assigned to any region of the tree. This extra flexibility leads nevertheless to a construction process that is more complex than the one underlying standard regression trees and that necessitates the inversion of a square matrix (for which we have theoretically provided a sufficient condition). In practice, we rely on the pseudo-inverse of

Uncertain data sets	BigMac	Ozone	Diabete	Abalone
Standard tree	22.28 (8.7)	19.37 (4.1)	60.47 (2.81)	2.54 (0.17)
Standard tree with uncertain prediction	21.68 (9.8)	17.35 (4.9)	57.92 (3.43)	2.40 (0.15)
Uncertain tree	19.28 (13.4)	16.87 (6.0)	58.56 (3.45)	2.34 (0.20)
Standard RF, $\tau = 500$	19.25 (7.8)	15.72 (3.0)	59.55 (4.39)	2.64 (0.18)
Uncertain RF, $\tau = 15$	18.06 (9.3)	15.49 (3.7)	55.66 (4.31)	1.98 (0.12)

Table 3: Average RMSE based on 5-fold cross-validation for the 4 modified data sets. Standard deviations are given in parentheses. On each data set, each observation is modified by a noise generated from the product of a Rademacher variable and a uniform variable coming from the interval $[\frac{\hat{\sigma}_{x_j}}{10}, \frac{\hat{\sigma}_{x_j}}{4}]$. For each uncertain tree-based method, σ_U is to half of the empirical standard deviations of the observed (modified) input variables.

this matrix. The experimental results fully justify the approach we have proposed and show that uncertain regression trees improve the results of standard regression trees and Random Forests on several benchmark data sets. A similar conclusion is drawn on artificial uncertain data sets obtained from the standard ones by introducing additional uncertainty in the form of a uniform noise.

The methodology developed in this study can also be adapted to the case where some input data are categorical. We plan to work on such an adaptation in the future, by considering different types of uncertainties.

Furthermore, as mentioned before, the vector of variances for modelling uncertainties can easily be learned by grid search on validation sets (typically using cross-validation). One can expect by doing so that the results would further improve. We have set this vector in our experiments to values that we believe are reasonable, so as to show that our approach is robust in the sense that it yields good results even when σ_U is set *a priori*. We nevertheless plan to run additional experiments to learn this vector. As a grid search can be easily parallelized, this learning should not impact the running time of the algorithm.

Lastly, exploring the extension of our method to Random Forests or boosting trees, as in [10, 20], is another promising research direction. We also intend to explore the use of other empirical loss functions, as the quantile loss used in the definition of quantile Random Forests or quantile boosting trees [9, 14, 15, 16].

References

- [1] S. Behnel, R. Bradshaw, C. Citro, L. Dalcin, D. S. Seljebotn and K. Smith, *Cython: the best of both worlds*, CiSE (2011) 13(2), pp. 31-3.
- [2] D.A. Belsley, E. Kuh and R.E. Welsch, *Regression diagnostics: Identifying Influential Data and Sources of Colinearity*, Wiley (1980), pp. 244-261.
- [3] L. Breiman, J. Friedman, C. J. Stone and R. A. Olshen, *Classification and regression trees*, CRC press (1984).
- [4] L. Breiman, *Random forests*, Machine learning 45(1)(2001), pp. 5–32.
- [5] P.A. Cornillon, A. Guyader, F. Husson, N. Jegou, J. Josse, M. Kloareg, E. Matzner-Lober, and L. Rouvière, *R for Statistics*, CRC Press (2012).
- [6] R. Genuer, J.M. Poggi, and C. Tuleau-Malot. *Variable selection using random forests*. Pattern Recognition Letters 31, no. 14 (2010), pp. 2225-2236.
- [7] R. Genuer, J.-M. Poggi, C. Tuleau-Malot, N. Villa-Vialaneix, *Random Forests for Big Data*. Big Data Research, 9(2017), pp. 28-46.
- [8] B. Efron, T. Hastie, I. Johnstone and R. Tibshirani, *Least angle regression*, Ann. Statist. 32,2 (2004), pp. 407–499.
- [9] N. Fenske, T. Kneib and T. Hothorn, *Identifying risk factors for severe childhood malnutrition by boosting additive quantile regression*, Journal of the American Statistical Association, 106(494) (2011), pp.494–510.
- [10] Y. Freund, R. Schapire and N. Abe, *A short introduction to boosting*, Journal-Japanese Society For Artificial Intelligence 14(5) (1999), pp.771–780.
- [11] Y. Gal, and G. Zoubin, *Dropout as a Bayesian approximation: Representing model uncertainty in deep learning* In international conference on machine learning (2016), pp. 1050-1059.
- [12] S. Gey and J.M. Poggi, *Boosting and instability for regression trees*. Computational statistics & data analysis, (2006) 50(2) pp 533–550.
- [13] T. Hastie, R. Tibshirani and J. Friedman, *The Elements of Statistical Learning*, Springer Series in Statistics (2009).
- [14] B. Krieglner and R. Berk, *Boosting the quantile distribution: A cost-sensitive statistical learning procedure*, Preprint (2007).
- [15] ———, *Small area estimation of the homeless in Los Angeles: An application of cost-sensitive stochastic gradient boosting*, The Annals of Applied Statistics, 4(3)(2010), pp.1234–1255.
- [16] N. Meinshausen, *Quantile regression forests*, Journal of Machine Learning Research, 7(2006), pp. 983–999.
- [17] M.S. Reis, and P. M. Saraiva. *Integration of data uncertainty in linear regression and process optimization*. AIChE journal 51.11 (2005), pp. 3007-3019.
- [18] W.J. Nash, T. L. Sellers, S. R. Talbot, A. J. Cawthorn

and B. B. Ford *The Population Biology of Abalone (*Haliotis* species) in Tasmania. I. Blacklip Abalone (*H. rubra*) from the North Coast and Islands of Bass Strait*, Sea Fisheries Division, Technical Report No. 48 (ISSN 1034-3288) (1994).

- [19] R. Quinlan, *Combining Instance-Based and Model-Based Learning*, In Proceedings on the Tenth International Conference of Machine Learning (1993), pp. 236-243, University of Massachusetts, Amherst. Morgan Kaufmann.
- [20] G. Ridgeway, *Generalized Boosted Models: A guide to the gbm package*, (2007).
- [21] G. Varoquaux, L. Buitinck, G. Louppe, O. Grisel, F. Pedregosa, and A. Mueller, *Scikit-learn: Machine Learning Without Learning the Machinery*, GetMobile: Mobile Comp. and Comm. 19, 1 (2015), pp. 29-33.
- [22] A. Verikas, A. Gelzinis, and M. Bacauskiene, *Mining data with random forests : A survey and results of new tests*. Pattern Recognition, 44(2)(2011), pp. 330-349.
- [23] S. Zheng *QBoost: Predicting quantiles with boosting for regression and binary classification*. Expert Systems with Applications. (2012) Feb 1;39(2), pp. 1687-97.