



HAL
open science

A REVIEW ON QUANTILE REGRESSION FOR STOCHASTIC COMPUTER EXPERIMENTS

Léonard Torossian, Victor Picheny, Robert Faivre, Aurélien Garivier

► **To cite this version:**

Léonard Torossian, Victor Picheny, Robert Faivre, Aurélien Garivier. A REVIEW ON QUANTILE REGRESSION FOR STOCHASTIC COMPUTER EXPERIMENTS. 2019. hal-02006032v2

HAL Id: hal-02006032

<https://hal.science/hal-02006032v2>

Preprint submitted on 30 Oct 2019 (v2), last revised 20 Jan 2020 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A REVIEW ON QUANTILE REGRESSION FOR STOCHASTIC COMPUTER EXPERIMENTS

Léonard Torossian

MIAT, Université de Toulouse, INRA
and Institut de Mathématiques de Toulouse
leonard.torossian@inra.fr

Victor Picheny

PROWLER.io, 72 Hills Road, Cambridge
victor@prowler.io

Robert Faivre

MIAT, Université de Toulouse, INRA
robert.faivre@inra.fr

Aurélien Garivier

Univ. Lyon, ENS de Lyon
aurelien.garivier@ens-lyon.fr

October 30, 2019

ABSTRACT

We report on an empirical study of the main strategies for quantile regression in the context of stochastic computer experiments. To ensure adequate diversity, six metamodels are presented, divided into three categories based on order statistics, functional approaches, and those of Bayesian inspiration. The metamodels are tested on several problems characterized by the size of the training set, the input dimension, the signal-to-noise ratio and the value of the probability density function at the targeted quantile. The metamodels studied reveal good contrasts in our set of experiments, enabling several patterns to be extracted. Based on our results, guidelines are proposed to allow users to select the best method for a given problem.

1 Introduction

1.1 Stochastic experiment setting

Computer simulation models are now essential for performance evaluation, quality control and uncertainty quantification to assess decisions in complex systems. These computer simulators generally model systems depending on multiple input variables that can be divided into two categories: the controllable variables and the uncontrollable variables.

For example, in pharmacology, the optimal drug dosage depends on the drug formulation but also on the targeted individual (genetics, age, sex) and environmental interactions. The shelf life and performance of a manufacturing device depend on its design, but also on its environment and on some uncertainties during to the manufacturing process. The plant growth and yield depend on the genes of the plant and on the gardening techniques but also on the weather and potential diseases.

Evaluating the influence of the controllable and uncontrollable variables directly on the real-life problems can be costly and tedious. One solution is to encapsulate the systems into a computer simulation model which would reduce the cost and the time required for each test (see [42] and reference therein for computer experiments applied to clinical trials, see [36] for a computer simulation model applied to industrial design and [95, 19] for computer experiments applied to crop production).

In such computer simulation models, the links between the inputs and outputs may be too complex to be fully understood or to be formulated in a closed form. In this case, the system can be considered as a black box and formalized by an unknown function: $f : \mathcal{X} \times \Omega \rightarrow \mathbb{R}$, where $\mathcal{X} \subset \mathbb{R}^d$ denotes the compact space of controllable variables, and Ω denotes a probability space representing the uncontrollable variables. Note that some black boxes have their outputs in a multidimensional space but this aspect is beyond the scope of the present paper.

Based on the standard constraints encountered in computer experiments, throughout this paper, we assume that the function is only accessible through pointwise evaluations $f(x, \omega)$; no structural information is available regarding f ; in addition, we take into account the fact that evaluations may be expensive, which drastically limits the number of possible calls to f .

If the space Ω is small enough (or highly structured), it may be possible to work directly on the space $\mathcal{X} \times \Omega$ (see [44] for example). Often, however, Ω is too complex and working on the joint space is intractable. This is the case for intrinsic stochastic simulators (see [50, 56] for examples of biological systems, where the stochasticity is driven by stochastic equations such as the Fokker-Planck or the chemical Langevin equations), or for simulators associated with a very large space Ω (see [19] for the crop model SUNFLO that considers 5 weather indicators on 130 days, *i.e.* Ω is a space of dimension 650).

In this paper we consider the case where Ω is too complex. We assume that Ω has any structure and its contribution is considered as random. In contrast to deterministic systems, for any fixed x , $f(x, \cdot)$ is considered as a random variable of distribution \mathbb{P}_x ; hence, such systems are often referred to as stochastic black boxes.

In order to understand the behavior of the system of interest or to take optimal decisions, information is needed about \mathbb{P}_x . An intuitive approach is to use a simple Monte-Carlo technique and evaluate $f(x, \omega_1), \dots, f(x, \omega_n)$ to extract statistical moments, the empirical cumulative distribution function, etc. Unfortunately, such a stratified approach is not efficient when evaluating f is expensive.

Instead, we focus on *surrogate models* (also referred to as *metamodels* or *statistical emulators*), which are appropriate approaches in a small data setting associated with a regularity hypothesis (with respect to \mathcal{X}) concerning the targeted statistics. Among the vast choice of surrogate models [88, 98], the most popular ones include regression trees, Gaussian processes, support vector machines and neural networks. In the framework of stochastic black boxes, the standard approach consists in estimating the conditional expectation of f . This case has been extensively treated in the literature and many applications, including Bayesian optimization [79], have been developed. However, the conditional expectation is risk-neutral, whereas pharmacologists, manufacturers, asset managers, data scientists and agronomists may need to evaluate the worst case scenarios associated with their decisions.

Risk information can be introduced by using a surrogate expectation-variance model in which the distribution can be estimated by non-parametric kernels (see [40, 80] for instance) or via heteroscedastic Gaussian processes [45, 49]. However, such approaches usually imply that the shape of the distribution (e.g. normal, uniform, etc.) is the same for all $x \in \mathcal{X}$. Another possible approach would be to learn the whole distribution \mathbb{P}_x with no strong structural hypotheses [63, 38, 32], but this requires a large number of evaluations of f . Here, we focus on the conditional quantile estimation of order τ , a flexible way to tackle cases in which the distribution of $f(x, \cdot)$ varies markedly in spread and shape with respect to $x \in \mathcal{X}$, and a classical risk-aware tool in decision theory [73].

1.2 Paper Overview

Many metamodels originally designed to estimate conditional expectations have been adapted to estimate the conditional quantile. However, despite extensive literature on estimating the quantile in the presence of spatial structure, few studies have reported on the constraints associated with stochastic black boxes. The performance of a metamodel with high dimension input is treated in insufficient details, performance based on the number of points has rarely been tackled and, to our knowledge, dependence on specific aspects of the quantile functions has never been studied. The aim of the present paper is to review quantile regression methods under standard constraints related to the stochastic black box framework, so as to provide information on the performance of the selected methods, and to recommend which metamodel to use depending on the characteristics of the computer simulation model and the data.

A comprehensive review of quantile regression is of course beyond the scope of the present work. We limit our review to the approaches that are best suited for our framework, while ensuring the necessary diversity of metamodels. In particular, we have chosen six metamodels that are representative of three main categories: approaches based on statistical order (K-nearest neighbors [KN] regression [7] and random forest [RF] regression [60]), functional or frequentist approaches (neural networks [NN] regression [18] and regression in reproducing kernel Hilbert space [RK] [90]), and Bayesian approaches based on Gaussian processes (Quantile Kriging [QK] [66] and the variational Bayesian [VB] regression [1]). Each category has some specificities in terms of theoretical basis, implementation and complexity. We begin this presentation by describing the methods in full in sections 3, 4 and 5.

In order to identify the relevant areas of expertise of the different metamodels, an original benchmark system is designed based on four toy functions and an agronomical model [19]. The dimension of the problems ranges from 1 to 9 and the number of observations from 40 to 2000. Particular attention is paid to the performance of each metamodel according to the size of the learning set, the value of the probability density function at the targeted quantile $\tilde{f}(\cdot, q_\tau)$

and the dimension of the problem. Sections 6 and 7 describe the benchmark system and detail its implementation, with particular focus on the tuning of the hyperparameters of each method. Full results and discussion are to be found in Sections 8 and 9, respectively.

2 Quantile emulators and design of experiments

We first provide the necessary definitions, objects and properties related to the quantile. The quantile of order $\tau \in (0, 1)$ of a random variable Y can be defined either as the (generalized) inverse of a cumulative distribution function (CDF), or as the solution to an optimization problem:

$$q_\tau = \min \{q \in \mathbb{R} : F(q) \geq \tau\} = \arg \min_{q \in \mathbb{R}} \mathbb{E}[l_\tau(Y - q)], \quad (1)$$

$F(\cdot)$ is the CDF of Y and

$$l_\tau(\xi) = (\tau - \mathbb{1}_{(\xi < 0)})\xi, \quad \xi \in \mathbb{R} \quad (2)$$

is the so-called pinball loss [47] (Figure 1). In the following, we only consider situations in which F is continuous.

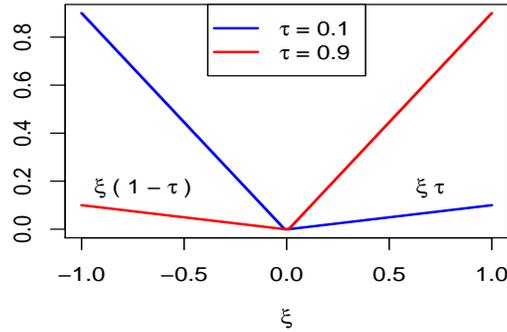


Figure 1: Pinball loss function with $\tau = 0.1$ and $\tau = 0.9$.

Given a finite observation set $\mathcal{Y}_n = (y_1, \dots, y_n)$ composed of i.i.d samples of Y , the empirical estimator of q_τ can thus be introduced in two different ways:

$$\hat{q}_\tau = \min \{y_i \in \mathcal{Y}_n : \hat{F}(y_i) \geq \tau\} \quad (3)$$

or

$$\hat{q}_\tau = \arg \min_{q \in \mathbb{R}} \frac{1}{n} \sum_{i=1}^n l_\tau(y_i - q), \quad (4)$$

where \hat{F} denotes an estimator of the CDF function. In (3) \hat{q}_τ coincides with an order statistic. For example, if \hat{F} is the empirical CDF function then $\hat{q}_\tau = y_{([n\tau])}$, where $[n\tau]$ represents the smallest integer greater than or equal to $n\tau$ and $y_{(k)} = \mathcal{Y}_n(k)$ is the k -th smallest value in the sample $\{y_1, \dots, y_n\}$. The estimators (4) and (3) may coincide, but are in general not equivalent.

Similarly to (1), the conditional quantile of order $\tau \in (0, 1)$ can be defined in two equivalent ways:

$$q_\tau(x) = \min \{q : F(q|X = x) \geq \tau\} = \arg \min_{q \in \mathbb{R}} \mathbb{E}[l_\tau(Y_x - q)], \quad (5)$$

where Y_x is a random variable of distribution \mathbb{P}_x and $F(\cdot|X = x)$ is the CDF of Y_x .

In a quantile regression context, one only has access to a finite observation set $\mathcal{D}_n = \{(x_1, y_1), \dots, (x_n, y_n)\} = (\mathcal{X}_n, \mathcal{Y}_n)$ with \mathcal{X}_n a $n \times d$ matrix. Estimators for (5) are either based on the order statistic as in (3) (section 3), or on a minimizer of the pinball loss as in (4) (sections 4 and 5). Throughout this work, the observation set \mathcal{D}_n is fixed (we do not consider a dynamic or sequential framework). Following the standard approach used in computer experiments, the training points x_i are chosen according to a space-filling design [20] over a hyperrectangle. In particular, we assume

that there are no repeated experiments: $x_i \neq x_j, \forall i \neq j$; most of the methods chosen in this survey (KN, RF, RK, NN, VB) work under that setting.

However, as a baseline approach, one may decide to use a stratified experimental design $\mathcal{D}_{n',r}$ with r i.i.d samples for a given $x_i, i = 1, \dots, n'$, extract pointwise quantile estimates using (3) and fit a standard metamodel to these estimates. The immediate drawback is that for the same budget ($n' \times r = n$) such experimental designs cover much less of the design space than a design with no repetition. The QK method is based on this approach.

3 Methods based on order statistics

A simple way to compute a quantile estimate is to take an order statistic of an i.i.d. sample. A possible approach is to emulate such a sample by selecting all the data points in the neighborhood of the query point x , and then by taking the order statistic of this subsample as an estimator for the conditional quantile. One may simply choose a subsample of \mathcal{D}_n based on a distance defined on \mathcal{X} : this is what the K -nearest neighbors approach does. It is common to use KN based on the Euclidean distance but of course any other distance can be used, such as Mahalanobis [97] or weighted Euclidean distance [31]. Alternatively, one may define a notion of neighborhood using some space partitioning of \mathcal{X} . That includes all the decision tree methods [16], in particular regression trees, bagging or random forest [60].

3.1 K -nearest neighbors

The K -nearest neighbors method was first proposed for the estimation of conditional expectations [86, 87]. Its extension to the conditional quantile estimation can be found in [7].

3.1.1 Quantile regression implementation

Define \mathcal{X}_{test} as the set of query points. KN works as follows: for each $x^* \in \mathcal{X}_{test}$ define $\mathcal{X}^K(x^*)$ the subset of \mathcal{X}_n containing the K points that are the closest to the query point x^* . Define $\mathcal{Y}_K^{x^*}$ the associated outputs, and define $\hat{F}^K(y|X = x^*)$ as the associated empirical CDF. Following (3), the conditional quantile of order τ can be defined as the statistical order

$$\hat{q}_\tau(x^*) = \mathcal{Y}_K^{x^*}([K\tau]). \tag{6}$$

Algorithm 1 details the implementation of the KN method.

Algorithm 1 K -nearest neighbors

- 1: **Inputs:**
 $\mathcal{D}_n, \tau, K, \mathcal{X}_{test}$
 - 2: **for** each point in $x^* \in \mathcal{X}_{test}$ **do**
 - 3: Compute all the distances between x^* and \mathcal{X}_n
 - 4: Sort the computed distances
 - 5: Select the K -nearest points from x^*
 - 6: $\hat{q}_\tau(x^*) = \mathcal{Y}_K^{x^*}([K\tau])$
 - 7: **end for**
-

3.1.2 Computational complexity

For a naive implementation of such an estimator, one needs to compute $n \times N_{new}$ distances, where N_{new} is the number of query points, hence for a cost in $O(nN_{new}d)$. Moreover, sorting n distances in order to extract the K nearest points has a cost in $O(nN_{new} \log n)$. Combining the two operations implies a complexity of order

$$O(nN_{new}d) + O(nN_{new} \log n).$$

Note that some algorithms have been proposed in order to reduce the computational time, for example by using GPUs [35] or by using tree search algorithms [4].

3.2 Random forests

Random forests were introduced by Breiman [15] for the estimation of conditional expectations. They have been used successfully for classification and regression, especially with problems where the number of variables is much larger than the number of observations [28].

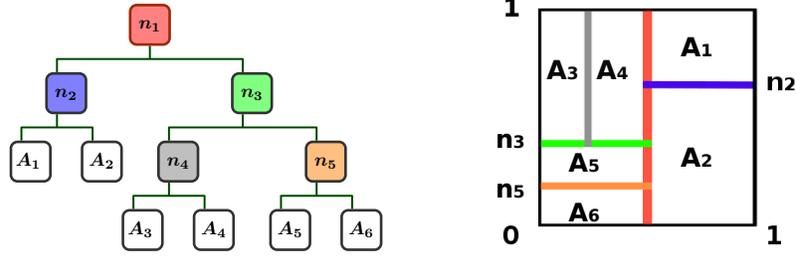


Figure 2: Left: a partitioning tree T . The nodes n_i ($1 \leq i \leq 5$) represent the splitting points, the A_i 's ($1 \leq i \leq 6$) represent the leaves. Right: $\mathcal{X} = [0, 1]^2$ as partitioned by T . The regression tree prediction is constant on each leaf A_i .

3.2.1 Overview

The basic element of random forests is the *regression tree* T , a simple regressor built via a binary recursive partitioning process. Starting with all data in the same partition *i.e.* \mathcal{X} , the following sequential process is applied. At each step, the data is split into two, so that \mathcal{X} is partitioned in a way that it can be represented by a tree as it is presented Figure 2.

Several splitting criteria can be chosen (see [43]). In [60], the splitting point x_S is the data point that minimizes

$$C(x_S) = \sum_{x_i \leq x_S} (y_i - \bar{Y}_L)^2 + \sum_{x_j > x_S} (y_j - \bar{Y}_R)^2, \quad (7)$$

where \bar{Y}_L and \bar{Y}_R are the mean of the left and right sub-populations, respectively. Equation (7) applies when the x 's are real-valued. In the multidimensional case, the dimension d_S in which the split is performed has to be selected. The split then goes through x_S and perpendicularly to the direction d_S . There are several rules to stop the expansion of T . For instance, the process can be stopped when the population of each cell is inferior to a minimal size *nodesize*: then, each node becomes a terminal node or leaf. The result of the process is a partition of the input space into hyperrectangles $R(T)$. Like the KN method, the tree-based estimator is constant on each neighborhood. The hope is that the regression trees automatically build neighborhoods from the data that should be adapted to each problem.

Despite their simplicity of construction and interpretation, regression trees are known to suffer from a certain rigidity and a high variance (see [14] for more details). To overcome this drawback, regression trees can be used with ensemble methods like bagging. Instead of using only one tree, bagging creates a set of tree $\mathcal{T}_N = \{T^1, \dots, T^N\}$ based on a bootstrap version $\mathcal{D}_{N,n} = \left\{ \left((x_{1t}, y_{1t}), \dots, (x_{nt}, y_{nt}) \right) \right\}_{t=1}^N$ of \mathcal{D}_n . Then the final model is created by averaging the results among all the trees.

Bagging reduces the variance of the predictor, as the splitting criterion has to be optimized over all the input dimensions, but computing (7) for each possible split is costly when the dimension is large. The random forest algorithm, a variant of bagging, constructs an ensemble of weak learners based on $\mathcal{D}_{N,n}$ and aggregates them. Unlike plain bagging, at each node evaluation, the algorithm uses only a subset of \tilde{d} covariables for the choice of the split dimensions. Because the \tilde{d} covariables are randomly chosen, the result of the process is a random partition $R(t)$ of \mathcal{X} constructed by the random tree T^t .

3.2.2 Quantile prediction

We present the extension proposed in [60] for conditional quantile regression. Let us define $\ell(x^*, t)$ the leaf obtained from the tree t containing a query point x^* and

$$\omega_i(x^*, t) = \frac{\mathbb{1}_{\{x_i \in \ell(x^*, t)\}}}{\#\{j : x_j \in \ell(x^*, t)\}}, \quad i = 1, \dots, n$$

$$\bar{\omega}_i(x^*) = \frac{1}{N} \sum_{t=1}^N \omega_i(x^*, t).$$

The $\bar{\omega}_i(x^*)$'s represent the weights illustrating the ‘‘proximity’’ between x^* and x_i . In the classical regression case, the estimator of the expectation is:

$$\hat{\mu}(x^*) = \sum_{i=1}^n \bar{\omega}_i(x^*) y_i. \quad (8)$$

In [60] the conditional quantile of order τ is defined as in (3) with the CDF estimator defined as

$$\hat{F}(y|X = x^*) = \sum_{i=1}^n \bar{\omega}_i(x^*) \mathbb{1}_{\{y_i \leq y\}}. \quad (9)$$

Algorithm 2 details the implementation of the RF method.

Algorithm 2 Random forest

```

1: Training
2: Inputs:
    $\mathcal{D}_n, N, \tilde{d}, m_s$ 
3: for each of the  $N$  trees do
4:   Uniformly sample with replacement  $n$  points in  $\mathcal{D}_n$  to create  $\mathcal{D}_{t,n}$ .
5:   Consider the cell  $R = \mathcal{X}$ .
6:   while any cell of the tree contains more than  $m_s$  observations do
7:     for the cells containing more than  $m_s$  observations do
8:       Uniformly sample without replacement  $\tilde{d}$  covariables in  $1, \dots, d$ .
9:       Compute the cell point among the  $\tilde{d}$  covariables that minimizes (7).
10:      Split the cell at this point perpendicularly to the selected covariable.
11:     end for
12:   end while
13: end for
14: Prediction
15: Inputs:
    $\mathcal{X}_{test}, \tau$ 
16: for each point in  $x^* \in \mathcal{X}_{test}$  do
17:   Compute  $\bar{\omega}_i(x^*)$ ,  $i = 1 \dots, n$ 
18:    $\hat{F}(y|X = x^*) = \sum_{i=1}^n \bar{\omega}_i(x^*) \mathbb{1}_{\{y_i \leq y\}}$ 
19:    $\hat{q}_\tau(x^*) = \inf \{y_i : \hat{F}(y_i|X = x^*) \geq \tau\}$ 
20: end for

```

3.2.3 Computational complexity

Assuming that the value of (7) can be computed sequentially for consecutive thresholds, the RF computation burden lies in the search of the splitting point that implies sorting the data. Sorting n variables has a complexity in $O(n \log n)$. Thus, at each node the algorithm finds the best splitting points considering only $\tilde{d} \leq d$ covariables (classically $\tilde{d} = d/3$). This implies a complexity of $O(\tilde{d}n \log n)$ per node. In addition, the depth of a tree is generally upper bounded by $\log n$. Then the computational cost of building a forest containing N trees under the criterion (7) is

$$O(N\tilde{d}n \log^2(n))$$

[55, 99]. One may observe that RF are easy to parallelize and that contrary to KN the prediction time is very small once the forest is built.

4 Approaches based on functional analysis

Functional methods search directly for the function mapping the input to the output in a space fixed beforehand by the user. With this framework, estimating any functional S of the conditional distribution implies selecting a loss l (associated to S) and a function space \mathcal{H} . Thus, the estimator $\hat{S} \in \mathcal{H}$ is obtained as the minimizer of the empirical risk \mathcal{R}_e associated to l , *i.e.*

$$\hat{S} \in \arg \min_{s \in \mathcal{H}} \mathcal{R}_e[s] = \arg \min_{s \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n l(y_i - s(x_i)). \quad (10)$$

The functional space \mathcal{H} must be chosen flexible enough to extract some signal from the data. In addition, \mathcal{H} needs to have enough structure to make the optimization procedure feasible (at least numerically). In the literature, several formalisms such as linear regression [78], spline regression [58], support vector machine [96], neural networks [9] or deep neural networks [76] use structured functional spaces with different levels of flexibility.

However, using a too large \mathcal{H} can lead to overfitting, *i.e.* return predictors that are good only on the training set and generalize poorly. Overcoming overfitting requires some *regularization* (see [77, 104, 105] for instance), defining for example the regularized risk

$$\mathcal{R}_{r,e}[s] = \frac{1}{n} \sum_{i=1}^n l(y_i - s(x_i)) + \lambda \|s\|^\beta, \quad (11)$$

where $\lambda \in \mathbb{R}^+$ is a penalization factor, $\beta \in \mathbb{R}^+$ and $\|\cdot\|$ is either a norm for some methods (Section 4.2) or a measure of variability for others (Section 4.1). The parameter λ plays a major role, as it allows to tune the balance between bias and variance.

Classically, squared loss is used: it is perfectly suited to the estimation of the conditional expectation. Using the pinball loss (Eq. 2) instead allows to estimate quantiles. In this section we present two approaches based on Equation (11) with the pinball loss. The first one is regression using artificial neural networks (NN), a rich and versatile class of functions that has shown a high efficiency in several fields. The second approach is the generalized linear regression in reproducing kernel Hilbert spaces (RK). RK is a non-parametric regression method that has been much studied in the last decades (see [84]) since it appeared in the core of learning theory in the 1990's [77, 96].

4.1 Neural Networks

Artificial neural networks have been successfully used for a large variety of tasks such as classification, computer vision, music generation, and regression [9]. In the regression setting, feed-forward neural networks have shown outstanding achievements. Here we present quantile regression neural network [18] which is an adaptation of the traditional feed-forward neural network.

4.1.1 Overview

A feed-forward neural network is defined by its number of hidden layers H , its numbers of neurons per layer $J_h, 1 \leq h \leq H$, and its activation functions $g_h, h = 1, \dots, H$. Given an input vector $x \in \mathbb{R}^d$ the information is fed to the hidden layer 1 composed of a fixed number of neurons J_1 . For each neuron $N_i^{(1)}, i = 1, \dots, J_1$, a scalar product (noted $\langle \cdot, \cdot \rangle$) is computed between the input vector $x = (x_1, \dots, x_d) \in \mathbb{R}^d$ and the weights $w_i^{(1)} = (w_{i,1}^{(1)}, \dots, w_{i,d}^{(1)}) \in \mathbb{R}^d$ of the $N_i^{(1)}$ neurons. Then a bias term $b_i^{(1)} \in \mathbb{R}$ is added to the result of the scalar product. The result is composed with the activation function g_1 (linear or non-linear) which is typically the sigmoid or the ReLu function [76] and the result is given to the next layer where the same operation is processed until the information comes out from the output layer. For example, the output of a 3-layers NN at x^* is given by

$$s(x^*) = g_3 \left(\sum_{j=1}^{J_2} g_2 \left(\sum_{i=1}^{J_1} g_1 (\langle w_i^{(1)}, x^* \rangle + b_i^{(1)}) w_{j,i}^{(2)} + b_j^{(2)} \right) w_{1,j}^{(3)} + b^{(3)} \right). \quad (12)$$

The corresponding architecture can be found in Figure 3.

The architecture of the NN defines \mathcal{H} . Finding the right architecture is a very difficult problem which will not be treated in this paper. However a classical implementation procedure consists of creating a network large enough (able to overfit) and then using techniques such as early stopping, dropout, bootstrapping or risk regularization to avoid overfitting [82]. In [18], the following regularized risk is used:

$$\mathcal{R}_{r,e}[s] = \frac{1}{n} \sum_{i=1}^n l(y_i - s(x_i)) + \lambda \sum_{j=1}^H \sum_{z=1}^{J_j} \|w_z^{(j)}\|^2. \quad (13)$$

4.1.2 Quantile regression

Minimizing Equation (13) (with respect to all the weights and biases) is in general challenging, as $\mathcal{R}_{r,e}$ is a highly multimodal function. It is mostly tackled using derivative-based algorithms and multi-starts (*i.e.* launching the optimization procedure M_s times with different starting points). In the case of quantile estimation, the loss function is

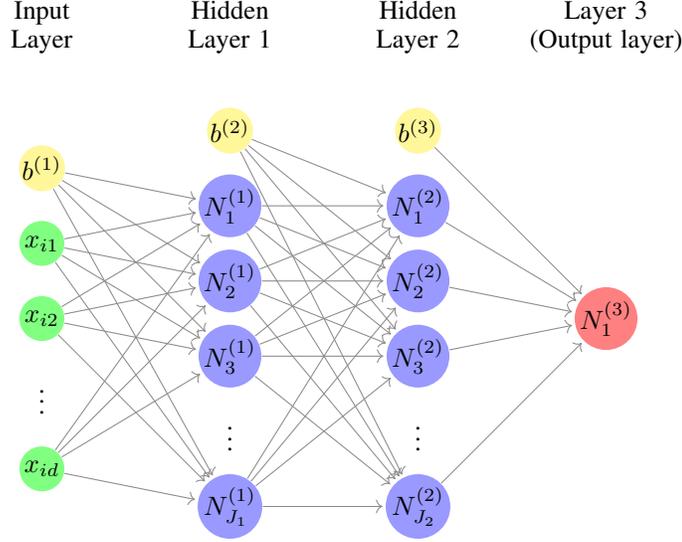


Figure 3: Architecture of 3-layer feedforward neural network.

non-differentiable at the origin, which may cause problems to some numerical optimization procedures. To address this issue, [18] introduced a smooth version of the pinball loss function, defined as:

$$l_r^\eta(\xi) = h^\eta(\xi)(\tau - \mathbf{1}_{\xi < 0}),$$

where

$$h^\eta(\xi) = \begin{cases} \frac{\xi^2}{2\eta} & \text{if } 0 \leq |\xi| \leq \eta \\ |\xi| - \frac{\eta}{2} & \text{if } |\xi| \geq \eta. \end{cases} \quad (14)$$

Note that if the optimizer is based on a first order method such as [46], then the transfer function does not require continuous derivatives. But using a second order method as it is done in the original paper implies the loss function to be twice differentiable with respect to the weights of the neural network. Then transfer functions such as logistic or hyperbolic tangent functions should be used over piecewise linear ones such as the ReLu or the PReLU functions [68].

Let us define \mathbf{w} the list containing the weights and bias of the network. To find \mathbf{w}^* , a minimizer of $\mathcal{R}_{r,e}$, the idea is to solve a series of problems using the smoothed loss instead of the pinball one with a sequence E_K corresponding to K decreasing values of η . The process begins with the optimization with the larger value η_1 . Once the optimization converges, the optimal weights are used as the initialization for the optimization with η_2 , and so on. The process stops when the weights based on $l_r^{\eta_K}$ are obtained. Finally, $\hat{q}_\tau(x^*)$ is given by the evaluation of the optimal network at x^* . Algorithm 3 details the implementation of the NN method.

4.1.3 Computational complexity

In [18] the optimization is based on a Newton method. Thus the procedure needs to inverse a Hessian matrix. Without sophistications, its cost is $O(s_{\text{pb}}^3)$ with s_{pb} the size of the problem *i.e.* the number of parameters to optimize. Note that using a high order method makes sense here because NN has few parameters (in contrast to deep learning methods). Moreover providing an upper bound on the number of iterations needed to reach an optimal point may be really hard in practice because of the non convexity of (13). In the non-convex case, there is no optimality guaranty and the optimization could be stuck in a local minima. However, it can be shown that the convergence near a local optimal point is at least super linear (see [12] Eq. (9.33)) and may be quadratic (if the gradient is small). It implies, for each η , the number of iterations until $\mathcal{R}_{r,e}^\eta(\mathbf{w}) - \mathcal{R}_{r,e}^\eta(\mathbf{w}^*) \leq \epsilon$ is bounded above by

$$\frac{\mathcal{R}_{r,e}^\eta(\mathbf{w}_0) - \mathcal{R}_{r,e}^\eta(\mathbf{w}^*)}{\gamma} + \log_2 \log_2(\epsilon_0/\epsilon),$$

with γ the minimal decreasing rate, $\epsilon_0 = 2M_\eta^3/L_\eta^2$, M_η the strong convexity constant of $\mathcal{R}_{r,e}^\eta$ near \mathbf{w}^* and L_η the Hessian Lipschitz constant (see [12] page 489). As $\log_2 \log_2(\epsilon_0/\epsilon)$ increases very slowly with respect to ϵ , it is

Algorithm 3 Neural network

-
- 1: **Training**
 - 2: **Inputs:**
 $\mathcal{D}_n, \tau, \lambda, H, (J_1, \dots, J_H), (g_1, \dots, g_H), E_K$
 - 3: **Initialize:**
Fix w_0 as the list containing the initial weights and biases.
 - 4: **for** $t = 1$ to K **do**
 - 5: $\epsilon \leftarrow E_K[t]$
 - 6: Starting the optimization procedure with w_0 and define

$$w_\tau^* = \arg \min_w \frac{1}{n} \sum_{t=1}^n l_\tau^\epsilon(y_i - \hat{q}^w(x_i)) + \lambda \sum_{j=1}^H \sum_{i=1}^J \|w_i^{(j)}\|^2$$

with $\hat{q}^w(\cdot)$ the output of the network with the weights w .

- 7: $w_0 \leftarrow w_\tau^*$
 - 8: **end for**
 - 9: **Prediction**
 - 10: **Inputs:**
 $\mathcal{X}_{test}, w_\tau^*, \lambda, H, (J_1, \dots, J_H), (g_1, \dots, g_H)$
 - 11: **for** each point in $x^* \in \mathcal{X}_{test}$ **do**
 - 12: Define $\hat{q}_\tau(x^*) = \hat{q}^{w_\tau^*}(x^*)$.
 - 13: **end for**
-

possible to bound the number of iterations N typically by

$$\frac{\mathcal{R}_{r,e}^\eta(w_0) - \mathcal{R}_{r,e}^\eta(w^*)}{\gamma} + 6.$$

That means, near an optimal point, the complexity is $O(L_\eta n (Jd)^3)$, with J the total number of neurons. Then using a multistart procedure implies a complexity of

$$O(M_s L_{\eta^*} n (Jd)^3),$$

with $L_{\eta^*} = \max_{\eta_1, \dots, \eta_K} L_\eta$.

4.2 Generalized linear regression

Regression in RKHS was introduced for classification via Support Vector Machine by [25, 41], and has been naturally extended for the estimation of the conditional expectation [30, 72]. Since, many applications have been developed (see [84, 77] for some examples), here we present the quantile regression in RKHS [90, 75].

4.2.1 RKHS introduction and formalism

Under the linear regression framework, S is assumed to be under the form $S(x) = x^T \alpha$, with α in \mathbb{R}^d . To stay in the same vein while creating non-linear responses, one can map the input space \mathcal{X} to a space of higher dimension \mathcal{H} (named the feature space), thanks to a feature map Φ . For example the feature space could be a polynomial space, in that case we are working with the spline framework [58]. For a large flexibility and few parameters, the feature space can even be chosen as an infinite dimensional space. In the following, $\Phi = (\phi_1, \phi_2, \phi_3, \dots)$ defines a feature map from \mathcal{X} to \mathcal{H} , where \mathcal{H} is the \mathbb{R} -Hilbert functional space defined as

$$\mathcal{H} = \left\{ s, s(x) = \sum_{j \in J} \alpha_j \phi_j(x), \text{ s.t. } \|s\|_{\mathcal{H}} < +\infty \right\},$$

$$\text{with } \|s\|_{\mathcal{H}} := \sqrt{\langle s, s \rangle_{\mathcal{H}}},$$

where J is the cardinality of a basis of \mathcal{H} . Under the hypothesis that S belongs to \mathcal{H} , S can be written as

$$S(x) = \sum_{j \in J} \alpha_j \phi_j(x). \tag{15}$$

Notice that without more hypothesis on \mathcal{H} , estimating S is difficult. In fact it is impossible to compute (15) directly because of the infinite sum. Thus, using the sample \mathcal{D}_n , a solution of (10) is not known and cannot be computed.

However, this issue can be tackled by the introduction of the regularized empirical risk

$$\mathcal{R}_{r,e}[s] = \frac{1}{n} \sum_{i=1}^n l(y_i - s(x_i)) + \frac{\lambda}{2} \|s\|_{\mathcal{H}}^2, \quad (16)$$

and the utilization of the RKHS formalism that is based on the *representer theorem* and the so-called *kernel trick* (see [77] for instance).

Let us first introduce the symmetric definite positive function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ such that:

$$k(x, x') = \langle \Phi(x'), \Phi(x) \rangle_{\mathcal{H}}. \quad (17)$$

Under this setting, \mathcal{H} is a RKHS with the reproducing kernel k , that means $\Phi(x) = k(\cdot, x) \in \mathcal{H}$ for all $x \in \mathcal{X}$ and the reproducing property

$$s(x) = \langle s, k(\cdot, x) \rangle_{\mathcal{H}} \quad (18)$$

holds for all $s \in \mathcal{H}$ and all $x \in \mathcal{X}$. It can be shown that working with a fixed kernel k is equivalent to working with its associated functional Hilbert space. Note that the kernel choice is based on kernel properties or assumptions made on the functional space. See for instance [84], chapter 4, for some kernel definitions and properties. In the following, \mathcal{H}_θ and k_θ denote respectively a RKHS and its kernel associated to the hyperparameters vector θ . $K_{x,x}^\theta \in \mathcal{R}^{n \times n}$ is the kernel matrix obtained via $K_{x,x}^\theta(i, j) = k_\theta(x_i, x_j)$.

From a theoretical point of view, the *representer theorem* implies that the minimizer \hat{S} of (16) lives in $\mathcal{H}_{|X}^\theta = \text{span}\{\Phi(x_i) : i = 1, \dots, n\}$ with $\|s\|_{\mathcal{H}_{|X}}^2 = \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j k_\theta(x_j, x_i)$. Combining this result to the definition (17) and the reproducing property (18), it is possible to rewrite \hat{S} as:

$$\hat{S}(x) = \sum_{i=1}^n \alpha_i k_\theta(x, x_i).$$

Hence, the original infinite dimensional problem associated to the formalism (15) becomes an optimization problem over n coefficients $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n) \in \mathbb{R}^n$. More precisely, finding \hat{S} is equivalent to minimize in α the quantity

$$\frac{1}{n} \sum_{i=1}^n l\left(y_i - \left(\sum_{j=1}^n \alpha_j k_\theta(x_i, x_j)\right)\right) + \frac{\lambda}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j k_\theta(x_j, x_i). \quad (19)$$

4.2.2 Quantile regression

Quantile regression in RKHS was introduced by [90], followed by several authors [52, 85, 23, 24, 75]. Quantile regression has two specificities compared to the general case. Firstly the loss l is defined as the pinball. Secondly, to ensure the quantile property, the intercept is not regularized. More precisely, we assume that

$$q_\tau(x) = g(x) + b \quad \text{with } g \in \mathcal{H}_\theta \quad \text{and } b \in \mathbb{R}.$$

and we consider the empirical regularized risk

$$\mathcal{R}_{r,e}[q] := \frac{1}{n} \sum_{i=1}^n l_\tau(y_i - q(x_i)) + \frac{\lambda}{2} \|g\|_{\mathcal{H}_\theta}^2. \quad (20)$$

Thus the *representer theorem* implies that \hat{q}_τ can be written under the form

$$\hat{q}_\tau(x^*) = \sum_{i=1}^n \alpha_i k_\theta(x^*, x_i) + b,$$

for a new query point x^* . Since (20) cannot be minimized analytically, a numerical minimization procedure is used. [25] followed by [90] introduced nonnegative variables $\xi^{(*)} \in \mathbb{R}^+$ to transform the original problem into

$$\mathcal{R}_{r,e}[q] := \frac{1}{n} \sum_{i=1}^n \tau \xi_i + (1 - \tau) \xi_i^* + \frac{\lambda}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j k_\theta(x_j, x_i),$$

subject to

$$y_i - \left(\sum_{j=1}^n \alpha_j k_\theta(x_i, x_j) + b \right) \leq \xi_i$$

and

$$\sum_{j=1}^n \alpha_j k_\theta(x_i, x_j) + b - y_i \leq \xi_i^*, \text{ where } \xi_i^*, \xi_i^* \geq 0.$$

Using a Lagrangian formulation, it can be shown (see [84] for instance) that minimizing $\mathcal{R}_{\tau, e}$ is equivalent to the problem:

$$\begin{aligned} \min_{\alpha \in \mathbb{R}^n} \quad & \frac{1}{2} \alpha^T K_{x,x} \alpha - \alpha^T \mathbf{y} \\ \text{s.t.} \quad & \frac{1}{\lambda n} (\tau - 1) \leq \alpha_i \leq \frac{1}{\lambda n} \tau, \forall 1 \leq i \leq n \\ \text{and} \quad & \sum_{i=1}^n \alpha_i = 0. \end{aligned} \tag{21}$$

It is a quadratic optimization problem under linear constraint, for which many efficient solvers exist.

The value of b may be obtained from the Karush-Kuhn-Tucker slackness condition or fixed independently of the problem. A simple way to do so is to choose b as the τ -quantile of $(y_i - \sum_{j=1}^n \alpha_j k_\theta(x_i, x_j))_{1 \leq i \leq n}$. Algorithm 4 details the implementation of the RK method.

Algorithm 4 RKHS regression

1: **Training**

2: **Inputs:**

$$\mathcal{D}_n, \tau, \lambda, k_\theta$$

3: **Initialize:**

$$\text{Compute the } n \times n \text{ matrix } K_{x,x}^\theta$$

4: **Optimization:** Select α^* as

$$\begin{aligned} \alpha^* = \arg \min_{\alpha \in \mathbb{R}^n} \quad & \frac{1}{2} \alpha^T K_{x,x}^\theta \alpha - \alpha^T \mathbf{y} \\ \text{s.t.} \quad & \frac{1}{\lambda n} (\tau - 1) \leq \alpha_i \leq \frac{1}{\lambda n} \tau, \forall 1 \leq i \leq n \\ \text{and} \quad & \sum_{i=1}^n \alpha_i = 0 \end{aligned}$$

5: Define b as the τ -quantile of $(y_i - \sum_{j=1}^n \alpha_j k_\theta(x_i, x_j))_{1 \leq i \leq n}$

6: **Prediction**

7: **Inputs:**

$$\mathcal{X}_{test}, \alpha^*, k_\theta$$

8: **for** each point in $x^* \in \mathcal{X}_{test}$ **do**

9: compute $K_{x^*,x}^\theta$

10: $\hat{q}_\tau(x_{test}) = K_{x^*,x}^\theta \alpha^* + b$

11: **end for**

4.2.3 Computational complexity

Let us notice two things. Firstly, the minimal upper bound complexity for solving (21) is $O(n^3)$. Indeed solving (21) without the constraints is easier and it needs $O(n^3)$. Secondly the optimization problem (21) is convex, thus the optimum is global.

There are two main approaches for solving (21), the interior point method [12] and the iterative methods like libSVM [21]. The interior point method is based on the Newton algorithm, one method is the barrier method (see [12] page 590). It is shown that the number of iterations N for reaching a solution with precision ϵ is $O(\sqrt{n} \log(\frac{n}{\epsilon}))$. Moreover

each iteration of a Newton type algorithm costs $O(n^3)$ because it needs to inverse a Hessian. Thus, the complexity of an interior point method for finding a global solution with precision ϵ is

$$O\left(n^{7/2} \log(n/\epsilon)\right).$$

On another hand, iterative methods like libSVM transform the main problem into a smaller one. At each iteration the algorithm solves a subproblem in $O(n)$. Contrary to the interior point methods, the number of iterations depends explicitly on the matrix $K_{x,x}^\theta$. [53] shows that the number of iterations is

$$O\left(n^2 \kappa(K_{x,x}^\theta) \log(1/\epsilon)\right),$$

where $\kappa(K_{x,x}^\theta) = \lambda_{\max}(K_{x,x}^\theta)/\lambda_{\min}(K_{x,x}^\theta)$. Note that $\kappa(K_{x,x}^\theta)$ depends on the type of the kernel, it evolves in $O(n^{s'})$ with $s' > 1$ an increasing value of the regularity of k_θ [22]. For more information about the eigenvalues of $K_{x,x}^\theta$ one can consult [13].

To summarize, it implies that the complexity of the libSVM method has an upper bound higher than the interior point algorithm. However, these algorithms are known to converge pretty fast. In practice, the upper bound is almost never reached, and thus the most important factor is the cost per iteration, rather than the number of iterations needed. This is the reason why libSVM is popular in this setting.

5 Bayesian approaches

Bayesian formalism has been used for a wide class of problems such as classification and regression [69], model averaging [11] and model selection [67].

Under the regression framework, a classical Bayesian model is defined as

$$y = S(x) + \varepsilon(x), \tag{22}$$

where S depends on parameters that have to be inferred, y is an observation and ε a noise term. Knowing \mathcal{D}_n , a classical goal is to estimate the posterior distribution $p(S|\mathcal{D}_n)$ at x^*

$$p(S|\mathcal{D}_n, x^*) = \int p(S|\zeta, \mathcal{D}_n, x^*) p(\zeta|\mathcal{D}_n) d\zeta,$$

with ζ a vector of parameters that characterizes S . The posterior $p(\zeta|\mathcal{D}_n)$ is a priori unknown and $p(S|\zeta, \mathcal{D}_n, x^*)$ is the model hypothesis.

The first Bayesian quantile regression framework was introduced in [101] where the authors worked under a linear framework, *i.e.* $q_\tau(x) = x^T \zeta$ but the linear hypothesis is too restrictive to treat the stochastic black box setting. [89] introduced a mixture modeling framework called Dirichlet process to perform nonlinear quantile regression. However the inference is performed with MCMC methods [37, 34], a procedure that is often costly. A possible alternative is the use of Gaussian process (GP). GPs are powerful in a Bayesian context because of their flexibility and their tractability (GP are only characterized by their mean m and covariance k_θ).

Using GPs, a possible approach is to use a joint modeling of the mean and variance, assuming that the distribution is Gaussian everywhere, and then to extract the quantiles of interest as it is done in [45, 49]. Nevertheless this strategy may introduce a high bias when the true output distribution is not Gaussian.

In this section we present QK and VB, two approaches that use a Gaussian process (GP) as a prior for q_τ . Contrary to the joint modelling approach, here the GP prior for q_τ does not imply any structure on the output distribution, which allows the creation of very flexible quantile models.

5.1 Quantile kriging

Kriging takes its origins in geostatistics and spatial data interpolation [26, 83]. Since the 2000's, kriging drew attention of the machine learning community (see [69] for some applications). In this section we present a very intuitive method that gives flexible quantile estimators based on data containing repetition and GPs [66].

5.1.1 Kriging introduction

Kriging is based on the hypothesis that

$$S(x) \sim \mathcal{GP}(m(x), k_\theta(x, x')) . \tag{23}$$

For every finite set (x_1, \dots, x_T) , the output $(S(x_1), \dots, S(x_T))$ is multivariate Gaussian. Here m is the mean of the process and k_θ is a kernel function also known as the covariance function. Note that without loss of generality, in the sequel we take $m = 0$ in order to simplify the computations and notations. The covariance function conveys many properties of the process, so its choice should depend on the assumptions made on S . A first classical assumption is that the GP is stationary, *i.e.* the correlation between two inputs does not depend on the location but only on the distance between the points. Then, different class of stationary kernels produce GPs with different regularities. The class of Matérn kernels is very convenient because it depends on a regularity hyperparameter that enables the end user to adapt its prior to his regularity assumptions (see [69] for more details). For example, the Matérn 5/2 kernel is defined as

$$k_\theta(x, x') = \rho^2 \left(1 + \sqrt{5} \|x - x'\|_\theta + \frac{5}{3} \|x - x'\|_\theta^2 \right) \exp(-\sqrt{5} \|x - x'\|_\theta), \quad (24)$$

where $\rho > 0$ and

$$\|x - x'\|_\theta^2 = (x - x')^T \Lambda_\theta (x - x'),$$

with Λ_θ a diagonal matrix with diagonal terms the inverses of the d squared length scales $\theta_i, i = 1, \dots, d$.

In addition to the first assumption, let us assume that ε in (22) is Gaussian white noise, that is:

$$y_i = S(x_i) + \varepsilon_i \quad \text{with} \quad \varepsilon_i \sim \mathcal{N}(0, \sigma_i^2). \quad (25)$$

The associated likelihood is Gaussian, *i.e.*

$$p(\mathcal{Y}_n | \mathcal{X}_n, \zeta) = \mathcal{N}(\mathcal{Y}_n | S, \text{diag}(\sigma^2)), \quad (26)$$

with $\sigma = (\sigma_1, \dots, \sigma_n)$. Then under the hypotheses (23) and (25), the marginal $p(\mathcal{Y}_n | \mathcal{X}_n)$ is Gaussian with zero mean and covariance $K_{x,x}^\theta + \text{diag}(\sigma^2)$. Following [69], $(\mathcal{Y}_n, S(x^*))$ is a Gaussian vector of zero mean and covariance

$$\mathbf{K} = \begin{pmatrix} K_{x,x}^\theta + \text{diag}(\sigma^2) & K_{x,x^*}^\theta \\ K_{x^*,x}^\theta & K_{x^*,x^*}^\theta \end{pmatrix}. \quad (27)$$

Because the distribution of $S(x^*)$ knowing \mathcal{Y}_n is still Gaussian (see [69] appendix A.2 for more details), it is possible to provide the distribution for Kriging regression model as

$$S(x^*) \sim \mathcal{GP}(\bar{S}(x^*), \mathbb{V}_S(x^*)) \quad (28)$$

with

$$\begin{aligned} \bar{S}(x^*) &= K_{x^*,x}^\theta (K_{x,x}^\theta + \text{diag}(\sigma^2))^{-1} \mathcal{Y}_n, \\ \mathbb{V}_S(x^*) &= k_\theta(x^*, x^*) - K_{x^*,x}^\theta (K_{x,x}^\theta + \text{diag}(\sigma^2))^{-1} K_{x,x^*}^\theta. \end{aligned}$$

As in Section 4.2, the covariance functions are usually chosen among a set of predefined ones (for example the Matérn 5/2 see Eq.24), that depend on a set of hyperparameters $\theta \in \mathbb{R}^{d+1}$. The best hyperparameter θ^* can be selected as the maximizer of the marginal likelihood. More precisely, since $p(\mathcal{Y}_n | S) = \mathcal{N}(S, \text{diag}(\sigma^2))$ it follows [69] that

$$\begin{aligned} p(\mathcal{Y}_n | \mathcal{X}_n, \theta) &= -\frac{1}{2} \mathcal{Y}_n^T (K_{x,x}^\theta + B)^{-1} \mathcal{Y}_n \\ &\quad - \frac{1}{2} \log |(K_{x,x}^\theta + B)| - \frac{n}{2} \log(2\pi), \end{aligned} \quad (29)$$

where $|K|$ is the determinant of the matrix K . Maximizing this likelihood with respect to θ is usually done using derivative-based algorithms, although the problem is non-convex and known to have several local maxima.

Different estimators of S may be extracted based on (28). Here \hat{S} is fixed as \bar{S}_{θ^*} . Note that this classical choice is made because the maximum a posteriori of a Gaussian distribution coincides with its mean.

5.1.2 Quantile kriging prediction

As q_τ is a latent quantity, the solution proposed in [66] is to consider the sample $\mathcal{D}_{n',r}$ corresponding to a design of experiments with n' different points that are repeated r times in order to obtain quantile observations. For each $x_i \in \mathcal{X}$, $1 \leq i \leq n'$, let us define:

$$y_{i,r} = (y_{i,1}, \dots, y_{i,r})$$

and

$$\mathcal{D}_{n',\tau,r} = \left((x_1, \hat{q}_\tau(x_1)), \dots, (x_n, \hat{q}_\tau(x_n)) \right), \text{ with } \hat{q}_\tau(x_i) = y_{i,([r\tau])}.$$

Algorithm 5 Quantile kriging

```

1: Training
2: Inputs:
    $\mathcal{D}_{n',r,\tau}, k_\theta$ 
3: Initialize:
   Compute the  $n' \times n'$  matrix  $K_{x,x}^\theta$ 
   Define the local estimators of the  $\tau$ -quantile
4: for  $i = 1$  to  $n'$  do
5:    $\hat{q}_\tau(x_i) \leftarrow y_{i,(\lceil r\tau \rceil)}$ 
6:   Estimate  $\sigma_i$  by bootstrap
7: end for
8: Define  $B = \text{Diag}(\sigma_1^2, \dots, \sigma_{n'}^2)$  and compute  $(K_{x,x}^\theta + B)^{-1}$ 
9: Define the kernel hyperparameters  $\theta^*$  as  $\arg \max_\theta$  of
   
$$-\frac{1}{2} \mathcal{Q}_{n'}^T (K_{x,x}^\theta + B)^{-1} \mathcal{Q}_{n'} - \frac{1}{2} \log |K_{x,x}^\theta + B| - \frac{n}{2} \log(2\pi)$$

10: Inputs:
    $\mathcal{X}_{test}, \mathcal{Q}_{n'}, \theta^*, B$ 
11: for each point in  $x^* \in \mathcal{X}_{test}$  do
12:   compute  $K_{x^*,x}^{\theta^*}$ 
13:    $\hat{q}_\tau^{\theta^*}(x^*) = K_{x^*,x}^{\theta^*} (K_{x,x}^{\theta^*} + B)^{-1} \hat{q}_\tau$ 
14: end for

```

Following [66], let us assume that

$$\hat{q}_\tau(x_i) = q_\tau(x_i) + \varepsilon_i, \text{ with } \varepsilon_i \sim \mathcal{N}(0, \sigma_i^2). \quad (30)$$

Note that from a statistical point of view Assumption (30) is wrong but asymptotically consistent. The resulting estimator is

$$\hat{q}_\tau(x^*) = K_{x^*,x}^\theta (K_{x,x}^\theta + B)^{-1} \mathcal{Q}_{n'}, \quad (31)$$

with $\mathcal{Q}_{n'} = (\hat{q}_\tau(x_1), \dots, \hat{q}_\tau(x_{n'}))$ and $B = \text{diag}(\sigma_1^2, \dots, \sigma_{n'}^2)$.

There are several possibilities to evaluate the noise variances σ_i^2 . Here we choose to use a bootstrap technique (that is, generate bootstrapped samples of $y_{i,r}$, compute the corresponding $\hat{q}_\tau(x_i)$ values and take the variance over those values as the noise variance) but it is possible to use the central limit theorem as it is presented in [5]. The hyperparameters are selected based on (29) changing \mathcal{Y}_n by $\mathcal{Q}_{n'}$. Algorithm 5 details the implementation of the QK method.

5.1.3 Computational complexity

Optimizing (29) with a Newton type algorithm implies to inverse a $(d+1) \times (d+1)$ matrix. In addition, for each θ , obtaining the partial derivatives of (29) requires the computation of $(K_{x,x}^\theta + B)^{-1}$ [69]. Thus at each step of the algorithm, the complexity is $O(n'^3 + d^3)$. Assuming the starting point θ_{start} is close to an optimal θ^* , based on the same analysis as in section 4.1.3, the complexity to find θ^* is

$$O(L(d^3 + n'^3)),$$

with L the Hessian Lipschitz constant.

Finally, obtaining $\hat{q}_\tau^{\theta^*}$ from (31) implies inverting the matrix $K_{x,x}^{\theta^*} + B$ that is in $O(n'^3)$. So the whole complexity is

$$O(L(d^3 + n'^3) + n'^3).$$

5.2 Bayesian variational regression

Quantile kriging requires repeated observations to obtain direct observations of the quantile and make the hypothesis of Gaussian errors acceptable. Variational approaches allow us to remove this critical constraint, while setting a more realistic statistical hypothesis on ε . Starting from the decomposition of Eq. 22, $\varepsilon(x)$ is now assumed to follow a Laplace asymmetric distribution [102, 57], implying:

$$p(y|q, \tau, \sigma, x) = \frac{\tau(1-\tau)}{\sigma} \exp\left(-\frac{l_\tau(y-q(x))}{\sigma}\right), \quad (32)$$

with the priors on q and σ that has to be fixed.

Such assumption may be justified by the fact that minimizing the empirical risk associated to the pinball loss is equivalent to maximizing the asymmetric Laplace likelihood, which is given by

$$p(\mathcal{Y}_n|q_\tau, \mathcal{X}_n, \theta) = \prod_{i=1}^n \frac{\tau(1-\tau)}{\sigma} \exp\left(-\frac{l_\tau(y_i - q_\tau(x_i))}{\sigma}\right). \quad (33)$$

According to the Bayes formula, the posterior can be written as

$$p(q_\tau|\mathcal{D}_n) = \frac{p(\mathcal{Y}_n|\mathcal{X}_n, q_\tau) p(q_\tau)}{p(\mathcal{Y}_n|\mathcal{X}_n)}.$$

As the normalizing constant is independent of q_τ , considering only the likelihood and the prior is enough. We obtain

$$p(q_\tau|\mathcal{D}_n) \propto p(\mathcal{Y}_n|\mathcal{X}_n, q_\tau) p(q_\tau). \quad (34)$$

Because of the Laplace asymmetric likelihood, contrary to the classical kriging model, here the posterior distribution (34) is not Gaussian anymore. Thus it is not possible to provide an analytical expression for the regression model. To overcome this problem, [10] used a variational approach with an expectation-propagation (EP) algorithm [62], while [1] used a variational expectation maximization (EM) algorithm which was found to perform slightly better.

5.2.1 Variational EM algorithm

The EM algorithm was introduced in [27] to compute maximum-likelihood estimates. Since then, it has been widely used in a large variety of fields (see [59] for more details). Classically, the purpose of the EM algorithm is to find ζ that maximizes $p(\mathcal{Y}_n|\zeta)$ thanks to the introduction of the hidden variables z . However dealing with this classical formalism implies that $p(z|\mathcal{Y}_n, \zeta)$ is known or some sufficient statistics can be computed (see [94, 71] for more details), which is not always possible. Using the variational EM framework is a possibility to bypass this requirement [94], by approximating $p(z|\mathcal{Y}_n, \zeta)$ by a probability distribution \tilde{p} that factorizes under the form

$$\tilde{p}(z) = \prod_{j=1}^M \tilde{p}_j(z_j).$$

Starting from the log-likelihood $\log(p(\mathcal{Y}_n|\zeta))$, thanks to Jensen's inequality, it is possible to show that:

$$\log(p(\mathcal{Y}_n|\zeta)) \geq \mathcal{L}(\tilde{p}, \zeta) + \text{kl}(\tilde{p}||p),$$

where

$$\mathcal{L}(\tilde{p}, \zeta) = \int \tilde{p}(z) \log\left(\frac{p(\mathcal{Y}_n, z|\zeta)}{\tilde{p}(z)}\right) dz,$$

and kl is the Kullback-Leibler divergence:

$$\text{kl}(\tilde{p}||p) = - \int \tilde{p}(z) \log\left(\frac{p(z|\mathcal{Y}_n, \zeta)}{\tilde{p}(z)}\right) dz.$$

As presented on Figure 5.2.1, the EM algorithm can be viewed as a two-step optimization technique. The lower bound \mathcal{L} is first maximized with respect to \tilde{p} (E-step) so that to minimize $\text{kl}(\tilde{p}||p)$. Next the likelihood is directly maximized with respect to the parameter ζ (M-step).

Classically and in the following, the E-step optimization problem is analytically solved (see [94] for details about the computation). In this particular case, the quantities \tilde{p}_j are limited to conditionally conjugate exponential families. But note that different strategies have been developed to relax this assumption (generally it comes with an higher computational complexity) so that creating more flexible models (see the review [103] for instance).

5.2.2 Variational EM applied to quantile regression

Following [1], let us suppose that

$$\begin{aligned} q_\tau(x) &\sim \mathcal{GP}(m(x), k_\theta(x, x')) \\ \sigma &\sim \text{IG}(10^{-6}, 10^{-6}), \end{aligned}$$

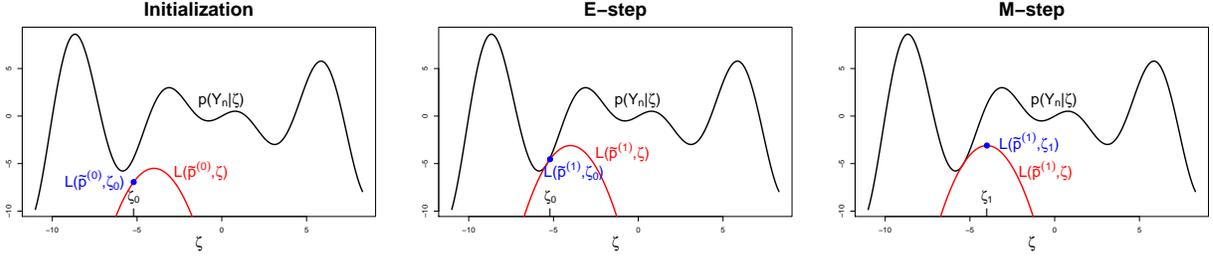


Figure 4: First steps of the variational EM algorithm.

with IG defining the inverse gamma distribution and without loss of generality, $m(\cdot) = 0$. Note that contrary to the formalism introduce with QK, here σ is taken as a random variable. To allow analytical computation, let us introduce an alternative definition of the Laplace distribution [57, 48]:

$$p(y_i|q_\tau, x_i, \sigma, \tau) = \int \mathcal{N}(y_i|\mu_i, \sigma y_i) \exp(-w_i) dw, \quad (35)$$

where $\mu_{y_i} = q_\tau(x_i) + \frac{1-2\tau}{\tau(1-\tau)}\sigma w_i$, $\sigma_{y_i} = \frac{2}{\tau(1-\tau)}\sigma^2 w_i$ and w_i is distributed according to an exponential law of parameter 1.

The distribution of q_τ at a new point x^* is given by averaging the output of all Gaussian models with respect to the posterior $p(q_\tau, \sigma, w|x, Y)$:

$$p(q_\tau|x^*, \mathcal{D}_n) = \int p(q_\tau|x^*, q_\tau, \sigma, w, \mathcal{D}_n) p(q_\tau, \sigma, w|\mathcal{D}_n) dq d\sigma dw. \quad (36)$$

Here the crux is to compute the posterior

$$p(q_\tau, \sigma, w|\mathcal{D}_n) \propto p(\mathcal{Y}_n|q_\tau, \sigma, w, \mathcal{X}_n) p(q_\tau, \sigma, w).$$

To do so, in [1] the authors use $z = (q_\tau(\mathcal{X}_n), w, \sigma)$ as hidden variables and $\zeta = \theta \in \mathbb{R}^{d+1}$ as parameters and the variational factorization approximation

$$p(q_\tau, \sigma, w|\mathcal{D}_n) \approx \tilde{p}(q_\tau, w, \sigma) = \tilde{p}(q_\tau) \tilde{p}(w) \tilde{p}(\sigma). \quad (37)$$

The EM algorithm provides a nice formalism here. Although the goal is to find θ such that $p(\mathcal{Y}_n|\mathcal{X}_n, \theta)$ is maximal, the algorithm estimates the underlying GP (i.e. $p(q_\tau, \sigma, w|\mathcal{D}_n)$) that is able to have a likelihood as large as possible. Then the estimated value $p(q_\tau, \sigma, w|\mathcal{D}_n)$ is plugged into (36) to provide the final quantity of interest.

E-step. Because the posteriors are conjugated, it is possible to obtain an analytical expression of the optimal distribution $\tilde{p}(q_\tau)$:

$$\tilde{p}(q_\tau|\mathcal{D}_n) \sim \mathcal{N}(\mu_\theta, \Sigma_\theta),$$

where

$$\mu_\theta = \Sigma_\theta \left(\mathbf{D}\mathcal{Y}_n - \frac{1-2\tau}{2} \mathbb{E}\left(\frac{1}{\sigma}\right) \mathbf{1} \right)$$

and

$$\Sigma_\theta = (\mathbf{D} + K_{x,x}^\theta)^{-1},$$

with

$$\mathbf{D} = \frac{\tau(1-\tau)}{2} \mathbb{E}\left(\frac{1}{\sigma^2}\right) \text{diag}\left(\mathbb{E}\left(\frac{1}{w_i}\right)\right)_{i=1,\dots,n}.$$

The posterior on w_i is a Generalized Inverse Gaussian **GIG**(1/2, α_i , β_i) with :

$$\alpha_i = \left(\frac{(1-2\tau)^2}{2\tau(1-\tau)} + 2 \right)$$

and

$$\beta_i = \frac{\tau(1-\tau)}{2} \mathbb{E}\left(\frac{1}{\sigma^2}\right) \left(y_i^2 - 2y_i \mathbb{E}(q_\tau(x_i)) + \mathbb{E}(q_\tau(x_i)^2) \right).$$

Due to numerical problems, in [1] the authors use the restriction $\tilde{p}(\sigma) = IG(a, b)$. Finding the best a, b is done numerically. Then finding the best a, b is equivalent to maximizing:

$$J(a, b) = (a - N - 10^{-6}) \log(b - \psi(a)) + (b - \gamma) \frac{a}{b} - \delta \frac{a(a+1)}{b^2} - a \log(b) + \log \Gamma(a),$$

with

$$\gamma = -\frac{1 - 2\tau}{2} \sum_{i=1}^n y_i - \mathbb{E}(q_\tau(x_i))$$

and

$$\delta = \frac{\alpha(1 - \tau)}{4} \sum_{i=1}^n \mathbb{E}\left(\frac{1}{w_i}\right) \left(y_i^2 - 2y_i \mathbb{E}(q_\tau(x_i)) + \mathbb{E}(q_\tau(x_i)^2)\right).$$

M-step. Ignoring terms that do not depend on θ , we obtain the lower bound:

$$\begin{aligned} \tilde{\mathcal{L}}(\theta) &= \int \tilde{p}(q_\tau|\theta) \tilde{p}(w) \tilde{p}(\sigma) \log p(y|q_\tau, w, \sigma) p(q_\tau|\theta) d\sigma dw dq_\tau \\ &\quad - \int \tilde{p}(q_\tau|\theta) \log \tilde{p}(q_\tau|\theta) dq_\tau \\ &= \frac{1}{2} \left(\mu_\theta^T \Sigma_\theta^{-1} \mu_\theta - \log |\mathbf{D}^{-1} + K_{x,x}^\theta| \right). \end{aligned} \quad (38)$$

The optimization of $\tilde{\mathcal{L}}$ with respect to θ is done using a numerical optimizer.

Recalling the goal is to compute (36), thanks to (37), we make the approximation:

$$p(q_\tau|x^*, \mathcal{D}_n) \approx \int p(q_\tau|x^*, q_\tau, \sigma, \mathcal{D}_n) \tilde{p}(q_\tau) \tilde{p}(\sigma) \tilde{p}(w) dq_\tau dw d\sigma.$$

Then we obtain

$$p(q_\tau|x^*, \mathcal{D}_n) \approx \mathcal{N}(\bar{q}_\tau(x^*), \mathbb{V}_q(x^*)),$$

where

$$\begin{aligned} \bar{q}_\tau(x^*) &= K_{x_*,x}^\theta K_{x,x}^{\theta-1} \mu_\theta \text{ and} \\ \mathbb{V}_q(x^*) &= k_\theta(x_*, x_*) - K_{x_*,x}^\theta K_{x,x}^{\theta-1} K_{x_*,x}^{\theta T} + K_{x_*,x}^\theta K_{x,x}^{\theta-1} \Sigma_\theta K_{x,x}^{\theta-1} K_{x_*,x}^{\theta T}. \end{aligned}$$

Algorithm 6 Bayesian variational regression

- 1: **Training**
 - 2: **Inputs:**
 $\mathcal{D}_n, \tau, k_{\theta_0}$
 - 3: **Initialize:**
Compute the $n \times n$ matrix $K_{x,x}^\theta$ and $K_{x,x}^{\theta-1}$
 $\theta = \theta_0$
 - 4: **for** $t = 1$ to n_{it} **do**
 - 5: **E-step**
 - 6: Compute $\Sigma_\theta, \mu_\theta, \alpha_i, \beta_i, w_i, (a, b)$
 - 7: **M-step**
 - 8: $\theta = \arg \max \frac{1}{2} \left(\mu_\theta^T \Sigma_\theta^{-1} \mu_\theta - \log |\mathbf{D}^{-1} + K_{x,x}^\theta| \right)$
 - 9: **end for**
 - 10: **Prediction**
 - 11: **Inputs:**
 $\mathcal{X}_{test}, \theta^* = \theta, \mu_{\theta^*}$
 - 12: **for** each point in $x^* \in \mathcal{X}_{test}$ **do**
 - 13: $\hat{q}_\tau(x^*) = K_{x^*,x}^{\theta^*} K_{x,x}^{\theta^*-1} \mu_{\theta^*}$
 - 14: **end for**
-

Finally, as explained in section 5.1.2, the quantile estimator \hat{q}_τ is selected as \bar{q}_τ . Algorithm 6 details the implementation of the VB method.

5.2.3 Computational complexity

E-step.

The complexity of this step is in

$$O(n^3).$$

In fact the algorithm computes Σ_θ that implies inverting a matrix of size $n \times n$.

M-step.

Optimizing $\tilde{\mathcal{L}}$ with a Newton type algorithm costs $O(n^3 + d^3)$ at each iteration (for details refer to the optimization description of (29)). Assuming the starting point θ_{start} is close to an optimal θ^* , based on the same analysis as in section 4.1.3, the whole complexity is in

$$O(L(d^3 + n^3)).$$

Overall complexity. At each iteration of the EM algorithm, the computation cost is $O(L(d^3 + n^3) + n^3)$. The final complexity is obtained by multiplying by the number of iterations n_{it} of the EM algorithm. Thus, the overall complexity is in

$$O(n_{\text{it}}(L(d^3 + n^3) + n^3)).$$

6 Metamodel summary and implementation

In this section we detail our implementation procedure. After providing a summary of the six metamodels in Table 1, we present the packages we used and the hyperparameters we chose (which hyperparameters we set and which hyperparameters we optimized). We then describe the procedure we used to optimize the hyperparameters (optimization strategies and evaluation metrics).

6.1 Summary of the models

Table 1 lists the analytical expressions of the six metamodels, along with the associated underlying quantity.

6.2 Packages and hyperparameter choices

Each method depends on many parameters that can be tuned to improve performance, for example the choice of the kernel function and the value of its parameters for RK, QK and VB or the penalization factor for RK and NN. Here, to limit the computational burden, we chose to optimize only the most critical ones. When possible, for the other parameters, we applied the arbitrary choices and values made by the authors of the original papers. Most changes were made to improve robustness. Below, we describe our experimental settings, also summarized in Table 2.

Nearest Neighbors. We set $d(\cdot)$ as the Euclidean distance and optimized only the size K of the neighborhood.

Random forest. In this case, the only hyperparameter we optimized was the maximum size of the leaves $m_s \in \mathbb{N}^*$. Regarding the number of trees, we noticed that the metamodel needs many more trees than are needed for the estimation of the expectation. In some problems, the metamodel needs up to 5,000 trees to stabilize the variance. Thus, in our experiments we set the number of trees at 10,000 in all cases. We set the number of dimensions considered for the split at the default choice $d/3$. The depth of the tree is not constrained and the splitting rule is based on Eq. 7. We used the R package *QuantregForest* [61].

Neural network. Based on [18], we set the number of hidden layers at one and the transfer function as the sigmoid. The optimization algorithm is a Newton method, we set E_K at $1/2^K$ with $K = 1, 2, 5, 10, 15, 20, 25, 30, 35$ and the number of multistarts to optimize the empirical risk at five. We optimized the number of neurons J_1 in the hidden unit and the regularization parameter $\lambda \in \mathbb{R}_+$. The metamodel is generated using the R package *qrnn* [18].

Regression in RKHS. The kernel was set as a Matérn 5/2. We optimized the length scale parameters $\theta \in \mathbb{R}_+^d$ and the regularization hyperparameter $\lambda \in \mathbb{R}_+$. Optimization (21) is done with the quadratic optimizer *quadprog* [93].

Quantile Kriging. The kernel was set as a Matérn 5/2. The number of repetitions was set according to the total number of observations (see Table 4). We optimized the length scale hyperparameter $\theta \in \mathbb{R}_+^d$ and variance

| Method | Definition of $\hat{q}_\tau(x^*)$ | Related quantity | Complexity |
|--------|--|---|---------------------------------------|
| KN | $y_{([K\tau])}(x^*)$ | The K -nearest points from x^* | $O(nN_{new}(d + \log n))$ |
| RF | $\inf\{y_i : \hat{F}(y_i X = x^*) \geq \tau\}$ | $\hat{F}(y_i X = x^*) = \sum_{i=1}^n \bar{w}_i(x^*) \mathbb{1}_{\{y_i \leq y\}}$ | $O(N\tilde{d}n \log^2 n)$ |
| NN | For a 3 layer NN $g_3(\sum_{j=1}^{J_2} g_2(\sum_{i=1}^{J_1} g_1(\langle w_i^{(h_1)}, x^* \rangle + b_i^{(h_1)}))w_j^{(h_2)} + b_j^{(h_2)})w^{(h_3)} + b^{(h_3)})$ | With $w_i^{(h_1)}, w_j^{(h_2)}, b_i^{(h_1)}, b_j^{(h_2)}, w^{(h_3)}, b^{(h_3)}$, $1 \leq i \leq J_1, 1 \leq j \leq J_2$, minimizing $\frac{1}{n} \sum_{t=1}^n l_\tau(y_t - \hat{q}_\tau(x_t)) + \sum_{j,i} \frac{\lambda}{J_j} \ w_i^{(h_j)}\ ^2$ | $O(M_s L_{\eta^*} n (Jd)^3)$ |
| RK | $\sum_{i=1}^n \alpha_i k_\theta(x^*, x_i) + b$ | With $\alpha = (\alpha_1, \dots, \alpha_n)$ minimizing $\frac{1}{2} \alpha^T K_{x,x} \alpha - \alpha^T \mathbf{y}$ s.t $\frac{\tau-1}{\lambda n} \leq \alpha_i \leq \frac{\tau}{\lambda n}, \forall 1 \leq i \leq n$ and $\sum_{i=1}^n \alpha_i = 0$ and b the τ -quantile of $(y_i - \sum_{j=1}^n \alpha_j k_\theta(x_i, x_j))_{1 \leq i \leq n}$ | $O(n^{7/2} \log(\frac{n}{\epsilon}))$ |
| QK | $K_{x^*,x}^\theta (K_{x,x}^\theta + B)^{-1} \hat{q}_\tau$ | Maximizing the likelihood: $p(\mathcal{Q}_n \mathcal{X}_n)$ $\hat{q}_\tau(x_i) = q_\tau(x_i) + \varepsilon_i \quad \varepsilon_i \sim \mathcal{N}(0, B_{ii})$ | $O(L(d^3 + n^3) + n^3)$ |
| VB | $K_{x^*,x}^\theta K_{x,x}^{\theta-1} \mu_\theta$ | Approached solution that maximize: $p(\mathcal{Y}_n \mathcal{X}_n)$ $y_i = q_\tau(x_i) + \varepsilon \quad \varepsilon \sim \text{ALP}(0, \sigma)$ | $O(n_{it}(L(d^3 + n^3) + n^3))$ |

Table 1: Summary of the metamodels, included the definition of the estimators, the associated numerical quantity and the related computation complexity. In this table L and L_{η^*} are Hessian Lipschitz constants, M_s the number of multistarts (on the weights at hyperparameters fixed), d the input dimension, N_{new} the size of the prediction set, n_{it} the number of iterations for the EM algorithm, $k_\theta(\cdot, \cdot)$ the kernel function, B a diagonal matrix, J the total number of neurons, \tilde{d} the number of covariables considered to find the best splitting point.

| Method | Hyperparameters |
|--------|--|
| KN | number of neighbors $K \in \mathbb{N}^*$ |
| RF | maximum size of the leaves $m_s \in \mathbb{N}^*$ |
| NN | regularization $\lambda \in \mathbb{R}_+$, $J_1 \in \mathbb{N}^*$ |
| RK | regularization $\lambda \in \mathbb{R}_+$, lengthscales $\theta \in \mathbb{R}_+^d$ |
| QK | length scale and variance $\theta \in \mathbb{R}_+^{d+1}$ |
| VB | length scale and variance $\theta \in \mathbb{R}_+^{d+1}$ |

Table 2: Hyperparameters optimized on our benchmark.

hyperparameter $\rho \in \mathbb{R}_+$. QK is implemented in the R package *DiceKriging* [74].

Variational regression. The kernel was set as a Matérn 5/2, the number of EM iterations n_{it} at 50. We optimized the length scale hyperparameter $\theta \in \mathbb{R}_+^d$ and variance hyperparameter $\rho \in \mathbb{R}_+$. The implementation is based on the Matlab code provided in [1].

6.3 Tuning the hyperparameters

In the previous section, we defined the hyperparameters we wanted to optimize for each method. In fact, once the type of metamodel is chosen, the quantile estimator is given by a function $\hat{q}_\Theta : \mathcal{X} \rightarrow \mathbb{R}$ where $\Theta \in \mathbb{R}^v$ are called hyperparameters and v is metamodel dependent. Hyperparameter optimization (also known as model selection) is an essential procedure when dealing with non-parametric estimators. Although \hat{q}_Θ may be very efficient on \mathcal{D}_n , the prediction may perform very poorly on an independent dataset \mathcal{D}'_p . The goal is to find the Θ that provides the best possible generalized estimator. In the following, we present the validation metric used to optimize the hyperparameter values and detail the hyperparameters optimization procedure associated with each method.

6.3.1 Metrics

In the standard conditional expectation estimation, the validation and performance metrics are both based on $\|\hat{m}_\Theta - m\|_{L^2}$, where \hat{m}_Θ is the estimator and m the targeted value. With the quantile estimation procedure the two metrics are no longer the same. The goal is to find \hat{q}_Θ such that $\|\hat{q}_\Theta - q\|_{L^2}$ is as small as possible. However, q is unobserved so the validation metric cannot be based on the L^2 norm. Here we present two metrics able to measure the generalization capacity of a quantile metamodel.

Bayesian metamodels (QK and VB) have their own validation metric, this is the likelihood function that can be maximized with respect to Θ . For quantile kriging, we use (29) while in the variational approach we use (38). The optimal hyperparameters are then:

$$\Theta_{mv}^* = \arg \max_{\Theta} p(\mathcal{Y}_n | \mathcal{X}_n, \Theta). \quad (39)$$

The second metric available for all metamodels is k -fold cross-validation associated with the pinball loss. The metric can be computed as follows. First, the data are split into k parts, then the model is trained on \mathcal{D}_{-j} , the training set without the j -th part. The performance is evaluated on the remaining part \mathcal{D}_j . As the quantile minimizes the pinball loss (on \mathcal{D}_j), the evaluation metric is

$$E_{cv}(\hat{q}_\tau^\Theta) = \frac{1}{k} \sum_{j=1}^k \frac{1}{n'_j} \sum_{i=1}^{n'_j} l_\tau((y_i - \hat{q}_\tau^\Theta(x_i))), \quad (40)$$

where n'_j is the number of observations in each fold. The optimal cross-validation hyperparameters are then:

$$\Theta_{cv}^* \in \arg \min_{\Theta} E_{cv}(\hat{q}_\tau^\Theta).$$

In our experiments, we chose $k = 5$ to limit the computational cost. However, we observed empirically that choosing a larger k did not substantially modify the performances of the metamodels. Although cross-validation is available for QK, we chose to stay in the spirit of the methods and to only use likelihood to select hyperparameters. Our choice is supported by [5] that does not show a clear improvement using cross validation instead of maximum likelihood techniques.

6.3.2 Hyperparameters optimization procedure

Both likelihood functions come with analytical derivatives, enabling the use of gradient-based algorithms. However, since both functions are multi-modal, multi-start techniques are necessary (and generally very efficient, see [39]) to avoid being trapped in local optima. To account for the increasing difficulty of the optimization task with the dimension while limiting the computational cost, we ran $n_{\text{start}} = 20d$ optimization procedures from different starting points θ_{start} and chose the set of starting points based on a *maximin* Latin hypercube design.

For QK, the BFGS algorithm is used to optimize (29). For VB, two derivative-based optimizers are used alternately for the E- and M-steps. Since each step may lead the algorithm toward a local minimum, we chose to apply the multi-start strategy in the entire EM procedure.

Optimization of the cross-validation metric (40) is done under the black box framework, since no structural, derivative or even regularity information is available. Hence, all optimizations are carried out using the branch-and-bound algorithm named Simultaneous Optimistic Optimization (SOO) [64]. SOO is a global optimizer, hence robust to local minima.

We used [91] to parallelize the computations.

6.3.3 Oracle metamodels

Each method presented in this paper is a trade-off between power and the difficulty of finding good hyperparameters. A good method should be powerful (i.e. provide flexible fits) but easy to tune. In order to assess the strengths and weaknesses of the hyperparameter tuning methods in addition to standard metamodels, we provide what we call the *oracle metamodels* for each problem. Instead of using the cross-validation or likelihood metric, the oracle tunings are directly based on the evaluation metric

$$E_{L^2}(\hat{q}) = \sum_{i=1}^{n_{\text{test}}} (\hat{q}^\Theta(x_i) - q(x_i))^2, \quad (41)$$

where n_{test} is the size of the test set. In a sense, they provide an upper bound on the performance of each method. This allows us to show which metamodels have the potential to tackle the problems and which are intrinsically too rigid or make poor use of information. In addition, this allows us to directly assess the quality of the validation procedure.

7 Benchmark design and experimental setting

Many factors can affect the efficiency of methods to estimate the right quantile. For our benchmark system, we considered five models or test cases to evaluate the performance of the six metamodels. We decided to focus primarily on the dimensionality of the problem, the number of training points available, the signal-to-noise ratio defined as

$$\text{SNR} = \frac{\mathbb{V}_X(\mathbb{E}(Y|X))}{\mathbb{E}_X(\mathbb{V}(Y|X))},$$

and the pdf value at the targeted quantile for test cases in which the distribution shape and the distribution spread (i.e. level of heteroscedasticity) can vary significantly. Our two objectives were to:

1. discover if there is a single best method for all factors variations considered or specific choices depending on the configuration at hand, and
2. assess the performance of the quantile regression, and in particular, the configurations for which the current state-of-the-art is insufficient.

A full 3D factorial experimental design was used to analyze the efficiency of the metamodels, the 3 factors being the test case (4 test cases), the number of training points (4 levels) and the quantile order (0.1, 0.5 and 0.9). We used part of this complete design to focus our analyses on the characteristics of the test cases (dimension, pdf shape and heteroscedasticity).

7.1 Test cases and numerical experiments

Test case 1 is a 1d toy problem on $[-1, 1]$ defined as

$$Y_x = 5 \sin(8x) + (0.2 + 3x^3)\xi,$$

with $\xi = \eta \mathbb{1}_{\eta \leq 0} + 7\eta \mathbb{1}_{\eta > 0}$ where $\eta \sim \mathcal{N}(0, 1)$.

The signal-to-noise ratio is $\text{SNR} \approx 0.5$, it is considered as small. The pdf value $\tilde{f}(x, q_\tau)$ varies substantially according to x for all q_τ . Indeed, on the interval $[-0.5, 0.3]$, for all values of q_τ , the pdf is very large (almost equals to $+\infty$) because the variance of the distribution on this interval is very small. In contrast, for $\tau = 0.9$ (resp. $\tau = 0.1$) the pdf is very small in the interval $[0.6, 1]$ (resp. $[-1, -0.6]$). In $[0.6, 1]$ the pdf of the 0.9-quantile is equal to the pdf of the 0.9-quantile of a normal distribution of variance $49(0.2 + 3x^3)^2$ that is, for example, approximately equal to 0.01 for $x = 0.9$. Because of these important variations, we consider the values of the pdf for all quantiles of interest as variable.

Test case 2 is a 2d toy problem on $[-5, 5] \times [-3, 3]$ based on the Griewank function [29], defined as

$$Y_x = G(x)\xi,$$

with

$$G(x) = \left[\sum_{i=1}^2 \frac{x_i^2}{4000} - \prod_{i=1}^2 \cos\left(\frac{x_i}{\sqrt{i}}\right) + 1 \right]$$

and $\xi = \eta \mathbb{1}_{\eta \leq 0} + 5\eta \mathbb{1}_{\eta > 0}$ where $\eta \sim \mathcal{N}(0, 1)$.

The signal-to-noise ratio is $\text{SNR} = 0$, it is consider as small. The pdf value $\tilde{f}(x, q_\tau)$ varies substantially according to q_τ . Indeed, for $\tau = 0.1$ (resp. $\tau = 0.9$) the pdf is small (resp. very small), more precisely the pdf of the 0.1-quantile (resp. 0.9-quantile) is equivalent to the pdf of the 0.1-quantile (resp. 0.9-quantile) of a normal distribution of variance G^2 (resp. $25G^2$) with G that takes values in $[0, 2]$. That implies $\tilde{f}(x, q_\tau)$ varies with respect to x as well. Note that close to $x = (0, 0)$ the pdf is very large because of the very small variance of the associated distribution. The pdf at the 0.5-quantile is equal to the pdf at the median of a normal distribution of variance G^2 , that we consider as large.

Test case 3 is a 1d toy problem based on the Michalewicz function [29] on $[0, 4]$, defined as

$$Y_x = -2 \sin(x) \sin^{30} \left(\frac{x^2}{\pi} \right) - \frac{0.1 \cos(\pi x/10)^3}{|-\sin(x) \sin^{30} \left(\frac{x^2}{\pi} \right) + 2|} \xi^2,$$

with $\xi = 3\eta \mathbb{1}_{\eta \leq 0} + 6\eta \mathbb{1}_{\eta > 0}$ where $\eta \sim \mathcal{N}(0, 1)$.

The signal-to-noise ratio is $\text{SNR} \approx 0.04$, we consider it as small. The pdf value $\tilde{f}(x, q_\tau)$ varies substantially according to q_τ and x . The conditional distribution of this problem is not a classical one but it is close to the distribution of $-\chi^2$ with one degree of freedom. It implies at x fixed, the pdf value increases with τ . For the 0.1-quantile (resp. 0.9-quantile) the mean of the pdf in the interval $[0, 2.5]$ is approximately 0.05 (resp. 9.1). According to x the pdf varies as well. For $x \in [3.5, 4]$ the variance of the conditional distribution is very small thus the pdf near the quantiles of interest is very large. The value of the pdf at the 0.5-quantile is between the value of the pdf at the 0.1 and 0.9 quantile. Thus we consider the pdf at $q_{0.1}$ as globally small and at $q_{0.5}$ and at $q_{0.9}$ as globally large

Test case 4 is a 9d toy problem based on the Ackley function [3] on $[-1, -0.7] \times [0, 1] \times [-0.7, -0.3] \times [0.5, 1] \times [-1, -0.5] \times [-3, -2.6] \times [-0.1, 0] \times [0, 0.1] \times [0, 0.8]$, defined as a function

$$Y_x = 30 \times A(x) + R(x) \times \xi$$

with

$$A(x) = a \exp \left(-b \sqrt{\frac{1}{9} \sum_{i=1}^9 x_i^2} \right) - \exp \left(\frac{1}{9} \sum_{i=1}^9 \cos(cx_i) \right) + a + \exp(1), \quad (42)$$

and

$$R(x) = 3A(x_2, x_3, \dots, x_9, x_1), \quad (43)$$

with $a = 10$, $b = 2 \times 10^{-4}$, $c = 0.9\pi$ and ξ follows a log-normal distribution of parameters $(0, 1)$.

The signal-to-noise ratio is $\text{SNR} \approx 2.3$, it is consider as large. The pdf value $\tilde{f}(x, q_\tau)$ varies substantially according to q_τ and x . Indeed, for $\tau = 0.1$ (resp. $\tau = 0.9$) the pdf is large (resp. very small). At the conditional distribution of this test case is a log-normal distribution of parameters $(\log R(x), 1)$. The expectations of the pdf value at different values of τ are $\mathbb{E}(\tilde{f}(\cdot, q_{0.1})) = 0.1$, $\mathbb{E}(\tilde{f}(\cdot, q_{0.5})) = 0.04$ and $\mathbb{E}(\tilde{f}(\cdot, q_{0.9})) = 0.005$. Thus the pdf at $q_{0.1}$ is consider as large, the pdf at $q_{0.5}$ and $q_{0.9}$ are consider as small and very small.

To provide a better intuition about this problem, we plotted what we call the marginals. For all dimensions except the j -th, the values of the input are fixed to $x_{-j} \in \mathbb{R}^8$ and the j -th dimension varies. Figure 6 represents the evolution of the quantiles w.r.t. the j -th dimension for two different x_{-j} and for $j = 1, \dots, 9$. In particular it shows that the difference between the 0.1 and 0.9-quantile depends significantly on x .

Note that on those four toy problems, the random term ξ is defined such that the resulting distribution of Y would be strongly asymmetric. As ξ is also multiplied by a factor that depends on x , the distribution of Y is also heteroscedastic. The first three toy problems are represented in Figure 5 and some illustrations of the fourth test case are available Figure 6.

Test case 5 is based on the agronomical model SUNFLO, a process-based model which was developed to simulate sunflower grain yield (in tons per hectare) as a function of climatic time series, environment (soil and climate), management practices and genetic diversity. The full description of the model is available in [19]. In the regression model we consider \mathcal{X} corresponding to nine macroscopic traits that characterize the sunflower variety. Although SUNFLO is a deterministic model, for each simulation the climatic time series are randomly chosen within a database containing 190 years of weather records, which makes the output stochastic (see also [65] for more details).

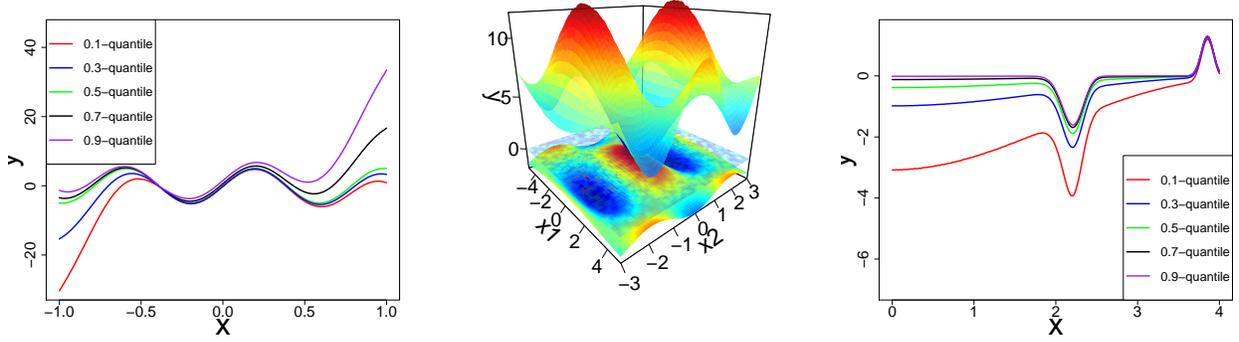


Figure 5: Illustration of the test cases (left: test case 1, center: test case 2, right: test case 3). For test case 1 and 3 the 0.1, 0.3, 0.5, 0.7, 0.9 -quantiles are represented. For test case 2, only the 0.1, 0.5, 0.9-quantiles are represented.

The signal-to-noise ratio is $SNR \approx 0.1$. The pdf value $\tilde{f}(x, q_\tau)$ varies substantially according to q_τ , more precisely $\mathbb{E}(\tilde{f}(\cdot, q_{0.1})) = 0.17$, $\mathbb{E}(\tilde{f}(\cdot, q_{0.5})) = 0.15$ and $\mathbb{E}(\tilde{f}(\cdot, q_{0.9})) = 0.05$. Thus the pdf at $q_{0.1}$ and at $q_{0.5}$ are consider as large, and the pdf at $q_{0.9}$ is consider as small.

In addition, the shape of the distribution varies significantly with x .

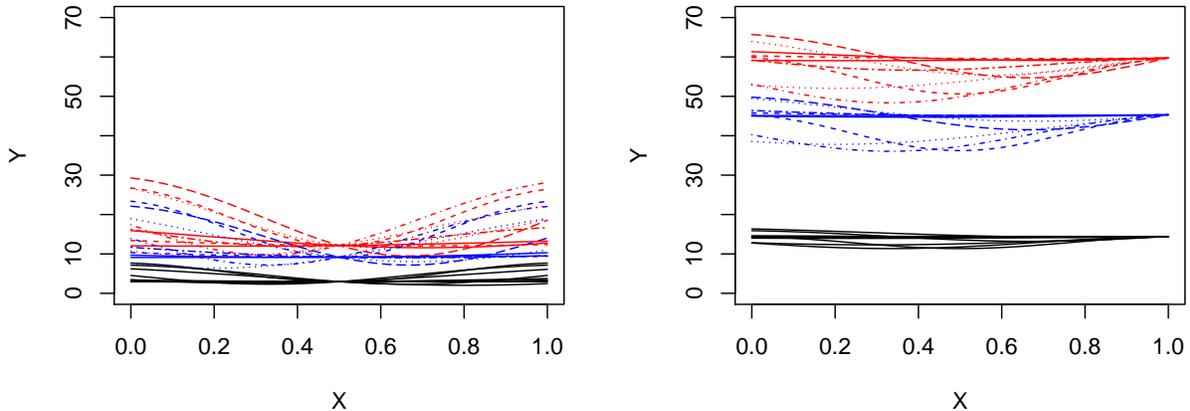


Figure 6: Illustration of some marginals of Test case 5. The conditional quantiles of order 0.9 (resp. 0.1) are represented in red (resp. in blue). The black curves represent the difference between the 0.1-conditional quantiles and the 0.9-conditional quantiles so that to measure the level of heteroscedasticity. To the right the noise level is low *i.e* between 0 and 10 while to the left the noise level is higher *i.e* between 10 and 20.

Numerical experiments. On all problems, we consider four sample sizes. Those sizes depend on the dimension and are empirically chosen so that the smallest size corresponds to the minimal information required by the metamodels to work and the largest size is chosen keeping in mind the potentially high cost of simulators. Besides, our focus is on computer experiments, where data sizes rarely exceed thousands of points. For the $1d$ problems, the points are generated on a uniform grid. For the $2d$ and $9d$ problems, the observations are taken on a Latin hypercube design optimized for a *maximin* criterion to ensure space-filling [33]. The same samples are used by all methods except QK, as it requires repetitions. For QK, the number of distinct points and number of repetitions depends on the budget. The different sample sizes are reported in Table 4. Finally, for each sample size and problem, 10 samples are drawn in order to assess robustness with respect to the data.

| | | | | | | |
|-------------------------------------|--------------|-------------|-------------|---------------------|-------------|--------|
| | | Test case 1 | Test case 2 | Test case 3 | Test case 4 | Sunflo |
| Dimension | | 1 | 2 | 1 | 9 | 9 |
| Heteroscedasticity | | very strong | very strong | very strong | strong | weak |
| Shape variation | | very strong | weak | weak | weak | strong |
| pdf value near the τ -quantile | $\tau = 0.1$ | variable | small | globally very small | large | large |
| | $\tau = 0.5$ | variable | large | small | small | large |
| | $\tau = 0.9$ | variable | very small | very large | very small | small |

Table 3: Summary of the characteristic of the problems.

| Dimension | Data size (no repetitions) | | | | Data size (with repetitions) | | | |
|-----------|----------------------------|-----|------|------|------------------------------|---------|----------|----------|
| | 1 | 40 | 80 | 160 | 320 | 5 (8) | 10 (8) | 10 (16) |
| 2 | 100 | 200 | 400 | 800 | 10 (10) | 20 (10) | 25 (16) | 40 (20) |
| 9 | 250 | 500 | 1000 | 2000 | 25 (10) | 50 (10) | 100 (10) | 100 (20) |

Table 4: Data sizes for the different problems. The number in parentheses are the number of repetitions for QK.

7.2 Structuration between the questions and the numerical setting

Factors. Three factors are explicit in our benchmark system: the number of training points, the problem dimension and the quantile level. The other factors depend on the characteristics of the problem concerned: shape variation, pdf value at the quantile, level of heteroscedasticity, signal-to-noise ratio. For all four test cases, we consider three quantile levels: 0.1, 0.5 and 0.9. Note that due to the asymmetry of the problems, learning for the 0.1 and 0.9 quantiles is not equivalent in terms of difficulty. Indeed, when the response is heteroscedastic (a variance/spread depending on x) and/or when the shape of \mathbb{P}_x varies in x , the pdf $\tilde{f}(x, q_\tau)$ may vary in x as well. Intuitively, quantiles with large pdf values are easier to learn, as the data points may be closer to them. Figure 7 illustrates this effect. Table 3 summarizes the characteristics of our design concerning the number of training points with respect to the dimension of the test case. To make our results easier to analyze, we divided the problems into groups that allow us to focus on subsets of factors.

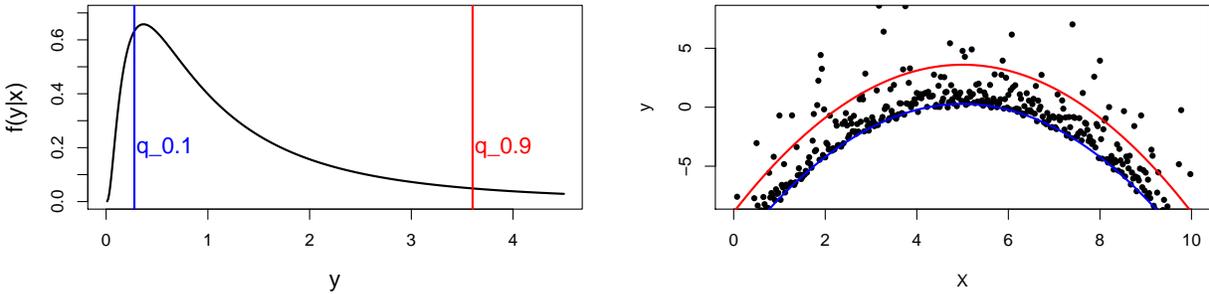


Figure 7: Left: log-normal density function with $\mu = 0$ and $\sigma = 1$. Right: a sample generated by the function $f(x) = \xi - (x - 5)^2/2$, with ξ following the density represented on the left. The 0.9- (resp. the 0.1-) quantile is represented in red (resp. in blue). One can notice that more information is available in areas with large pdf (i.e. for the 0.1-quantile) than areas with small pdf.

Focus 1: is there a universal winner? To provide a universal ranking of the methods, we use all test cases, training points and quantile levels. As highlighted in Table 3, we created a set of different problems representative of a large number of characteristics that could be met dealing with any quantile regression problem. Note that our benchmark system is slightly biased towards small-dimensional problems, since only three-fifths of the cases have a dimension higher than two.

Focus 2: what behavior with respect to the dimension, the number of training points, signal-to-noise ratio and pdf value? To analyze the effects of these factors on the performance of the methods, we combine toy problems 1, 2, 3,

4 and the SUNFLO model. Note that once the pdf value is taken as an explanation variable, toy problem 1 is excluded from the group because the pdf value near all the studied quantiles cannot be classified as large or small.

7.3 Performance evaluation and comparison metrics

Assessing the performance of quantile regression is not an easy task when only limited data are available. Here, since we are considering toy problems (except for SUNFLO, in that case the true quantile values are taken as the quantiles of the 190 years of weather records), the true quantile values can be approximated with precision, so we can evaluate the L^2 error for each emulator. The value of the criterion is provided by (41).

We chose $n_{\text{test}} = 250$ for the $1d$ problems and $n_{\text{test}} = 4000$ for the others. Now, since the problems vary in difficulty and in their response range (Figure 5), $E_{L^2}(\hat{q})$ cannot be aggregated directly over several problems or configurations. To do so, we normalize this error by the error obtained by a constant model (the constant being taken as the quantile of the sample):

$$E_{cq}(\hat{q}) = \sqrt{\frac{E_{L^2}(\hat{q})}{E_{L^2}(\text{CQ})}} \times 100, \quad (44)$$

where CQ stands for constant quantile.

As an alternative criterion, we consider the ranks of the metamodels based on their L^2 error. Although ranks do not provide information regarding the range of errors, they are insensitive to scaling issues, which makes aggregation between configurations more sensible. They allow us to assess whereas any method consistently outperforms others, regardless of overall performance.

8 Results

8.1 Focus 1: overall performance and ranks

First, we consider the overall performance and ranks, integrated over all runs. We have considered 5 test cases, for each test case we have generated 4 sizes of training sets, for each test case and sample size 10 samples are drawn and for each of these occurrences we have estimated 3 different conditional quantiles. Thus we have $5 \times 4 \times 10 \times 3 = 600$ experiments. They are shown as boxplots in Figure 8. Based on these ranks (Figure 8, left), VB appears to be the best solution since it is ranked either 1st or 2nd in 50% of the problems. RK is in second position, its median is the same as VB but it is generally ranked between 2nd and 3rd. In addition RK seems slightly less risky than VB in the sense that it is almost never ranked 5th or 6th. KN is the worst since its median rank is equal to five. However, all boxplots range

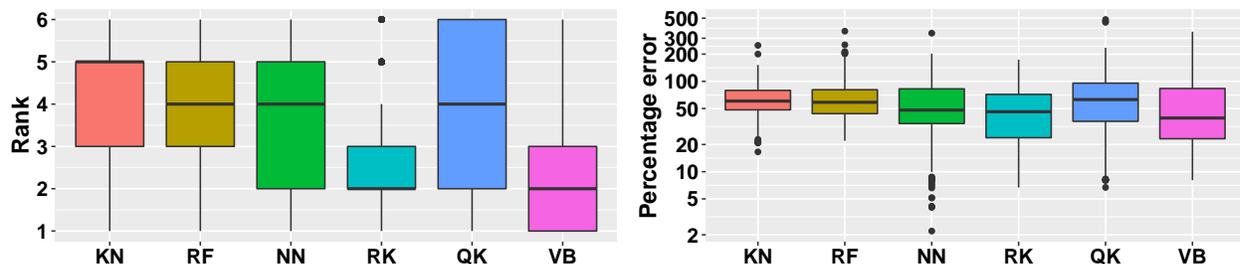


Figure 8: Boxplots of ranks and E_{cq} error over the entire benchmark. Note that for clarity, the right boxplots do not contain the error of the median for the toy problem 2.

between 1 and 6, indicating that no method is outperformed by another on all problems. This finding is reinforced by the performance boxplots (Figure 8, right), where all median performances are similar (VB and RK being the best and QK the worst which means QK may be very bad sometimes), and the variance is very large. Indeed, the errors range from 2% for NN (of the error achieved by a constant metamodel) to 500% for QK, all methods experiencing cases with more than 100% error (i.e. situations where they are worse than the constant metamodel).

8.2 Focus 2: dimension, number of training points and pdf value

8.2.1 Performance according to the constant quantile

In this section, we analyze the performance of the methods with respect to the pdf value and the number of points.

Sample size: Figure 9 shows the performances of the methods grouped according to the size of the sample. As expected, the performances increase with the size of the sample. For size 1 ($n \approx 50d$), the distribution of E_{CQ} of all the metamodels is almost centered around 100%, implying that these correspond to limit cases for quantile regression since the metamodels do not outperform the constant metamodel (although in some cases the error is as small as 40%). For size 4 ($n \approx 300d$), the median performance is roughly 50% (twice as accurate as the constant metamodel). BV, RK and especially NN experience situations with very accurate models. However, all the methods also experience bad performances (error greater than 100%) in the large sample regime. Unfortunately, from Figure 9 we can conclude that no method is sufficiently robust in all cases.

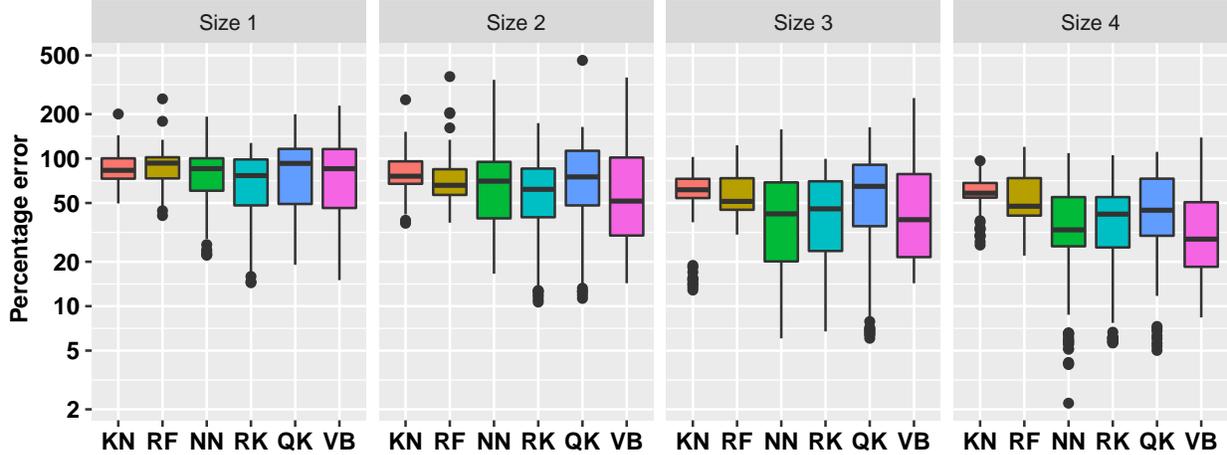


Figure 9: Error according to the sample size

Signal-to-noise ratio. Figure 11 groups performance with respect to the signal-to-noise ratio. According to the figure the performance depends to a great extent on the signal-to-noise ratio. Dealing with a high signal and considering the performance, there is a clear difference between the statistical methods (KN, RF) and the four others (RK, NN, VB, QK) while this difference is not visible in the small signal setting. The impact of the signal-to-noise ratio is strong on the performance. Dealing with a high signal, the median of the performance for NN, RK, QK, VB is close to 20% with a small variance while if the signal is small, the error is larger (the median is above 50% for all methods) and the variance is larger.

In the following we focus our interest only on cases with small signal-to-noise ratio. Indeed as the performance of RK, NN, QK and VB are close to each others, we think we do not have enough experiments to extract patterns.

Pdf value with small signal-to-noise ratio. Figure 11 groups performance with respect to sample size (either small, i.e. level 1 and 2 or large, i.e. level 3 and 4) and pdf value (according to Table 3). According to the Figure 10, the performance depends to a great extent on the pdf value in the neighborhood of the targeted quantile. With a small n , the pdf value has no significant impact on the median of the performance (except for VB) but it does have an impact on the lower bound of the error. More precisely, the median of the error does not depend on the pdf value in the case of a small pdf but sometimes the metamodel errors are sensibly smaller when the pdf is large. With large samples, both the median and the lower bound of the error depend on the pdf value. Metamodels may be very good when the pdf is large, for example 20 times better than the constant metamodel for NN whereas the error appears to have a lower bound when the pdf is small even with large n . In addition, for a problem with a small n and a large pdf, the performance is similar to the performance for problems with a large n and a small pdf (Figure 11, see the two columns in the center).

8.2.2 Rank in the context of a low signal-to-noise ratio

Pdf value. Figure 12 shows clearly that when the pdf is large, VB is the best model while when the pdf is small (and the problem is heteroscedastic), VB is less good than RN, RK and KN. This observation is supported by Figure 11 which reveals a strong contrast between the performance of the VB method. QK is poor in both cases, whereas RK

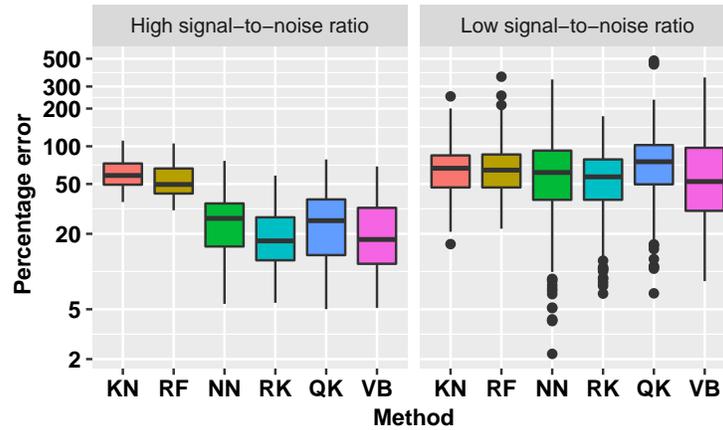


Figure 10: Error according to the signal-to-noise ratio.

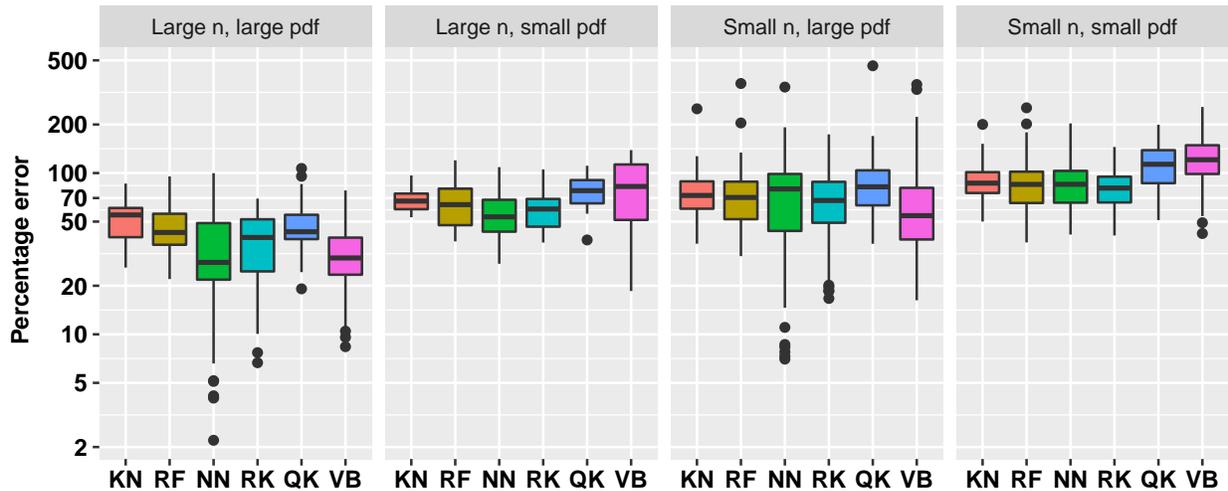


Figure 11: Error according to the size of the training set and the pdf value.

performs comparatively better with small pdf.

Sample size. Figure 13 shows that the number of points has a major impact on the ranking of some methods. The ranking of QK and VB is relatively insensitive to the size of the sample. The other methods are less distinguishable when the sample size is small than when it is large. With a small sample KN, RF, NN and RK are comparable, whereas when the sample size increases, NN and RK clearly outperform KN and RF. For the largest size, NN is slightly better than VB.

Dimension. Figure 14 groups performance based on dimension. The first contrast is the permutation between RK and NN. With a small dimension, RK is better than NN but the relative performance of NN increases w.r.t. the dimension. With small dimensions, RF and KN are comparable, but with high dimensions, RF outperforms KN.

High dimension, small pdf. Figure 15 shows an extreme case in which the pdf is small but the dimension is high. With a small n , the best method is clearly RF followed by RK and KN. VB and QK are not well ranked. With a larger n , as mentioned above, NN and VB are better but with large variance, while RF, RK and KN rank less well.

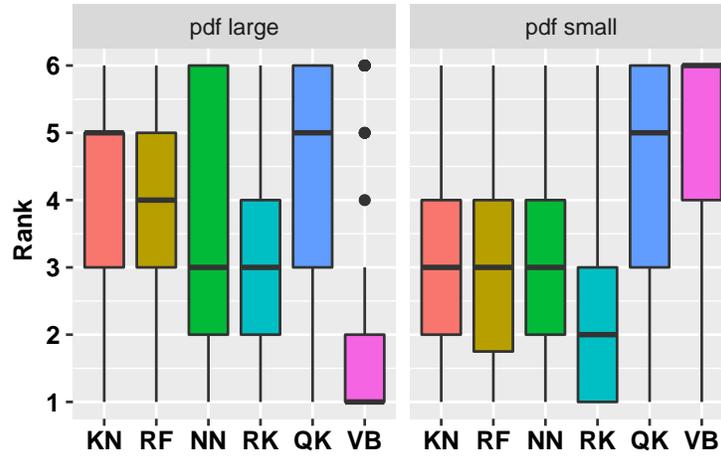


Figure 12: Rank according to the pdf value

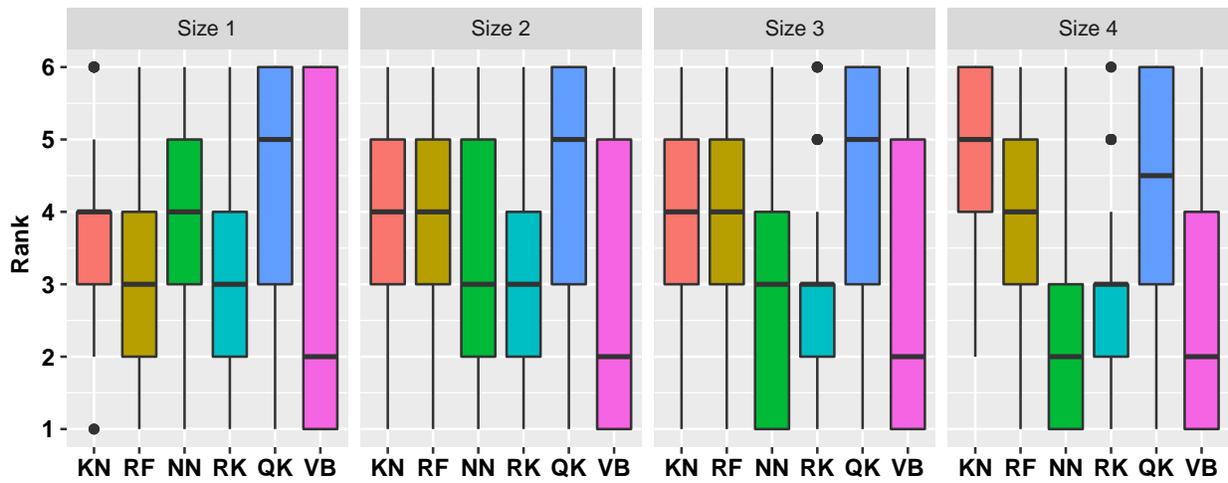


Figure 13: Rank according to the size of the sample

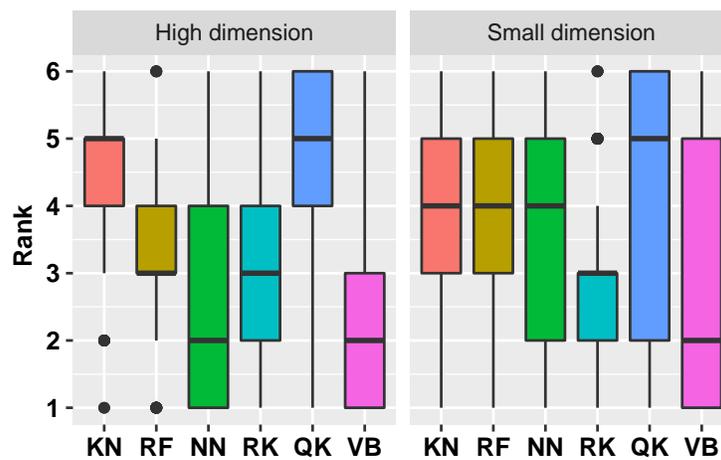


Figure 14: Rank according to the dimension

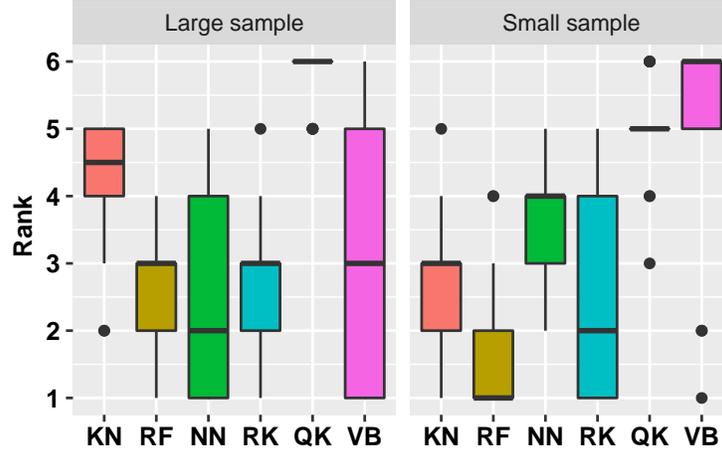


Figure 15: Rank associated to the case where little information is available, *i.e* high dimension and small pdf

9 Extensions and open questions

9.1 Effect of hyperparameter tuning

In the following we define

$$\Delta E(\hat{q}_\tau^\Theta) = E_{cq}(\hat{q}_\tau^\Theta) - E_{cq}(\hat{q}_\tau^{\Theta^*}),$$

the performance gap between the regular metamodel and its oracle performance (the loss in performance between actual hyperparameter tuning and oracle tuning). Figure 16 gives the average values of ΔE aggregated respectively over all problems and only aggregated over the problems with a large pdf, and considering the effect of dimension and sample size. In high dimension, the easiest methods to tune are KN, NN, and RF. In our study, KN and RF have a single hyperparameter to tune regardless of the dimension, and NN has two. This is clearly an advantage in terms of robustness in high dimension. The other methods are kernel-based and require the tuning of at least $d + 1$ hyperparameters. This consistently affects RK and QK, but affects VB only in the case of small pdf, while it is the most stable method in the other cases.

With a small dimension, all the methods have roughly the same number of hyperparameters. The most noticeable change compared to the case of a high dimension is the good performance of RK, while NN becomes comparatively the most difficult method to train.

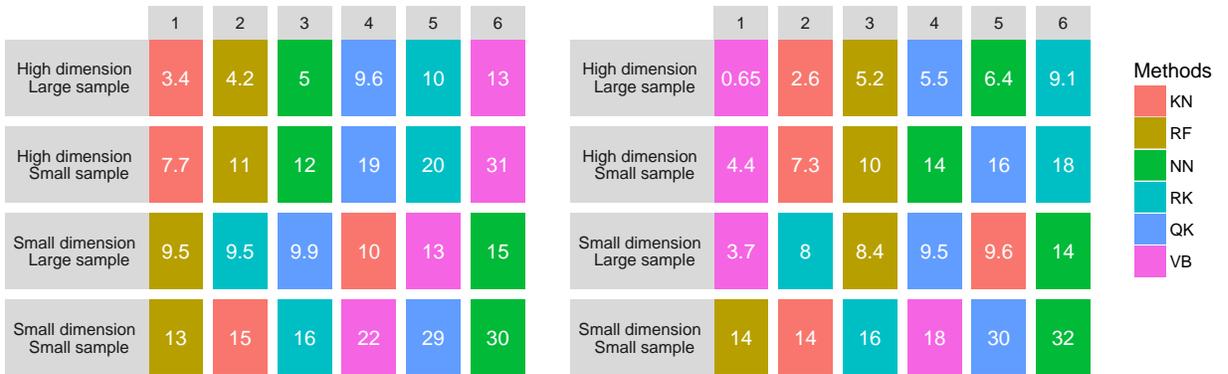


Figure 16: Average ΔE aggregated over all problems (left) and over the problems with a large pdf only (right), arranged in increasing order. For each method the rank is provide on the top of each figure.

9.2 On the methods behavior

Statistical order methods. As presented on Figure 10 this methods are not relevant on high signal-to-noise regime. A possible explanation is the difficulty to fit smooth variations with high amplitudes with a piecewise constant model.

Concerning the low signal-to-noise regime, it is clear from Figure 14 that KN performs poorly in a high dimension. This may be due to the irrelevance of the Euclidean distance when there is a significant increase in dimension. RF clearly outperforms KN in this situation, as it is able to produce better neighborhoods than the Euclidean distance. Overall, (compared with the other methods), RF performance increases with dimension. This may be due to the fact that it has fewer hyperparameters to tune.

Functional methods. As presented on Figure 10 this methods are relevant on high signal-to-noise regime.

Concerning the low signal-to-noise regime, Figure 13 shows that NN works poorly in a small sample setting, but it is one of the best methods when the sample is large. This result reflects the high flexibility of NN. Too much flexibility leads to overfitting when the sample is small. In contrast, when the number of points is large, NNs are able to fit the data very well (e.g. Figure 9, Size 4). According to Figures 9 and 13, RK is a robust method. Its robustness in both small and large data settings can be attributed in part to the selected kernel. If the selected kernel is sufficiently smooth (here continuous and derivable), the resulting metamodel cannot produce instable results. However, it seems (Figure 9, Size 3 and 4) that this lack of flexibility may affect the performance with an increase in the size of the data set. In this case, more flexible methods like NN may outperform RK. The contrast between RK and NN shown in Figure 14 can be explained by the level of difficulty associated with each method involved in finding good hyperparameters (as explained above).

Bayesian models. As presented on Figure 10 this methods are relevant on high signal-to-noise regime.

Dealing with low signal-to-noise ratio the QK method under-performs comparing to others. One possible explanation is the erroneous assumption in Equation (30) that the noise is centered, which is more critical for extreme quantiles. Another possible explanation lies in the small number of replica. The local inference (that uses statistical orders) is biased and in a low signal-to-noise regime it has high variance. In addition, the increasingly bad performance of QK with an increase in dimension (Figure 9) is likely a consequence of the fact that empty areas become larger in high dimensions.

VB is one of the best methods presented in our paper. Figures 12 and 16 show that VB is also the most dependent on the pdf value. When the pdf is large, it is the best method whereas when the pdf is small (and the shape of the distribution and/or the variance depend on the input), it may be the worst. The explanation lies in the philosophy of the model. In the case of RN and RK, the model complexity (i.e. smoothness) is almost entirely related to the regularity of parameter λ that is selected by cross-validation. Hence, the model cannot excessively overfit and cannot perform very poorly. With Bayesian methods, the regularization is included in the model hypothesis: in our setting, the quantile is assumed to be a Gaussian process with covariance function $k_\theta(\cdot, \cdot)$, so θ performs the regularization. We observed that if the local quantity of information (roughly the product of the number of points times the pdf value in the neighborhood of the quantile) is too small comparing to the information needed to fit well the quantile, the metamodel tends to interpolate the available data. When sufficient information is available, the optimization of the marginal likelihood provides θ values that allow a good trade-off between flexibility and smoothness. This is likely the reason why VB is easily beaten by RN and RK when the pdf is small.

9.3 Varying shape and heteroscedasticity.

If the shape of $\mathbb{P}_x(Y)$ or the variance of Y_x (heteroscedasticity) vary w.r.t. x , then $\tilde{f}(x, q_\tau)$ may vary in x . Figure 17 illustrates the ability of RF, RK and VB to estimate quantiles of a distribution with a strongly varying shape. In this problem, as depicted in Figure 17 (top row), the quantiles are not perfectly estimated but the metamodels provide good indications about the shape of the true distribution. However, as can be seen in Figure 17 (bottom row), the methods can present strong instabilities. Here, for a sample virtually indistinguishable from the one leading to accurate estimates, the median estimates largely overestimate the true values for large x values. Such instabilities can be partly imputed to the difficulty of the task. However, this is also because no method is actually designed to deal with strongly varying pdf, as we explain below.

An ideal method would be almost interpolant for a very large pdf but only loosely fit the data when the pdf is small. Indeed, if the pdf is very large then the output is almost deterministic, thus the metamodel should be as closer as possible to the data. In the small pdf case, a point does not provide a lot of information. Information should be

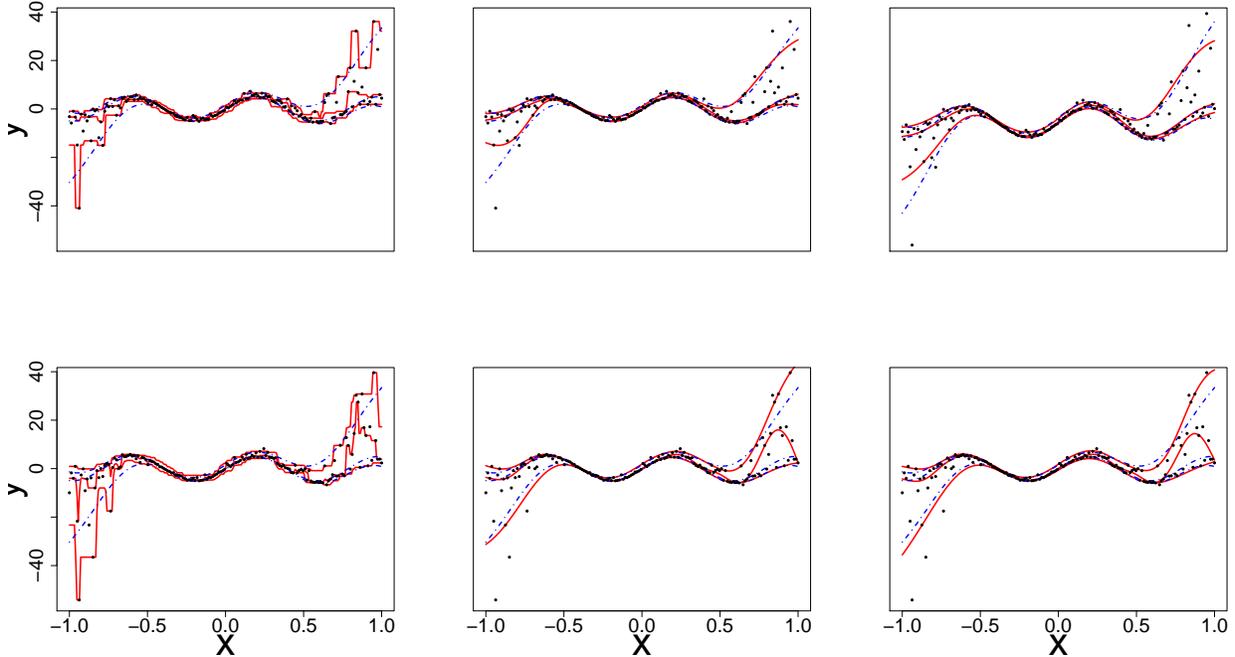


Figure 17: Quantiles estimates using RF (left), RK (middle), VB (right) for two 160-point samples (top and bottom rows, resp.) of the toy problem 1. Dots: observations; plain red lines: metamodels for the 0.1, 0.5, 0.9 quantile estimates; dotted blue lines: actual quantiles.

extracted from a group of points, that means the metamodel must not interpolate the data. However, most of the methods presented here rely on a single hyperparameter to tune the trade-off between data fitting and generalization: the number of neighbors for KN, the maximum size of the leafs for RF and the penalization factor for NN and RK. As a result, the selected hyperparameters are the ones that are best on average. Theoretically, this is not the case for the Bayesian approaches: QK accounts for it *via* the error variance σ_i^2 computed by bootstrap, and the weights w_i (Eq. 35) allow VB to attribute different "confidence levels" to the observations. However in practice, both methods fail to tune the values accurately, as we illustrate below. Figure 18 shows the three quantiles of toy problem 3 and their corresponding RF, RK and VB estimates. For $\tau = 0.1$ in particular, the pdf ranges from very small (x close to 0) to very large (x close to 4). Here, RF and RK use a trade-off that globally captures the trend of the quantile, but cannot capture the small hill in the case of large x . Inversely, VB perfectly fits this region but dramatic overfitting occurs on the rest of the domain.

Finally, Figure 19 illustrates that this is not an issue of hyperparameter tuning. For each method, we show the oracle estimate, a tuning that tends to underfit and another that tends to overfit. One can see that no tuning is entirely satisfactory, since capturing the region with high pdf leads to overfitting on the rest of the domain and vice-versa.

We believe that further research is necessary to obtain estimators that intrinsically account for strong heteroscedasticity and varying shape. One possible direction is the use of stacking, in the spirit of [81]. Under the stacking framework the final estimator could be

$$\hat{q}(x) = \sum_{i=1}^N g_i(x) \hat{q}_{\theta_i}(x),$$

where $\{\hat{q}_{\theta_i}\}_{1 \leq i \leq N}$ is a set of metamodel and $\{g_i(x)\}_{1 \leq i \leq N}$ is a set of weight functions. Choosing $\{\hat{q}_{\theta_i}\}_{1 \leq i \leq N}$ such that they correspond to different pdf values might provide more flexible estimates.

9.4 On the non-crossing of the quantile functions

While the quantile functions (for different quantile levels) may obviously never cross, unfortunately, their estimators may not always satisfy this property. This is a well-known issue against which none of the methods presented here is immune.

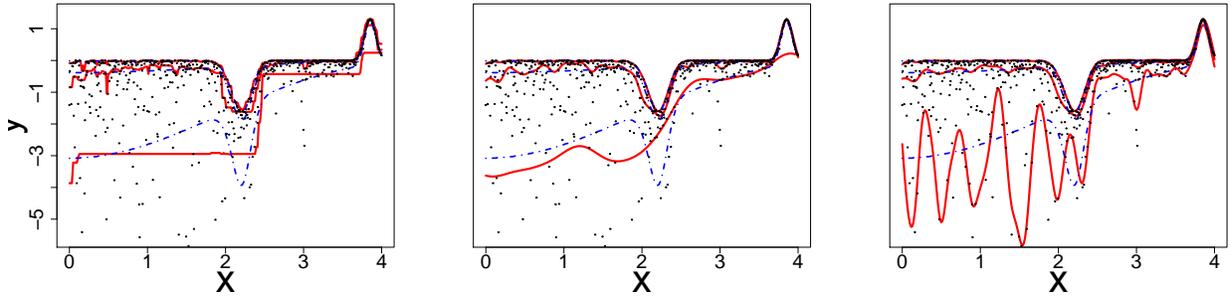


Figure 18: Quantiles estimates using RF (left), RK (middle), VB (right) for a 640-point sample of the toy problem 3. Dots: observations; plain red lines: metamodels for the 0.1, 0.5, 0.9 quantile estimates; dotted blue lines: actual quantiles.

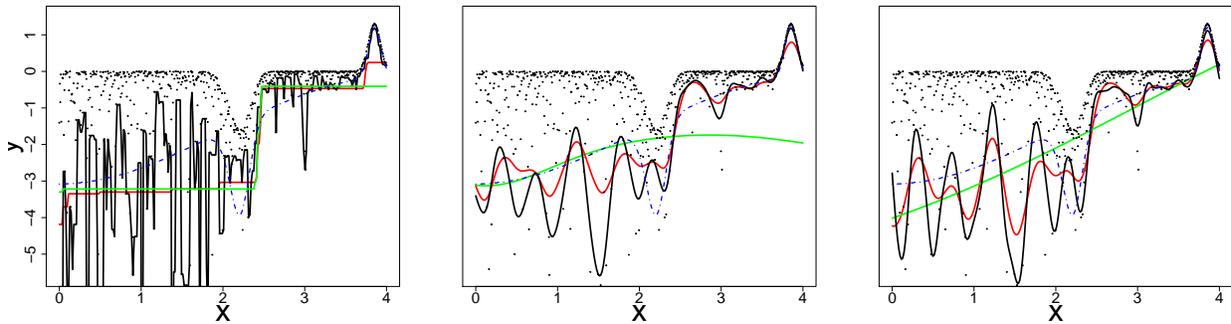


Figure 19: Metamodel responses for toy problem 3 and $\tau = 0.1$ with 640 training sample (left: RF, center: RK, right: VB) for different values of hyperparameters. The true 0.1-quantile is presented in dotted blue lines. In green and black two extreme metamodels associated to two extreme hyperparameter values, while in red the oracle metamodels are represented.

The neighborhood approaches first estimate the CDF, then extract the quantiles. If the hyperparameters are the same for all quantiles, crossing is impossible. However in our setting, different neighborhood sizes were used for different quantiles.

With functional analysis approaches, crossing may happen even if each quantile is built with the same hyperparameters. In the literature, authors have produced methods to address this issue. It could be reduced by the introduction of additional constraints in the model [90] or by the construction of a new model that intrinsically produces non-crossing curves [75]. However in both cases, the dimension of the optimization problem then increases significantly.

Finally the stochastic process approaches estimate each quantile in independent Gaussian processes, so crossing may occur. While the number of training points is larger than what we consider here (10^4 repetitions for each input point), [17] use GP and takes into account all the quantile orders at the same time and thus ensures non-crossing.

Another approach available for all the methods presented here is related to the rearrangement of curves or isotonic regression [6, 2]. The idea is to perform many quantile regressions with a large number of different values of τ or a large set of bootstrapped versions of \mathcal{D}_n and then to rearrange the curves, thereby obtaining the whole distribution and then extracting the quantiles that by definition do not cross.

In theory, adding non-crossing constraints and predicting several quantiles simultaneously could improve the quality of the estimates (in particular as it might add some robustness). However, in practice, it also makes the model more rigid (i.e. a single regularization hyperparameter for all quantiles), and preliminary experiments have shown no gain in accuracy compared to independent predictions, despite a considerably higher computational cost. Hence, multi-quantile predictors were not considered in our study.

9.5 Assessment of prediction accuracy

To assess the accuracy of the results and the confidence that we can have in the estimation it could be useful to provide confidence intervals for the predictor. From this point of view the methods are not equal. With Bayesian approaches, theoretical confidence intervals are provided with the models. More precisely as the output model is Gaussian and as it returns the mean and the variance, confidence intervals can be created. For example the 0.9-confidence interval is provided by

$$\text{CI}(x) = \hat{q}_\tau(x) \pm 1.96\sqrt{\mathbb{V}_q(x)}.$$

However as presented on Figure 20, while the confidence intervals obtained from QK seem useful, the VB model is clearly overconfident.

The statistical order methods consider that inside each neighborhood the samples are i.i.d. Based on that, it is possible to use Wilks' formula (see [70] for instance) or deviation inequality as presented in [92] to extract confidence intervals. For example, using Chernoff's inequality, for any $\eta > 0$, the confidence interval of order $1 - \eta$ is as $\text{CI}(x) = [L_K(x), U_K(x)]$ with

$$U_K(x) = \min \{q, \hat{F}^K(q|X = x) \geq \tau \text{ and } n \text{kl}(\hat{F}^K(q|X = x), \tau) \geq \log(2/\eta)\},$$

and

$$L_K(x) = \max \{q, \hat{F}^K(q|X = x) \leq \tau \text{ and } n \text{kl}(\hat{F}^K(q|X = x), \tau) \geq \log(2/\eta)\}.$$

Using this method enables the confidence intervals to be data dependent. For example in Figure 20 the confidence intervals are not symmetric. Note that while the confidence intervals obtained with this technique come with theoretical guarantees, they are very conservative and they depend a lot on the size of the training set.

Finally for all methods it is possible to use a bootstrap technique to create different regression models and then to aggregate them in order to create empirical confidence intervals. For example Figure 20 shows confidence intervals using a bootstrap technique for RK. The main drawback of such method is its computational cost.

To the best of our knowledge there are no other methods that sensibly improve this results. Thus there is still a room for improvement concerning quantile regression model assessment.

10 Summary and perspectives

10.1 General recommendations

In this presentation we have introduced six metamodels for quantile regression. In the first part of the paper we have provided a full description of the six metamodels, first focusing on their theoretical basis, then discussing their implementation procedure. This part of the paper have enabled us to highlight the similarities and differences of the methods so that providing critical perspectives on the state of the art. The second part of the paper focused on performance comparison according to the dimension of the problem, the size of the learning set, the signal-to-noise ratio and the value of the pdf at the targeted quantile. We have compared the presented methods on 4 toy problems in dimension 1, 2, 9 and on an agronomic model in dimension 9.

Figure 21 summarizes our findings. In a nutshell, when the signal-to-noise ratio is high RK, VB, QK and NN shows good results in our experimental setting but as soon as the signal-to-noise ratio decreases, quantile regression requires larger budgets and comparing the methods seems to be more complicated. Indeed, while the rule-of-thumb for computer experiments is a budget (i.e. number of experiments) 10 times the dimension (see [54] for instance), as we work on problems with low signal-to-noise ratio, we found that no method was able to provide a relevant quantile estimate with a number of observations less than 50 times the dimension. For larger budgets, no method works uniformly better than any other. NN and VB are best when the budget is large. When the budget is smaller, RF, RK, KN are best when the pdf is small in the neighborhood of the quantile, in other words, when little information is available. However, VB outperforms all the other methods when more information is available, that is, when the pdf is large in the neighborhood of the quantile.

10.2 Possible ways of improvement

In our benchmark, we generally followed the approaches as presented by their authors. However, most of them could be improved. The optimization scheme of NN is the computational bottleneck of the method, which makes it the most expensive method in our benchmark system. One possible improvement would be using the BFGS algorithm (see [51] for details about the BFGS algorithm applied to non-smooth functions) or the ADAM algorithm [46] to optimize

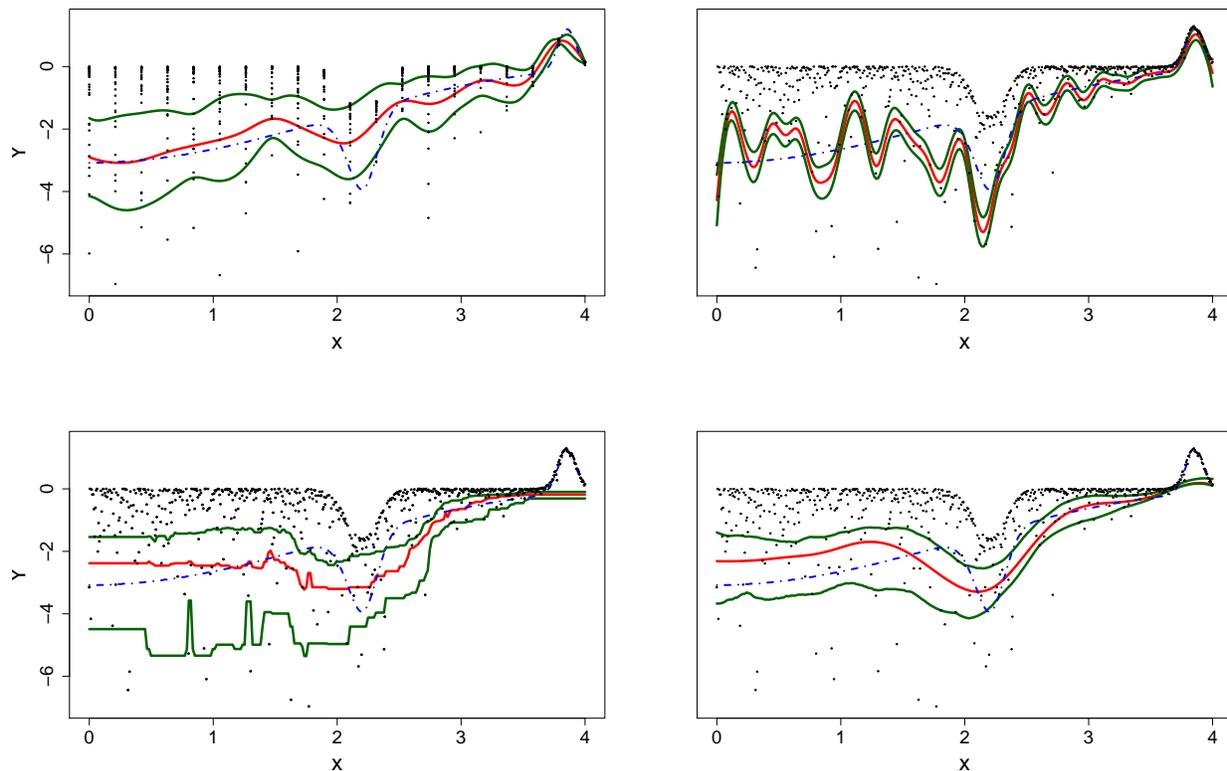


Figure 20: In red quantile metamodel, in blue the true 0.1-quantile and in green the 0.9-confidence intervals. Top left QK, top right VB, bottom left KN, bottom right RK.

directly the empirical risk associated to the pinball loss. A faster scheme would allow more restarts, and hence improve robustness.

Another improvement concerns the splitting criterion (7) of RF, which is not designed for the quantile but for the expectation. This could lead to poor estimates for problems where quantiles are weakly correlated with expectations. Defining an appropriate splitting criterion could significantly improve the performance of this method.

In our experiments, QK used a predefined number of sampling points that were heuristically defined as a trade-off between space-filling and pointwise quantile estimation accuracy. The performance of QK could be significantly improved by optimally tuning the ratio between the number of points and repetitions, in the spirit of [8].

The KN method can naturally be extended to a variant that uses the whole sample instead of the K nearest points. The weights associated to each point of the sample could be based on Gaussian or triangular kernel for example. This idea has been developed in [100] to estimate the conditional expectation but we think that it could be possible to adapt this approach to the estimation of the conditional quantile.

Finally, in practice, finding the best hyperparameters was the most difficult part of the proposed benchmark system. While this aspect is often toned down by authors, we believe hyperparameters tuning is a key practical aspect that remains a challenging problem in quantile regression.

Acknowledgments

This work is part of a Ph.D. of L. Torossian funded by INRA and Région Occitanie. The authors would like to thank Edouard Pauwels for the discussions on the topic of convex and non-convex optimization and Pierre Casadebaig for its suggestions and advices concerning the crop model sunflo.

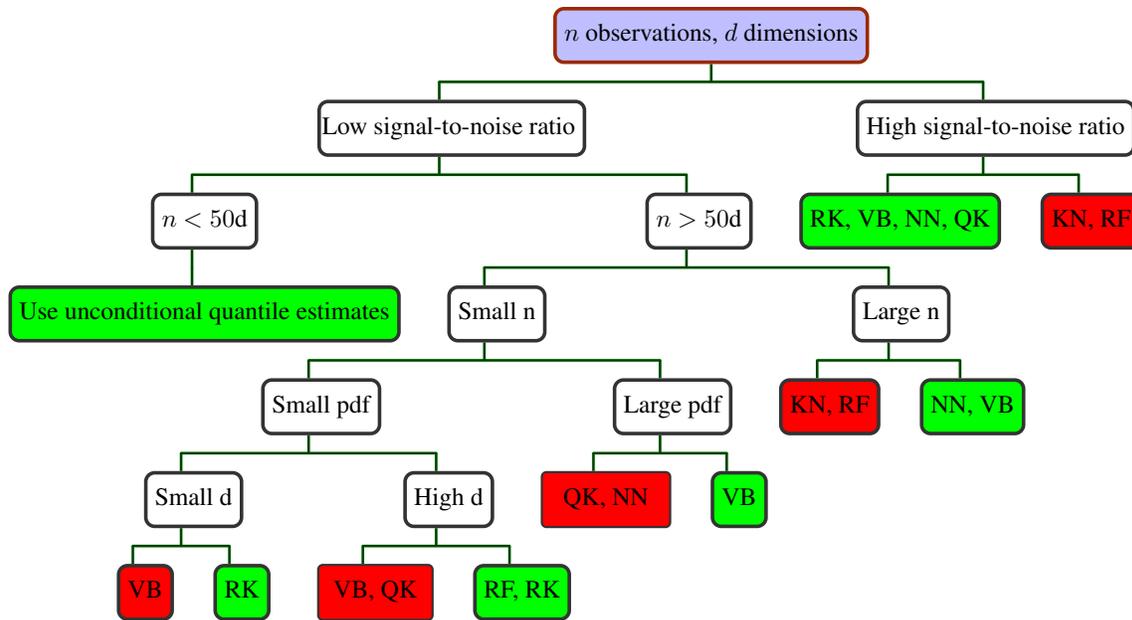


Figure 21: Method recommendation depending on the problem at hand (green: recommended methods, red: methods to avoid). KN: nearest-neighbors, RF: random forests, NN: neural networks, RK: RKHS regression, QK: quantile kriging, VB: variational Bayesian, unconditional quantile estimates: KN with $K = n$.

References

References

- [1] Sachinthaka Abeywardana and Fabio Ramos. Variational inference for nonparametric bayesian quantile regression. In *AAAI*, pages 1686–1692, 2015.
- [2] Jason Abrevaya. Isotonic quantile regression: asymptotics and bootstrap. *Sankhyā: The Indian Journal of Statistics*, pages 187–199, 2005.
- [3] DH Ackley. A connectionist machine for genetic hillclimbing, vol ume secs28 of, 1987.
- [4] Sunil Arya, David M Mount, Nathan S Netanyahu, Ruth Silverman, and Angela Y Wu. An optimal algorithm for approximate nearest neighbor searching fixed dimensions. *Journal of the ACM (JACM)*, 45(6):891–923, 1998.
- [5] François Bachoc. *Parametric estimation of covariance function in Gaussian-process based Kriging models. Application to uncertainty quantification for computer experiments*. PhD thesis, Université Paris-Diderot-Paris VII, 2013.
- [6] Alexandre Belloni, Victor Chernozhukov, Denis Chetverikov, and Iván Fernández-Val. Conditional quantile processes based on series or many regressors. *arXiv preprint arXiv:1105.6154*, 2011.
- [7] Pallab K Bhattacharya and Ashis K Gangopadhyay. Kernel and nearest-neighbor estimation of a conditional quantile. *The Annals of Statistics*, pages 1400–1415, 1990.
- [8] Mickael Binois, Robert B Gramacy, and Mike Ludkovski. Practical heteroskedastic gaussian process modeling for large simulation experiments. *Journal of Computational and Graphical Statistics*, (just-accepted):1–41, 2018.
- [9] Christopher M Bishop. *Neural networks for pattern recognition*. Oxford university press, 1995.
- [10] Alexis Boukouvalas, Remi Barillec, and Dan Cornford. Gaussian process quantile regression using expectation propagation. *arXiv preprint arXiv:1206.6391*, 2012.
- [11] George EP Box and George C Tiao. *Bayesian inference in statistical analysis*, volume 40. John Wiley & Sons, 2011.
- [12] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

- [13] Mikio L Braun. Accurate error bounds for the eigenvalues of the kernel matrix. *Journal of Machine Learning Research*, 7(Nov):2303–2328, 2006.
- [14] Leo Breiman. Bias, variance, and arcing classifiers. 1996.
- [15] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [16] Leo Breiman. *Classification and regression trees*. Routledge, 2017.
- [17] Thomas Browne, Bertrand Iooss, Loïc Le Gratiet, Jérôme Lonchamp, and Emmanuel Remy. Stochastic simulators based optimization by gaussian process metamodels—application to maintenance investments planning issues. *Quality and Reliability Engineering International*, 32(6):2067–2080, 2016.
- [18] Alex J Cannon. Quantile regression neural networks: Implementation in r and application to precipitation downscaling. *Computers & geosciences*, 37(9):1277–1284, 2011.
- [19] Pierre Casadebaig, Lydie Guilioni, Jérémie Lecoecur, Angélique Christophe, Luc Champolivier, and Philippe Debaeke. Sunflo, a model to simulate genotype-specific performance of the sunflower crop in contrasting environments. *Agricultural and forest meteorology*, 151(2):163–178, 2011.
- [20] Marco Cavazzuti. Design of experiments. In *Optimization Methods*, pages 13–42. Springer, 2013.
- [21] Chih-Chung Chang and Chih-Jen Lin. Libsvm: a library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)*, 2(3):27, 2011.
- [22] Ching-Hua Chang and Chung-Wei Ha. On eigenvalues of differentiable positive definite kernels. *Integral Equations and Operator Theory*, 33(1):1–7, 1999.
- [23] Andreas Christmann and Ingo Steinwart. Consistency of kernel-based quantile regression. *Applied Stochastic Models in Business and Industry*, 24(2):171–183, 2008.
- [24] Andreas Christmann and Ingo Steinwart. How svms can estimate quantiles and the median. In *Advances in neural information processing systems*, pages 305–312, 2008.
- [25] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.
- [26] Noel Cressie. The origins of kriging. *Mathematical geology*, 22(3):239–252, 1990.
- [27] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society. Series B (methodological)*, pages 1–38, 1977.
- [28] Ramón Díaz-Uriarte and Sara Alvarez De Andres. Gene selection and classification of microarray data using random forest. *BMC bioinformatics*, 7(1):3, 2006.
- [29] L C W Dixon and G P Szego. The global optimization problem. an introduction. *Toward global optimization*, 2:1–15, 1978.
- [30] Harris Drucker, Christopher JC Burges, Linda Kaufman, Alex J Smola, and Vladimir Vapnik. Support vector regression machines. In *Advances in neural information processing systems*, pages 155–161, 1997.
- [31] Sahibsingh A Dudani. The distance-weighted k-nearest-neighbor rule. *IEEE Transactions on Systems, Man, and Cybernetics*, (4):325–327, 1976.
- [32] Sam Efromovich. Dimension reduction and adaptation in conditional density estimation. *Journal of the American Statistical Association*, 105(490):761–774, 2010.
- [33] Kai-Tai Fang, Runze Li, and Agus Sudjianto. *Design and modeling for computer experiments*. Chapman and Hall/CRC, 2005.
- [34] Dani Gamerman and Hedibert F Lopes. *Markov chain Monte Carlo: stochastic simulation for Bayesian inference*. Chapman and Hall/CRC, 2006.
- [35] Vincent Garcia, Eric Debreuve, and Michel Barlaud. Fast k nearest neighbor search using gpu. In *Computer Vision and Pattern Recognition Workshops, 2008. CVPRW'08. IEEE Computer Society Conference on*, pages 1–6. IEEE, 2008.
- [36] EV Gijo and Johny Scaria. Product design by application of taguchi’s robust engineering using computer simulation. *International Journal of Computer Integrated Manufacturing*, 25(9):761–773, 2012.
- [37] Walter R Gilks, Sylvia Richardson, and David Spiegelhalter. *Markov chain Monte Carlo in practice*. CRC press, 1995.
- [38] Peter Hall, Jeff Racine, and Qi Li. Cross-validation and the estimation of conditional probability densities. *Journal of the American Statistical Association*, 99(468):1015–1026, 2004.

- [39] Nikolaus Hansen. Benchmarking the nelder-mead downhill simplex algorithm with many local restarts. In *Proceedings of the 11th Annual Conference Companion on Genetic and Evolutionary Computation Conference: Late Breaking Papers*, pages 2403–2408. ACM, 2009.
- [40] Xuming He. Quantile curves without crossing. *The American Statistician*, 51(2):186–192, 1997.
- [41] Marti A. Hearst, Susan T Dumais, Edgar Osuna, John Platt, and Bernhard Scholkopf. Support vector machines. *IEEE Intelligent Systems and their applications*, 13(4):18–28, 1998.
- [42] Ralf Herwig. Computational modeling of drug response with applications to neuroscience. *Dialogues in clinical neuroscience*, 16(4):465, 2014.
- [43] Hemant Ishwaran. The effect of splitting on random forests. *Machine Learning*, 99(1):75–118, 2015.
- [44] Janis Janusevskis and Rodolphe Le Riche. Simultaneous kriging-based estimation and optimization of mean response. *Journal of Global Optimization*, 55(2):313–336, 2013.
- [45] Kristian Kersting, Christian Plagemann, Patrick Pfaff, and Wolfram Burgard. Most likely heteroscedastic gaussian process regression. In *Proceedings of the 24th international conference on Machine learning*, pages 393–400. ACM, 2007.
- [46] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [47] Roger Koenker and Gilbert Bassett Jr. Regression quantiles. *Econometrica: journal of the Econometric Society*, pages 33–50, 1978.
- [48] Samuel Kotz, Tomasz Kozubowski, and Krzysztof Podgorski. *The Laplace distribution and generalizations: a revisit with applications to communications, economics, engineering, and finance*. Springer Science & Business Media, 2012.
- [49] Miguel Lázaro-Gredilla and Michalis Titsias. Variational heteroscedastic gaussian process regression. 2011.
- [50] Jinzhi Lei. Stochastic modeling in systems biology. *J. Adv. Math. Appl*, 1(1):76–88, 2012.
- [51] Adrian S Lewis and Michael L Overton. Nonsmooth optimization via bfgs. *Submitted to SIAM J. Optimiz*, pages 1–35, 2009.
- [52] Youjuan Li, Yufeng Liu, and Ji Zhu. Quantile regression in reproducing kernel hilbert spaces. *Journal of the American Statistical Association*, 102(477):255–268, 2007.
- [53] Nikolas List and Hans Ulrich Simon. Svm-optimization and steepest-descent line search. In *Proceedings of the 22nd Annual Conference on Computational Learning Theory*, 2009.
- [54] Jason L Loepky, Jerome Sacks, and William J Welch. Choosing the sample size of a computer experiment: A practical guide. *Technometrics*, 51(4):366–376, 2009.
- [55] Gilles Louppe. Understanding random forests: From theory to practice. *arXiv preprint arXiv:1407.7502*, 2014.
- [56] Michael Ludkovski and Jarad Niemi. Optimal dynamic policies for influenza management. *Statistical Communications in Infectious Diseases*, 2(1), 2010.
- [57] Kristian Lum, Alan E Gelfand, et al. Spatial quantile multiple regression using the asymmetric laplace process. *Bayesian Analysis*, 7(2):235–258, 2012.
- [58] Lawrence C Marsh and David R Cormier. *Spline regression models*, volume 137. Sage, 2001.
- [59] Geoffrey McLachlan and Thiriyambakam Krishnan. *The EM algorithm and extensions*, volume 382. John Wiley & Sons, 2007.
- [60] Nicolai Meinshausen. Quantile regression forests. *Journal of Machine Learning Research*, 7(Jun):983–999, 2006.
- [61] Nicolai Meinshausen and Maintainer Nicolai Meinshausen. The quantregforest package. 2007.
- [62] Thomas P Minka. Expectation propagation for approximate bayesian inference. In *Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence*, pages 362–369. Morgan Kaufmann Publishers Inc., 2001.
- [63] Vincent Moutoussamy, Simon Nanty, and Benoît Pauwels. Emulators for stochastic simulation codes. *ESAIM: Proceedings and Surveys*, 48:116–155, 2015.
- [64] Rémi Munos. Optimistic optimization of a deterministic function without the knowledge of its smoothness. In *Advances in neural information processing systems*, pages 783–791, 2011.

- [65] Victor Picheny, Ronan Trépos, and Pierre Casadebaig. Optimization of black-box models with uncertain climatic inputs-application to sunflower ideotype design. *PLoS one*, 12(5):e0176815, 2017.
- [66] Matthew Plumlee and Rui Tuo. Building accurate emulators for stochastic simulations via quantile kriging. *Technometrics*, 56(4):466–473, 2014.
- [67] Adrian E Raftery. Bayesian model selection in social research. *Sociological methodology*, pages 111–163, 1995.
- [68] Prajit Ramachandran, Barret Zoph, and Quoc V Le. Searching for activation functions. *arXiv preprint arXiv:1710.05941*, 2017.
- [69] Carl Edward Rasmussen and Christopher KI Williams. Gaussian processes for machine learning. 2006. *The MIT Press, Cambridge, MA, USA*, 38:715–719, 2006.
- [70] Rolf D Reiss and Ludger Ruschendorf. On wilks’ distribution-free confidence intervals for quantile intervals. *Journal of the American Statistical Association*, 71(356):940–944, 1976.
- [71] Christian Robert and George Casella. *Monte Carlo statistical methods*. Springer Science & Business Media, 2013.
- [72] Roman Rosipal and Leonard J Trejo. Kernel partial least squares regression in reproducing kernel hilbert space. *Journal of machine learning research*, 2(Dec):97–123, 2001.
- [73] Marzena Rostek. Quantile maximization in decision theory. *The Review of Economic Studies*, 77(1):339–371, 2010.
- [74] Olivier Roustant, David Ginsbourger, and Yves Deville. Dicekriging, diceoptim: Two r packages for the analysis of computer experiments by kriging-based metamodeling and optimization. 2012.
- [75] Maxime Sangnier, Olivier Fercoq, and Florence d’Alché Buc. Joint quantile regression in vector-valued rkhs. In *Advances in Neural Information Processing Systems*, pages 3693–3701, 2016.
- [76] Johannes Schmidt-Hieber. Nonparametric regression using deep neural networks with relu activation function. *arXiv preprint arXiv:1708.06633*, 2017.
- [77] Bernhard Schölkopf. The kernel trick for distances. In *Advances in neural information processing systems*, pages 301–307, 2001.
- [78] George AF Seber and Alan J Lee. *Linear regression analysis*, volume 329. John Wiley & Sons, 2012.
- [79] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.
- [80] Jooyong Shim, Changha Hwang, and Kyung Ha Seok. Non-crossing quantile regression via doubly penalized kernel machine. *Computational Statistics*, 24(1):83–94, 2009.
- [81] Joseph Sill, Gábor Takács, Lester Mackey, and David Lin. Feature-weighted linear stacking. *arXiv preprint arXiv:0911.0460*, 2009.
- [82] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [83] Michael L Stein. *Interpolation of spatial data: some theory for kriging*. Springer Science & Business Media, 2012.
- [84] Ingo Steinwart and Andreas Christmann. *Support vector machines*. Springer Science & Business Media, 2008.
- [85] Ingo Steinwart, Andreas Christmann, et al. Estimating conditional quantiles with the help of the pinball loss. *Bernoulli*, 17(1):211–225, 2011.
- [86] Charles J Stone. Nearest neighbor estimators of a nonlinear regression function. In *Computer Science and Statistics: 8th Annual Symposium on the Interface*, pages 413–418, 1975.
- [87] Charles J Stone. Consistent nonparametric regression. *The annals of statistics*, pages 595–620, 1977.
- [88] C.B. Storlie and J.C. Helton. Multiple predictor smoothing methods for sensitivity analysis: Description of techniques. *Reliability Engineering and System Safety*, 93:28–54, 2008.
- [89] Matthew A Taddy and Athanasios Kottas. A bayesian nonparametric approach to inference for quantile regression. *Journal of Business & Economic Statistics*, 28(3):357–369, 2010.
- [90] Ichiro Takeuchi, Quoc V Le, Timothy D Sears, and Alexander J Smola. Nonparametric quantile estimation. *Journal of Machine Learning Research*, 7(Jul):1231–1264, 2006.

- [91] Ole Tange. Gnu parallel 2018. 2018.
- [92] Léonard Torossian, Aurélien Garivier, and Victor Picheny. \mathcal{X} -Armed Bandits: Optimizing Quantiles, CVaR and Other Risks. In *Asian Conference on Machine Learning*, pages 252–267, 2019.
- [93] Berwin A Turlach and Andreas Weingessel. quadprog: Functions to solve quadratic programming problems. *CRAN-Package quadprog*, 2007.
- [94] Dimitris G Tzikas, Aristidis C Likas, and Nikolaos P Galatsanos. The variational approximation for bayesian inference. *IEEE Signal Processing Magazine*, 25(6):131–146, 2008.
- [95] A Van Maanen and X-M Xu. Modelling plant disease epidemics. *European Journal of Plant Pathology*, 109(7):669–682, 2003.
- [96] Vladimir Vapnik. *The nature of statistical learning theory*. Springer science & business media, 2013.
- [97] Ghislain Verdier and Ariane Ferreira. Adaptive mahalanobis distance and k -nearest neighbor rule for fault detection in semiconductor manufacturing. *IEEE Transactions on Semiconductor Manufacturing*, 24(1):59–68, 2011.
- [98] N. Villa-Vialaneix, M. Follador, M. Ratto, and A. Leip. A comparison of eight metamodeling techniques for the simulation of N_2O fluxes and N leaching from corn crops. *Environmental Modelling & Software*, 34:51–66, 2012.
- [99] Ian H Witten, Eibe Frank, Mark A Hall, and Christopher J Pal. *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann, 2016.
- [100] Kai Yu, Liang Ji, and Xuegong Zhang. Kernel nearest-neighbor algorithm. *Neural Processing Letters*, 15(2):147–156, 2002.
- [101] Keming Yu and Rana A Moyeed. Bayesian quantile regression. *Statistics & Probability Letters*, 54(4):437–447, 2001.
- [102] Keming Yu and Jin Zhang. A three-parameter asymmetric laplace distribution and its extension. *Communications in Statistics-Theory and Methods*, 34(9-10):1867–1879, 2005.
- [103] Cheng Zhang, Judith Butepage, Hedvig Kjellstrom, and Stephan Mandt. Advances in variational inference. *IEEE transactions on pattern analysis and machine intelligence*, 2018.
- [104] Peng Zhao and Bin Yu. On model selection consistency of lasso. *Journal of Machine learning research*, 7(Nov):2541–2563, 2006.
- [105] Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.