

Ensemble Multi-Quantiles: Adaptively Flexible Distribution Prediction for Uncertainty Quantification

Xing Yan , Yonghua Su , and Wenxuan Ma 

Abstract—We propose a novel, succinct, and effective approach for distribution prediction to quantify uncertainty in machine learning. It incorporates adaptively flexible distribution prediction of $\mathbb{P}(y|X = x)$ in regression tasks. This conditional distribution’s quantiles of probability levels spreading the interval (0,1) are boosted by additive models which are designed by us with intuitions and interpretability. We seek an adaptive balance between the structural integrity and the flexibility for $\mathbb{P}(y|X = x)$, while Gaussian assumption results in a lack of flexibility for real data and highly flexible approaches (e.g., estimating the quantiles separately without a distribution structure) inevitably have drawbacks and may not lead to good generalization. This ensemble multi-quantiles approach called EMQ proposed by us is totally data-driven, and can gradually depart from Gaussian and discover the optimal conditional distribution in the boosting. On extensive regression tasks from UCI datasets, we show that EMQ achieves state-of-the-art performance comparing to many recent uncertainty quantification methods. Visualization results further illustrate the necessity and the merits of such an ensemble model.

Index Terms—Uncertainty quantification, distribution prediction, adaptive flexibility, non-Gaussianity, ensemble multi-quantiles.

I. INTRODUCTION

MACHINE learning algorithms have become very powerful in recent years due to the successes of deep neural networks. Advanced deep learning models have achieved state-of-the-art performance in many tasks under the evaluation criteria of prediction accuracy [1]. However, people criticize deep learning for not only its opaque in making predictions but also its over-confidence on the predictions [2], [3]. Over-confidence may occur when deep learning has a wrong estimation of the uncertainty. We thereby need giving a good measurement of the uncertainty level for the prediction for each input of the model, e.g., giving a standard deviation or a prediction interval (that the label will fall into with a fixed probability) accompanied with the prediction value in a regression task.

Manuscript received 28 October 2022; revised 11 June 2023; accepted 15 June 2023. Date of publication 20 June 2023; date of current version 3 October 2023. This work was supported in part by the Fundamental Research Funds for the Central Universities, and the Research Funds of Renmin University of China under Grant 21XNLG09, in part by the Young Scientists Fund of National Natural Science Foundation of China under Grant 62206300, and in part by Public Computing Cloud of Renmin University of China. Recommended for acceptance by M. Sugiyama. (Corresponding author: Wenxuan Ma.)

The authors are with the Institute of Statistics and Big Data, Renmin University of China, Beijing 100872, China (e-mail: xingyan@ruc.edu.cn; 2019000154@ruc.edu.cn; mawenxuan@ruc.edu.cn).

Digital Object Identifier 10.1109/TPAMI.2023.3288028

Uncertainty quantification lies in a crucial position in machine learning, especially in real applications of deep learning when people must make decisions relying on it. A wrong estimate of uncertainty will make decisions highly risky. Good uncertainty estimation quantifies when we can trust the model’s predictions. For example, in disease detection [4], the model can defer to human experts when it believes the uncertainty is high. An autopilot can transfer the authority to the human operator when the prediction is risky. Besides, good uncertainty quantification can benefit us in detecting and dealing with out-of-distribution samples [5], whose uncertainty should be quantified to be very high. Another representative case is in financial risk management, decision-makers need to set a new margin rate or to take a hedging action if he/she believes the market uncertainty is high, or, the risk measure defined by some distributional characteristic of the market variable (e.g., the change of S&P 500 index in a future period) is high. Such distributional characteristics include volatility, kurtosis, Value-at-Risk, expected shortfall, etc., all of which are important and commonly-used risk measures in financial markets [6], [7], [8].

It has been documented in many existing works [3], [5], [9], [10] that deep learning may produce poor uncertainty estimation. This is partially because there is a lack of ground truth for the uncertainty. Thus, solving this important problem is facing enormous challenges. For classification tasks where the softmax technique is used, it is straightforward to take the prediction score of the class as uncertainty estimation. But researchers rejected this practice, e.g., in [9], [11]. Although the uncertainty quantification in classification tasks needs much attention, in this paper, we focus on the regression problem whose uncertainty estimation is challenging too. In regression, (aleatoric) uncertainty quantification can be formulated as the problem of estimating some distributional characteristics of the conditional distribution $\mathbb{P}(y|X = x)$, where \mathbf{X} is the feature vector and y is the continuous label.

Recently, a variety of methods have been developed for uncertainty quantification for neural networks. A direct approach is to assume a Gaussian distribution for $\mathbb{P}(y|X = x)$ and try to estimate its standard deviation as a function of x , using a neural network trained with maximum likelihood and some special scheme, see example works such as [5], [12], [13], [14]. This simple idea has been proven very successful when good training and the ensemble are deployed. Besides, Bayesian methods can naturally return the density function of $\mathbb{P}(y|X = x)$, such as in [9], [10], [15]. Bayesian methods are sometimes hard to implement and computationally expensive [5], and they require

the prior distribution to be correctly specified (but usually, priors of convenience are assumed). Conformal prediction [16], [17] is also a popular framework which separates a calibration set from the training set and quantifies uncertainty with a conformity score. Many variants [18], [19], [20] were proposed to make conformal prediction more flexible and adaptive to heteroscedasticity. Some recent approaches estimate the conditional distribution by designing some novel loss function between the assumed $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$ and the observed data, such as in [21], [22]. The ensemble over multiple Gaussians [23] or prediction intervals [24] was proposed to obtain better uncertainty estimation, with the parallel ensemble manner (not the boosting-like in this paper). To summarize, there are many different types of methods for uncertainty quantification in machine learning.

A. Motivations

There is still another class of methods that quantify uncertainty by estimating quantiles [20], [25], [26], [27], [28] to construct prediction intervals, which is also the main focus of this paper. Most existing works employ the quantile regression loss from [29], [30] to do the model training. They generally will produce a prediction interval $[q_\tau(x), q_{1-\tau}(x)]$ formed by the τ -quantile and $(1 - \tau)$ -quantile of $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$. However, we focus on the situation where many $[q_\tau(x), q_{1-\tau}(x)]$ need to be predicted for many different τ (for quantifying uncertainty with different confidences). From another perspective, those quantiles can form an approximation of $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$. The focus now becomes the conditional distribution prediction.

In this case, if no distribution structure is assumed, it is equivalent to assuming an extremely flexible distribution for $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$. It corresponds to the non-parametric approach in statistics context. The opposite extreme is assuming a commonly-used and well-understood distribution structure such as Gaussian and then predicting its parameters. We argue that both the two extremes are potentially harmful for better uncertainty estimation.

In the former approach, the resulting implied density function (can be obtained from the quantiles predicted) will display a chaotic function graph with no distribution shape, even if we let the training dataset be large. One can refer to Fig. 4 for a clear illustration. Some methods developed a neural-network representation for $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$ (e.g., modeling $q_\tau(x)$ as a neural network taking x and $\tau \in (0, 1)$ as inputs), which also incorporate great flexibility in distribution prediction. We argue that too much flexibility does not necessarily lead to performance improvement because: i) the training will become more difficult; and ii) it may be easier to overfit; and iii) it will lead to strange-shaped or even invalid density function with no interpretability. Unnecessary flexibility may be harmful, and has been added into the modeling of $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$ in some existing methods. We will show that seeking a balance between the flexibility and the distribution structure will produce substantially better distribution prediction. Actually, by proposing a novel ensemble multi-quantiles (EMQ) method, we successfully obtain adaptively flexible distribution prediction. It can find the balance adaptively and thus approximate $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$ well, as we will show in the rest of the paper.

Oppositely, in the other approach, assuming a succinct distribution structure for $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$ has obvious shortcomings. It will suffer from the lack of flexibility, and the model performance will heavily rely on the correctness of the choice of distribution structure. For convenience, usually Gaussian is chosen. However, conditional distributions in real data may exhibit multi-modality, asymmetry, and heavy tails, and may even be heterogeneous within a dataset. The mixture of Gaussians in Bayesian approaches or model ensembles does not alleviate this kind of deficiency substantially. Other distribution structures such as t -distribution, skew- t , and Generalized beta, may have complicated mathematical expressions for their density functions, increasing the difficulty of training a model with them with maximum likelihood loss because the gradient with respect to the distribution parameters will be difficult to obtain. Additionally, the choice among them relies on specific domain knowledge.

Another motivation of this paper is to solve the quantile-crossing issue when estimating multi-quantiles of $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$. Although there have been many methods proposed in statistics literature [31], [32], [33], [34], [35], they scarcely solve the unnecessary flexibility issue we mention above, yet may cause other problems such as the computational burden/difficulty if they are constraint optimization-based. Besides constraint optimization, post-processing or re-sorting is another imperfect solution to the quantile-crossing issue. We believe that more natural solutions without additional effort to deal with such an issue are absolutely necessary. Actually, the proposed EMQ naturally overcomes this issue with nice designs. Our solution is a built-in solution without additional effort such as constraint optimization or post-processing, as shown in the rest of the paper.

B. Our Contributions

In this paper, we tackle the problem of conditional distribution prediction in regression tasks through proposing an ensemble learning approach named EMQ. Multi-quantiles are boosted by additive models with intuitive designs, and can be used for uncertainty quantification with excellent performance. To summarize, our contributions are as follows:

- i) We propose a novel, succinct, and effective approach for multi-quantiles ensemble learning, in which a strong learner aims at learning a Gaussian distribution for $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$ and the consequent weak learners are used for tuning or improving the quantiles. The design of the whole ensemble approach is intuitive and interpretable, and aims at adaptively seeking a balance between the distribution structure (such as Gaussian) and the flexibility.
- ii) The quantile-crossing issue is overcome naturally in our approach without additional effort.
- iii) In the experiments, the proposed EMQ achieves start-of-the-art performance over a large range of datasets under the criteria of calibration and sharpness, compared to many recent uncertainty quantification methods including Gaussian assumption-based, Bayesian methods, quantile regression-based, and traditional tree models.
- iv) Through the experiments, we find that EMQ can indeed do adaptively flexible distribution prediction, especially when

we adopt the adaptive T_{ada} strategy. The ensemble steps in EMQ work correctly as we except. In addition, the variant EMQW adopts the weighted quantile loss and achieves superior performance on the tail sides.

v) Empirically, we find that some datasets have Gaussian-distributed $\mathbb{P}(y|\mathbf{X} = x)$ while others have non-Gaussian ones with significantly different density shapes such as sharp peak, asymmetry, long tail, multi-modality, etc., all of which are successfully captured by our approach. Even in one dataset, this variety may occur, which verifies the complexity of real data and the necessity and merits of our approach.

Code and data are released in the link.¹

II. RELATED WORKS

We summarize some recent popular works that are directly related to our research, which will be the competing methods in our experiments. They can be divided into three categories: variance networks, Bayesian, and quantile regression-based.

Variance Networks: It is straightforward to let the neural network output a variance to represent the uncertainty. Heteroskedasticity Neural Network (HNN) is such a method that outputs both the mean (the prediction) and the variance. Gaussian maximum likelihood loss is used to learn the network parameters. [13] suggested to separately and alternately train a mean network and a variance network, for better variance estimation. We will follow this suggestion in our implementation of HNN in the experiments. Deep Ensemble is the method who will train multiple HNNs with different initializations and aggregate the means and variances obtained (see [5]). Despite being popular, these two methods are Gaussian assumption-based and are not suitable for more complicated non-Gaussian cases.

Bayesian: Bayesian methods compute the predictive posterior distribution as the uncertainty estimation by weighting the likelihood with the posterior of network parameters. MC Dropout [9] uses dropout as a Bayesian approximation to overcome the computational challenge. Concrete Dropout [36] is a dropout variant which allows the dropout probability to be tuned continuously using gradient methods. As stated in original papers, these two Bayesian methods are popular for large computer vision models. Nevertheless, they do not show superior performance on tabular data. In addition, similar to variance networks, they only output a variance as the uncertainty estimation, without more details on the conditional distribution.

Quantile Regression-Based: Quantile regression has a long history in statistics literature [29], [30], [37]. The loss function in quantile regression is $L_\tau(y, q) = (y - q)(\tau - \mathbb{I}_{\{y < q\}})$, where y is the observed label, q is the τ -quantile predicted by a model given x , and τ is a specified probability level. If a model can output multiple quantiles of different levels, the losses for different τ can be summed and then minimized. We call this approach vanilla quantile regression (Vanilla QR). [26] proposed simultaneous quantile regression (SQR), a neural network model which takes both the feature vector x and the level τ as inputs, and outputs the τ -quantile, and minimizes the above pinball loss

with $\tau \sim \text{Uniform}(0, 1)$. An alternative to the pinball loss is the interval score proposed in [38]. Let (l, u) be the $(1 - \tau)$ centered prediction interval, where l and u stand for the $\frac{\tau}{2}$ -quantile and $(1 - \frac{\tau}{2})$ -quantile we are looking for. This yields the interval score

$$S_\tau(l, u; y) = (u - l) + \frac{2}{\tau}(l - y)\mathbb{I}_{\{y < l\}} + \frac{2}{\tau}(y - u)\mathbb{I}_{\{y > u\}}. \quad (1)$$

Besides, [27] proposed a combined loss considering both calibration and sharpness terms and first simultaneously optimized the interval score. [28] proposed the pinball loss with an orthogonality regularization by considering the independence between the interval length and the violation.

Other Methods of Probabilistic Forecasting: [39] adopted a non-parametric method, the linear spline, or a piecewise-linear function to represent a conditional quantile function. The parameters of the spline are the output of a recurrent neural network with the features as its input. It may need many pieces in such a piecewise-linear function to approximate a smooth quantile function well, e.g., the quantile function of a simple Gaussian. In financial econometrics literature, [40], [41] proposed a skewed generalized t -distribution, which has a highly complicated mathematical form and of course is computationally unfriendly, especially in the conditional setting where the gradient of the loss will be difficult to obtain. [42] adopted a new parametric quantile function which transforms the Gaussian. It has four parameters that control the location, scale, right tail, and left tail respectively. The new quantile function has an S-shaped Q-Q plot against Gaussian, thus with heavier tails, but with limited flexibility.

A. Aleatoric and Epistemic Uncertainty

Aleatoric uncertainty [43] refers to the randomness of the outcome, or, the inherently stochastic and irreducible random effects. It does not vanish as the sample size goes to infinity. Epistemic uncertainty [44], or model uncertainty, refers to the uncertainty caused by a lack of knowledge. It is the uncertainty due to model limitations. Generally, Bayesian and ensemble methods can estimate both epistemic and aleatoric uncertainty [13]. The variance networks and quantile-based methods focus on the aleatoric uncertainty. In this paper, we focus on the aleatoric uncertainty only too, as what we propose is intrinsically a quantile-based approach. Both types of uncertainty are important and challenging to estimate. For more about epistemic uncertainty estimation, please refer to [26], [45], etc.

III. METHODOLOGY

In this section, we elaborate the proposed ensemble multi-quantiles (EMQ) approach for uncertainty quantification. Throughout the paper, we use \mathbf{X} and \mathbf{y} to denote the random variables of the features and the continuous label. Given the feature vector $\mathbf{X} = x$, the goal is to predict the conditional τ_k -quantile of the variable \mathbf{y} , denoted as $q_k(x)$, where $k = 1, \dots, K$ and $0 < \tau_1 < \dots < \tau_K < 1$. Here τ_1, \dots, τ_K are probability levels specified by us that densely distribute in the interval $(0, 1)$, such as 0.01, 0.02, \dots , 0.98, 0.99, denoted as the set $\Pi =$

¹<https://github.com/xingyan-fml/emq>

$\{\tau_k\}_{k=1}^K$. These K conditional quantiles form an approximation of the conditional distribution $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$. The learning will be conducted given a dataset $\{(x_i, y_i)\}_{i \in D_{tr}}$. For concision, we sometimes drop the letter x and simplify the notation $q_k(x)$ to q_k , and suppose x is given and remains the same in this situation.

A. The Framework

Our general idea is to iteratively update the quantile predictions in multiple steps in a boosting manner, i.e., updating q_1^t, \dots, q_K^t as the step $t = 0, 1, \dots, T$ with additive models and taking the final-step predictions q_1^T, \dots, q_K^T as the desired ones. To be more specific, denoting $Q^t(x) = [q_1^t(x), \dots, q_K^t(x)]^\top$, we propose a new ensemble framework

$$Q^0(x) = G_0(F_0(x; \Theta_0)), \quad (2)$$

$$Q^t(x) = Q^{t-1}(x) + G(F_t(x; \Theta_t), Q^{t-1}(x)), \quad t = 1, \dots, T. \quad (3)$$

F_0 and F_t are neural network models with learnable parameters Θ_0 and Θ_t . G_0 and G are fixed transformations that convert the neural network outputs to the quantile predictions, as we will introduce later.

Iteratively, at every step $t = 0, \dots, T$, a learning objective is set by us to train the model F_t with F_0, \dots, F_{t-1} fixed

$$\min_{\Theta_t} \mathbb{E}_{\mathbb{P}(\mathbf{X}, \mathbf{y})} [L(\mathbf{y}, Q^t(\mathbf{X}))]. \quad (4)$$

L is a proper loss function between y and $Q^t(x)$. Given the probability level set $\Pi = \{\tau_k\}_{k=1}^K$, the loss L is chosen as the sum of quantile regression losses

$$L(y, Q^t(x)) = \sum_{k=1}^K L_{\tau_k}(y, q_k^t(x)), \quad (5)$$

$$L_{\tau}(y, q) = (y - q)(\tau - \mathbb{I}_{\{y < q\}}). \quad (6)$$

Given the training dataset, the empirical version of the learning objective that we should optimize is

$$\min_{\Theta_t} \frac{1}{|D_{tr}|} \sum_{i \in D_{tr}} L(y_i, Q^t(x_i)) + \lambda_t R(\Theta_t). \quad (7)$$

$R(\Theta_t)$ is a regularization term and λ_t is a trade-off hyperparameter. The above framework is general and needs to be elaborated further. In the following, we carefully design every component of it, particularly the choices for F_0, G_0, F_t, G and the dealing of regularization term $\lambda_t R(\Theta_t)$.

B. Intuitions

The Initial Step $t = 0$: In traditional boosting methods, weak learners are boosted iteratively to form a strong learner. The weak learner at $t = 0$ may be a constant (such as in GBDT) or have no difference to those at $t > 0$. Here we face a different situation. Because we are motivated by i) the necessary adaptive flexibility of the distribution that departs from Gaussian, and ii) the fact that assuming Gaussian distribution can yield acceptable uncertainty estimation sometimes. Thus, starting from Gaussian

will be a natural and reasonable idea. Consequently, we aim to let the quantile predictions $Q^0(x)$ at $t = 0$ be as good as possible under the Gaussian assumption. So, F_0 and G_0 will be designed to align with this goal. F_0 cannot be treated as a weak learner now. This will make the following learning at $t > 0$ relatively easy.

The Ensemble Steps $t > 0$: After learning a Gaussian-distributed $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$, imposing consecutive moderate modifications or tuning on the distribution is a natural idea, and is also consistent with the motivations of this paper. So, when $t > 0$, we set F_t as a relatively weak learner that aims to moderately update the quantile predictions $Q^{t-1}(x)$ to better ones $Q^t(x)$. This updating should not be totally free or non-parametric, because preserving partially the old distribution is necessary for the purpose of adaptive flexibility. F_t and G will be designed accordingly. The transformer G which is a crucial part in the framework builds a bridge between the weak learner F_t and the quantile differences $Q^t(x) - Q^{t-1}(x)$ in two adjacent steps, as shown later.

The Regularization Term: Traditional machine learning models may adopt explicit regularization terms added into the learning objective, to decrease the model complexity and reduce the risk of over-fitting. For deep neural networks, the regularization terms can be in implicit forms, such as batch normalization, dropout, and early stopping, all of which are popular strategies. In our framework, both F_0 and F_t are neural networks. Thus, we adopt the following regularization strategies instead of the explicit term $\lambda_t R(\Theta_t)$: i) restrict the number of layers and the hidden layer sizes in F_t when $t > 0$; ii) restrict the output dimension of F_t for $t \geq 0$ (with the assistance of non-learnable G_0 or G to map the output to the quantiles); and iii) adopt those strategies mentioned above such as early stopping. Finally and most importantly, to further avoid over-fitting and to seek the balance between the flexibility and the distribution structure, we use a regularization strategy of adaptively deciding the number of ensemble steps T , called the adaptive T strategy.

The Monotonic Property of Quantiles: As we mentioned in Section I-A, a natural and built-in solution to the quantile-crossing issue is required in our approach. Because $Q^0(x)$ obeys a Gaussian distribution, the monotonic property is of course satisfied. We suppose $Q^{t-1}(x)$ satisfies the monotonic property, i.e., $q_1^{t-1} < \dots < q_K^{t-1}$. As long as we have an appropriate transformer G whose output is the difference $Q^t(x) - Q^{t-1}(x)$, we can yield monotonically increasing $Q^t(x)$: $q_1^t < \dots < q_K^t$. Thereby, we can solve the quantile-crossing issue naturally without additional effort. Again we see G is a crucial part in our framework. We will design an appropriate G that meets this goal.

To summarize, starting from Gaussian at $t = 0$, we gradually and moderately update q_1^t, \dots, q_K^t to be better, and meanwhile ensure that q_1^t, \dots, q_K^t are strictly increasing with respect to k . The relatively strong learner F_0 and the weak learners F_t , $t > 0$ will be trained. Appropriate G_0 and G will be designed and useful regularization strategies will be adopted, including the adaptive T strategy. The resulting final-step predictions q_1^T, \dots, q_K^T are monotonically increasing and can form a good

approximation of the conditional distribution $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$. We will elaborate all of these in the following.

C. The Initial Step

At $t = 0$, Gaussian assumption gives $\mathbb{P}(\mathbf{y}|\mathbf{X} = x) \sim \mathcal{N}(\mu(x), \sigma(x)^2)$. We choose F_0 in (2) as a neural network which maps the feature vector x to the parameters of the Gaussian $\mu(x), \sigma(x)$

$$[\mu(x), \sigma(x)]^\top = F_0(x; \Theta_0). \quad (8)$$

The conditional τ_k -quantile of \mathbf{y} given $\mathbf{X} = x$ is then $q_k^0(x) = \mu(x) + \sigma(x)\Phi^{-1}(\tau_k)$, where $\Phi^{-1}(\tau_k)$ is the τ_k -quantile of standard Gaussian. Thus, G_0 in (2) is

$$\begin{aligned} Q^0(x) &= G_0(\mu(x), \sigma(x)) \\ &= \mu(x) + \sigma(x)[\Phi^{-1}(\tau_1), \dots, \Phi^{-1}(\tau_K)]^\top. \end{aligned}$$

Given the training set $\{(x_i, y_i)\}_{i \in D_{tr}}$ and the intuition of implicit regularizations, the learning objective at $t = 0$ is $\min_{\Theta_0} \frac{1}{|D_{tr}|} \sum_{i \in D_{tr}} L(y_i, Q^0(x_i))$, as in (7). After this, $Q^0(x)$ may be a coarse approximation of the conditional distribution $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$. It is not flexible enough to express the multimodality, asymmetry, heavy tails, etc. Carefully picking a non-Gaussian assumption relying on the expertise is far away from a good solution. So in our approach, we improve the predictions in ensemble steps.

Notice that our framework is flexible to different application domains, because at $t = 0$, it is not limited to Gaussian assumption only. One can have other choices as needed in specific tasks, for example, replacing standard Gaussian by the standard exponential distribution or some other.

D. The Ensemble Steps

Suppose $Q^{t-1}(x)$, or $q_1^{t-1}, \dots, q_K^{t-1}$, obtained at step $t - 1$ is known, now we give the model details in (3) at the ensemble step $t \geq 1$. It is obvious that q_1^0, \dots, q_K^0 are strictly increasing because they obey a Gaussian distribution, so we always assume $q_1^{t-1}, \dots, q_K^{t-1}$ are strictly increasing. Specifically, when $t \geq 1$, to generate strictly increasing and better predictions q_1^t, \dots, q_K^t , we add a new quantity to each q_k^{t-1}

$$q_k^t = q_k^{t-1} + g_k^t(\lambda_k^t), \quad k = 1, \dots, K, \quad (9)$$

where $\lambda_k^t \in (-1, 1)$ is a quantity that we need to determine later. $g_k^t(\cdot)$ is a continuous function we set that is strictly increasing over $[-1, 1]$ and satisfies

$$g_k^t(0) = 0, \quad (10)$$

$$g_k^t(1) = r_k^{t-1} := \frac{q_{k+1}^{t-1} - q_k^{t-1}}{2}, \quad (11)$$

$$g_k^t(-1) = l_k^{t-1} := \frac{q_{k-1}^{t-1} - q_k^{t-1}}{2}. \quad (12)$$

In this paper, we choose a piecewise linear $g_k^t(\cdot)$ to satisfy the above conditions

$$g_k^t(\lambda) = (-l_k^{t-1} + \mathbb{I}_{\{\lambda > 0\}} \cdot (r_k^{t-1} + l_k^{t-1})) \lambda. \quad (13)$$

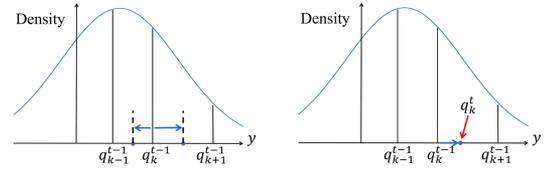


Fig. 1. The illustration of the ensemble step $t \geq 1$. q_k^t is selected in the range between the middle point of $(q_{k-1}^{t-1}, q_k^{t-1})$ and the middle point of $(q_k^{t-1}, q_{k+1}^{t-1})$. By this design, q_1^t, \dots, q_K^t are ensured to be strictly increasing if $q_1^{t-1}, \dots, q_K^{t-1}$ are strictly increasing.

To be more intuitive, one can easily verify that the left and right bounds of q_k^t given $\lambda_k^t \in (-1, 1)$ are the middle points of the intervals $(q_{k-1}^{t-1}, q_k^{t-1})$ and $(q_k^{t-1}, q_{k+1}^{t-1})$ respectively, and besides, $q_k^t = q_k^{t-1}$ when $\lambda_k^t = 0$. This design ensures that q_k^t is strictly increasing with respect to k because obviously,

$$q_k^t < q_k^{t-1} + g_k^t(1) = \frac{q_{k+1}^{t-1} + q_k^{t-1}}{2} = q_{k+1}^{t-1} + g_{k+1}^t(-1) < q_{k+1}^t. \quad (14)$$

The idea here is to move q_k^{t-1} towards a better prediction q_k^t , but with only a small distance to i) remain the monotonic property, and ii) control the flexibility added in step t (see the description later). We illustrate this idea in Fig. 1 as well. When $k = 0$ or $k = K + 1$, we naturally set q_0^{t-1} or q_{K+1}^{t-1} to be the left/right boundary of the support set of $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$ if the boundary is finite. Now $g_1^t(\cdot)$ or $g_K^t(\cdot)$ is well-defined such that $q_1^{t-1} + g_1^t(-1)$ or $q_K^{t-1} + g_K^t(1)$ is a feasible bound. Otherwise, then $q_0^{t-1} = -\infty$ or $q_{K+1}^{t-1} = +\infty$, we can now replace q_0^{t-1} by $-B$ or replace q_{K+1}^{t-1} by $+B$ where B is a large positive constant to make $g_1^t(\cdot)$ or $g_K^t(\cdot)$ well-defined.

To make this step t complete, we set λ_k^t as a cubic polynomial function of τ_k followed by a tanh for transforming to $(-1, 1)$, i.e., $\lambda_k^t = \tanh(\sum_{l=0}^3 a_l^t(\tau_k)^l)$, where a_l^t is an output of the model $F_t(x; \Theta_t)$

$$[a_0^t, a_1^t, a_2^t, a_3^t]^\top = F_t(x; \Theta_t). \quad (15)$$

So far, the function G in (3) has been defined. We avoid letting λ_k^t be totally free and be an output of $F_t(x; \Theta_t)$ directly, which will lead to totally K outputs for F_t (i.e., $K = 99$). The cubic polynomial works as an implicit regularization that puts restrictions on λ_k^t , $k = 1, \dots, K$. This setting controls the flexibility added in step t and is crucially important and useful for adaptively balancing the distribution structure and the flexibility for $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$.

From another perspective, our approach will give a smooth density function for $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$ at each step t . Oppositely, letting λ_k^t be totally free will give a strange or chaotic density function that cannot be interpreted, which is also the common drawback of all over-flexibility methods, as we mentioned in Section I-A. Back to our method, when the training dataset $\{(x_i, y_i)\}_{i \in D_{tr}}$ is given, the learning objective at step t is $\min_{\Theta_t} \frac{1}{|D_{tr}|} \sum_{i \in D_{tr}} L(y_i, Q^t(x_i))$ with implicit regularizations. We summarize the whole procedure of our approach in Algorithm 1.

Algorithm 1: The Ensemble Multi-Quantiles (EMQ) Approach for Uncertainty Quantification.

Input: the training dataset $\{(x_i, y_i)\}_{i \in D_{tr}}$, and K probability levels of interest $0 < \tau_1 < \dots < \tau_K < 1$.
Hyper-parameters: the number of ensemble steps T ; and the hyper-parameters of the neural network models F_0 and $F_t, t \geq 1$.
Output: the learnable parameters Θ_0 and $\Theta_t, t \geq 1$ of the models F_0 and F_t .

1: Let $[\mu(x), \sigma(x)]^\top = F_0(x; \Theta_0)$ and $Q^0(x) = G_0(\mu(x), \sigma(x)) = \mu(x) + \sigma(x)[\Phi^{-1}(\tau_1), \dots, \Phi^{-1}(\tau_K)]^\top$, solve

$$\min_{\Theta_0} \frac{1}{|D_{tr}|} \sum_{i \in D_{tr}} L(y_i, Q^0(x_i)), \quad (16)$$

where the loss function L can be found in (5). $\Phi^{-1}(\tau_k)$ is the τ_k -quantile of standard Gaussian.

2: **for** $t = 1, \dots, T$ **do**
 3: Let $[a_0^t, a_1^t, a_2^t, a_3^t]^\top = F_t(x; \Theta_t)$, and $\lambda_k^t = \tanh(\sum_{l=0}^3 a_l^t (\tau_k)^l)$. Let $Q^t(x) = Q^{t-1}(x) + [g_1^t(\lambda_1^t), \dots, g_K^t(\lambda_K^t)]^\top$, where both $g_k^t(\cdot)$ and λ_k^t condition on x . The definition of $g_k^t(\cdot)$ can be found in (13). Then solve

$$\min_{\Theta_t} \frac{1}{|D_{tr}|} \sum_{i \in D_{tr}} L(y_i, Q^t(x_i)). \quad (17)$$

4: **end for**
 5: **return** the learned parameters $\Theta_t, t = 0, 1, \dots, T$. Given a new test input x , the step- T predictions $Q^T(x)$ are the final predictions of the quantiles of $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$. The computing process of $Q^T(x)$ is described as above.

E. Adaptive T Strategy

From Algorithm 1, the only hyper-parameter that is newly introduced into the EMQ approach is the number of ensemble steps T . Such a hyper-parameter is certainly crucial for the performance of the approach. Indeed, T functions as a trade-off between the distribution structure and the flexibility, because more ensemble steps will lead to more flexible distribution that is farther away from Gaussian. Here we introduce the adaptive T strategy which determines T adaptively and hence can help EMQ find the balance.

The strategy works analogously to the early stopping in deep learning, thus is a kind of implicit regularization. In this strategy, we split the training dataset D_{tr} into a new training dataset and a validation set (can be the same validation set used in the early stopping when training $F_t, t \geq 0$). The performance of the predictions $Q^t(x)$ on the validation set is tracked for all $t \geq 0$ and $t \leq T_{max}$. Under some criteria, the EMQ approach terminates the ensemble steps when $Q^t(x)$ starts to deteriorate. T_{ada} is set to be the step t at which $Q^t(x)$ is the best. The EMQ model will finally have T_{ada} ensemble steps and $Q^{T_{ada}}(x)$ is the final quantile predictions. To further reduce the risk of over-fitting, in

Algorithm 2: The Adaptive T Strategy.

Input: the inputs of the EMQ approach in Algorithm 1; and an additional validation set.

Hyper-parameters: the number of maximum ensemble steps T_{max} ; two numbers of steps $t_1 > t_2$.

Output: the adaptive number of ensemble steps T_{ada} .

1: **for** $t = 0, \dots, T_{max}$ **do**
 2: Perform step t of EMQ and evaluate $Q^t(x)$: $e_t = \text{ECE of } Q^t(x) \text{ on the validation set.}$
 3: **if** $t \geq t_1$ and $\text{avg.}(e_{t-t_2+1:t}) > \text{avg.}(e_{t-t_1+1:t-t_2})$ **then**
 4: $t' = t$
 5: **Break**
 6: **end if**
 7: **end for**
 8: $T_{ada} = \arg \min_{\{t: 0 \leq t \leq t'\}} e_t$.
 9: **return** T_{ada} . Given a new test input x , the final quantile predictions of EMQ are $Q^{T_{ada}}(x)$.

the strategy, we evaluate $Q^t(x)$ on the validation set with a metric that is different from the quantile loss L in (5) used for training $F_t, t \geq 0$. For example, we can use the metric ECE (expected calibration error, computed by calibrating for every $\tau_k \in \Pi$ and taking the average). This adaptive T strategy is summarized in Algorithm 2.

F. EMQW

In many circumstances, people are interested in tail-side prediction intervals such as the 90%-interval, 80%-interval, etc. But the quantiles at the tail sides are more difficult to predict, so we extend EMQ to a new version which places more emphasis at the tail sides than EMQ does. The only difference between this new EMQW method and the original EMQ is that EMQW uses a weighted quantile loss function instead of (5)

$$L^W(y, Q^t(x)) = \sum_{k=1}^K w_{\tau_k} L_{\tau_k}(y, q_k^t(x)), \quad (18)$$

$$L_\tau(y, q) = (y - q)(\tau - \mathbb{I}_{\{y < q\}}), \quad (19)$$

$$w_\tau = 1/\mathbb{E}_{y \sim \mathcal{N}(0,1)} [L_\tau(y, \Phi^{-1}(\tau))]. \quad (20)$$

Φ is the cumulative distribution function of $\mathcal{N}(0, 1)$. This aims to let the losses for different τ_k have a similar magnitude, because $\mathbb{E}_{y \sim \mathcal{N}(0,1)} [L_\tau(y, \Phi^{-1}(\tau))]$ may be quite different for different τ . This was also suggested by [34], [46]. Except for this difference, EMQW and EMQ share the same designs and settings. We will show in the experiments that EMQW is a better choice if tail-side quantiles should be considered, and we will use it to compare to other methods.

G. Computational Complexity

At step 0, the computational complexity of training F_0 is the same as that of training an ordinary multi-layer perception with a 2-dimensional output layer. We denote the computational cost

of training F_0 as C_0 . At step $t > 0$, the network model F_t has a much smaller structure than F_0 does. So, training F_t takes much less time. If the computational cost of training F_t is C_1 (then $C_1 \ll C_0$), the total computational cost of training EMQ or EMQW will be $C_0 + T_{ada}C_1$, and at most $C_0 + T_{max}C_1$. We will see in the experimental section that setting $T_{max} = 40$ is enough for most datasets. For space complexity, as we should store the predicted quantiles at every step, the additional space we need except for that consumed in training neural networks is $O(|D_{tr}| \cdot |\Pi| \cdot T_{ada})$, and at most $O(|D_{tr}| \cdot |\Pi| \cdot T_{max})$ ($|\Pi| = 99$ in our experiments). It is noticeable that training T_{ada} networks requires only the same space as training F_0 does because F_0 is the biggest one.

IV. EXPERIMENTS

In this section, we implement the EMQ approach and its variant EMQW, and compare it with both recent popular and classical methods on a wide range of datasets. The extensive evaluations reveal that EMQW achieves state-of-the-art performance. We also visualize the learning process and learning results of EMQW for better interpretation, and demonstrate that empirically it works as we expect. We show that it captures a variety of distribution shapes existing in real data. Consequently, we can conclude that the proposed approach owns the merits we state before.

A. Datasets and Evaluation Settings

We collect 20 UCI datasets [47] and show their information in Table I. The datasets are from different domains and can be used for comprehensive regression studies. For every dataset, we extract a short name from its full name. The numbers of samples vary from 10,000 to 515,345 and the feature dimensions vary from 5 to 465. The 20 datasets cover both medium-size and large-size datasets, with both low-dimensional and high-dimensional features.

For evaluation, we randomly select 20% from the full dataset as the testing set, and do the training and testing N_{te} times with different training/testing splits, for averaging the evaluation metrics. We set $N_{te} = 5$ for the datasets with sample size less than 100,000 and $N_{te} = 1$ for the rest: GPU, Query, Wave, Air, and Year for reducing the computational burden. In addition, the proportion of the testing set is raised from 20% to 50% for the two largest datasets Air and Year, for further reducing the computational burden. All features and labels are normalized to have sample mean 0 and sample variance 1.

B. Evaluation Metrics

To evaluate the predicted quantiles of $\mathbb{P}(y|\mathbf{X} = x)$, we use both calibration and sharpness. The probability level set is $\Pi = \{\tau_k\}_{k=1}^K = \{0.01, 0.02, \dots, 0.98, 0.99\}$ in this paper, with totally $K = 99$ levels. Given the testing set $\{(x_i, y_i)\}_{i \in D_{te}}$ and a machine learning model that is ready to predict, we temporarily denote the predicted τ -quantile of $\mathbb{P}(y|\mathbf{X} = x_i)$ as $q_\tau(x_i)$, for any $i \in D_{te}$ and $\tau \in \Pi$. Three evaluation metrics (times 100 in our tables) are used:

TABLE I
THE INFORMATION OF THE 20 UCI DATASETS USED

Dataset Name	Full Name	Sample Size	Feature Dimension
Grid	Electrical Grid Stability	10,000	12
Naval	Conditional Based Maintenance [48]	11,933	15
Appliance	Appliances energy prediction [49]	19,735	27
Election	Real-time Election [50]	21,643	23
Steel	Steel Industry Energy Consumption	35,040	6
Facebook1	Facebook Comment Volume 1 [51]	40,949	52
Facebook2	Facebook Comment Volume 2 [52]	81,312	52
PM2.5	Beijing PM2.5 [53]	41,757	10
Bio	Physicochemical Properties	45,730	9
Blog	BlogFeedback [54]	52,397	276
Consumption	Power consumption of Tetouan city [55]	52,416	5
Video	Online Video Characteristics	68,784	16
GPU	GPU kernel performance [56]	241,600	14
Query	Query Analytics Workloads [57], [58]	199,843	6
Wave	Wave Energy Converters	287,999	48
Air	Beijing Multi-Site Air-Quality Data [59]	383,585	14
Year	YearPredictionMSD	515,345	90
Relative	Relative location of CT slices	53,500	379
UjiIndoorloc	UJIIndoorLoc [60]	19,937	465
Cup98	KDD Cup 1998	95,412	326

i) *EICE (Expected Interval Calibration Error)*: Generally in uncertainty quantification, the construction of intervals is required. We compute the calibration error of the interval constructed by each pair of quantiles, e.g., $[q_{0.1}(x_i), q_{0.9}(x_i)]$, and take the average over all intervals. To be short,

$$EICE = \frac{1}{49} \sum_{k=1}^{49} \left| 1 - 2\tau_k - \frac{1}{|D_{te}|} \sum_{i \in D_{te}} \mathbb{I}_{\{q_{\tau_k}(x_i) < y_i < q_{\tau_{100-k}}(x_i)\}} \right|. \quad (21)$$

ii) *EIS (Expected Interval Sharpness)*: Sharpness measures the length of the intervals and is an evaluation metric supplement to calibration, because good calibration can be obtained by unsatisfactory intervals such as those from the unconditional distribution of y . We define

$$EIS = \frac{1}{49} \sum_{k=1}^{49} \frac{1}{|D_{te}|} \sum_{i \in D_{te}} |q_{\tau_{100-k}}(x_i) - q_{\tau_k}(x_i)|. \quad (22)$$

iii) *TICE (Tail Interval Calibration Error)*: Researchers and practitioners may be more interested in the 90%-interval, 80%-interval, etc. So we use only the intervals constructed by quantiles at the tail sides, and define

$$TICE = \frac{1}{|\Pi_{ta}|} \sum_{\tau \in \Pi_{ta}} \left| 1 - 2\tau - \frac{1}{|D_{te}|} \sum_{i \in D_{te}} \mathbb{I}_{\{q_\tau(x_i) < y_i < q_{1-\tau}(x_i)\}} \right|, \quad (23)$$

where $\Pi_{t_a} = \{0.05, 0.1, 0.15, 0.2\}$. It corresponds to the average calibration of 90%, 80%, 70%, and 60%-intervals.

C. Methods and Implementations

We implement EMQ and two variants, as well as 8 popular existing deep learning-based methods for uncertainty quantification for comparisons. In them, HNN and Deep Ensemble are variance network models. MC Dropout and Concrete Dropout are Bayesian methods. Vanilla QR, QRW, and SQR are quantile regression models. Interval Score is a recent method related to quantile regression. The implementation details of every method are summarized as follows.

EMQ: In EMQ implementation, we set the only additional hyper-parameter $T_{max} = 40$ if the dataset has a feature dimension less than 300, and $T_{max} = 200$ otherwise. For the network structure, at the step $t = 0$ for F_0 , we use 3 hidden layers with layer sizes $[8, 16, 4] \times d$ where d is the input dimension. The output dimension is 2 at $t = 0$. At the step $t > 0$ for F_t , we use 2 hidden layers with layer sizes $[16, 8]$, which are much smaller than those in F_0 . Now the output dimension is 4 at $t > 0$. For activation functions, we use tanh for all hidden layers and linear for all output layers. Softplus is applied additionally for ensuring a positive output if necessary.

EMQ₀: We also implement a special case of EMQ in which we set $T_{max} = 0$, named EMQ₀. It is equivalent to assuming $\mathbb{P}(y|\mathbf{X} = x)$ to be Gaussian all along, without further ensemble steps. We can use it to prove the advantages of EMQ. The hyper-parameters of EMQ₀ are set the same as EMQ (at the step $t = 0$).

EMQW: Except for the loss function, EMQW and EMQ share the same hyper-parameters. We aim to learn better intervals at the tail sides, and compare with other competing methods using this EMQW.

HNN: We train a mean network and a variance network alternately twice with negative log-likelihood loss, as suggested by [13]. Two hidden layers with layer sizes $[8, 4] \times d$ are used for reducing the computational cost, because HNN is the building block of Deep Ensemble whose computational complexity is the highest among all methods as it requires training multiple HNNs. Only alternately training twice is also due to this reason. For the mean network, the ReLU activation is used on hidden layers and linear on output layer, and for the variance network, tanh is used on hidden layers and Softplus on output layer.

Deep Ensemble: We train 5 HNNs with different initialization seeds for ensemble. The hyper-parameters of every HNN are exactly the same as described above. A mixture of Gaussian combining the obtained multiple Gaussian distributions is assumed and its mean and variance are computed to form an approximated Gaussian, as done in [5].

MC Dropout: The network in MC Dropout has 3 hidden layers with sizes $[8, 16, 4] \times d$ and the dropout probability is 0.1, as suggested in [9]. ReLU is applied on hidden layers and linear on output layer. Mean squared error is used for training the network. When predicting, 1,000 predictions for every input under the training mode are obtained for computing the mean and variance.

Concrete Dropout: The network in Concrete Dropout has 2 hidden layers with sizes $[8, 4] \times d$. We use the official code

released in [36]. If in some cases the network cannot give a valid training loss, we change the hidden layer sizes to be the suggested [100,80] in the referenced code.

Vanilla QR: In vanilla quantile regression, we use a network having 3 hidden layers with sizes $[8, 16, 4] \times d$, which outputs the 99 conditional quantiles directly. ReLU is applied on hidden layers and linear on output layer. The quantile loss in (5) is used to train the network.

QRW: In this new method called quantile regression with weights (QRW), we use the same network as in Vanilla QR. The only difference is that the weighted quantile loss in (18) is used instead of the un-weighted version in (5). Similarly, this will give more emphasis on the tail sides than Vanilla QR does.

SQR: In simultaneous quantile regression [26], we use the hidden layer sizes $[8, 16, 4] \times d$ for the network. ReLU is applied on hidden layers and linear on output layer. The quantile loss of a random τ of every data point is used to update the network parameters at each descent step. After training, the quantiles of any levels can be predicted.

Interval Score: In Interval Score method, we use the hidden layer sizes $[8, 16, 4] \times d$ for the network too. It outputs the 99 conditional quantiles directly as well. The only difference to Vanilla QR is that now the interval score in (1) is adopted to train the network.

Common Settings: For all the above methods, we use some common settings for all neural network training. The batch size is 128 and the learning rate is 0.01. We randomly select 20% from the training set as the validation set, which is used in the early stopping strategy. The maximum number of epochs is 1,000. All the optimization is done with Adam [61] implemented by PyTorch [62]. Initializations of all neural networks are done by Xavier Normal method [63] with specific seeds for reproducibility.

D. Performance of EMQ Family

We first examine the three models in EMQ family: EMQ₀, EMQ, and EMQW. Intuitively, EMQ₀ assumes conditional Gaussian and lacks flexibility, EMQ adds ensemble steps for seeking a balance between the distribution structure and the flexibility, and EMQW places more weights on the tail-side quantiles. We report their performance under the three evaluation metrics respectively in Tables II, III, and IV.

From the results reported, we find that EMQ significantly outperforms EMQ₀ on both calibration and sharpness (EMQ beats EMQ₀ on 15 datasets in Table II, and beats EMQ₀ on 18 in Table III), verifying that the ensemble steps in EMQ successfully add appropriate flexibility into the conditional distributions learned. EMQW performs further better than EMQ on calibration, or on both EICE and TICE metrics. It indeed gives better tail-side intervals, as we expect. On the sharpness metric EIS, EMQW performs slightly worse than EMQ, but relatively the results are close. Overall, we think EMQW is the best choice in the three, and in the following experiments, we will use EMQW as our model choice when comparing to other methods.

TABLE II
EICE RESULTS OF THE THREE MODELS IN EMQ FAMILY

Dataset	EMQ ₀	EMQ	EMQW
Grid	5.33	5.93	*3.33
Naval	18.95	*2.47	2.87
Appliance	2.27	2.26	*1.99
Election	27.22	24.62	*22.21
Steel	18.77	9.92	*8.81
Facebook1	8.18	4.88	*3.03
Facebook2	7.57	*3.26	3.74
PM2.5	1.69	1.90	*1.05
Bio	1.85	1.98	*1.10
Blog	11.57	*2.90	3.34
Consumption	*1.46	1.52	2.05
Video	9.93	6.33	*3.85
GPU	3.61	3.54	*1.67
Query	*2.15	2.15	3.50
Wave	*3.34	3.34	3.80
Air	1.37	*0.84	1.21
Year	*0.54	1.12	1.77
Relative	6.78	6.48	*3.86
Ujiindoorloc	7.42	4.37	*3.11
Cup98	27.07	*4.46	13.87

EMQW wins on 11 datasets, EMQ wins on 5, and EMQ₀ wins on 4. * is the best among three.

TABLE III
EIS RESULTS OF THE THREE MODELS IN EMQ FAMILY

Dataset	EMQ ₀	EMQ	EMQW
Grid	21.52	*20.26	20.94
Naval	19.39	*9.13	10.26
Appliance	80.41	*78.23	88.67
Election	8.42	*6.34	8.66
Steel	2.24	*1.45	1.51
Facebook1	21.63	*20.70	24.73
Facebook2	22.68	*19.50	23.26
PM2.5	51.77	*50.74	54.35
Bio	*69.91	70.08	80.51
Blog	45.04	*36.49	42.60
Consumption	92.63	*90.62	90.98
Video	*32.43	51.92	60.72
GPU	6.20	5.82	*5.81
Query	*9.35	9.35	9.44
Wave	*10.00	10.00	12.09
Air	28.25	*28.06	29.90
Year	107.60	*103.81	106.32
Relative	39.25	37.55	*20.56
Ujiindoorloc	53.30	33.11	*22.47
Cup98	39.47	*27.25	39.57

EMQW wins on 3 datasets, EMQ wins on 13, and EMQ₀ wins on 4.

E. Performance of Competing Methods

We compare the proposed EMQW to other competing methods in Tables V, VI, and VII which show EICE, EIS, and TICE results respectively. The competing methods include variance networks, Bayesian methods, quantile regression-based models, etc., which are extensive and highly representative.

From EICE results in Table V, we can find that EMQW wins the best on 13 out of 20 datasets, and wins the second best or above on 16 out of 20 datasets. On some datasets where EMQW wins the best, such as Facebook1, Facebook2, Bio, Blog, Video, GPU, Wave, and Cup98, the improvements over other methods are large, producing a decrease rate of about 50% or even 80%. Conclusively, EMQW produces superior state-of-the-art

TABLE IV
TICE RESULTS OF THE THREE MODELS IN EMQ FAMILY

Dataset	EMQ ₀	EMQ	EMQW
Grid	7.21	8.49	*4.72
Naval	18.94	*2.30	2.66
Appliance	2.97	3.59	*2.00
Election	18.60	18.63	*14.69
Steel	18.93	11.42	*10.59
Facebook1	7.61	4.74	*3.43
Facebook2	7.42	*3.43	3.62
PM2.5	2.61	3.14	*1.74
Bio	2.57	3.29	*1.30
Blog	9.32	*3.23	4.43
Consumption	*1.87	2.18	2.86
Video	8.31	6.34	*4.70
GPU	3.45	3.29	*1.34
Query	3.40	3.40	*2.41
Wave	*3.29	3.29	4.61
Air	2.14	*1.32	1.96
Year	*0.68	1.72	1.95
Relative	6.74	5.73	*4.81
Ujiindoorloc	6.68	4.23	*3.91
Cup98	18.72	*5.17	10.16

EMQW wins on 12 datasets, EMQ wins on 5, and EMQ₀ wins on 3.

performance on the predicted quantiles spreading the whole distribution.

From EIS results in Table VI, we can find that some methods give unusual sharpness results. MC Dropout and SQR give very small sharpness. However, their calibration results are unacceptable. Concrete Dropout has no good performance on sharpness, and on calibration as well. Excluding these, the remaining methods have comparable performance on the sharpness. Thus, for the remaining methods, the two calibration metrics EICE and TICE are more important now. Our EMQW achieves an appropriate balance between the calibration and sharpness.

From TICE results in Table VII, we can find that EMQW wins the best on 10 out of 20 datasets, and wins the second best or above on 16 out of 20 datasets. On some datasets where EMQW wins the best, such as Facebook1, Facebook2, Blog, GPU, Wave, and Cup98, the improvements over other methods are large too, producing a decrease rate close to or more than 50%. It indicates that EMQW produces rather good intervals of 90%, 80%, 70%, and 60% probabilities. From another perspective, QRW also places more weights on the tail-side quantiles, however, its performance is far worse than EMQW. QRW assumes no distribution structure and has great flexibility, verifying that the over-flexibility is harmful and the balance between distribution structure and the flexibility in EMQW is beneficial and necessary.

F. Performance of Ensemble Tree Models

Because the proposed EMQW is a type of ensemble (boosting) model, it will be interesting and helpful to compare it to traditional ensemble tree models such as GBDT [64] and LightGBM [65], both of which are industry-level machine learning models that are extremely popular and useful. They are thought to be the most suitable models for tabular datasets, such as the datasets used in this paper. In both GBDT and LightGBM, one can specify quantile loss as the loss function and thereby predict

TABLE V
EICE RESULTS OF NINE METHODS

Dataset \ Method	EMQW	HNN	Deep Ensemble	MC Dropout	Concrete Dropout	Vanilla QR	QRW	SQR	Interval Score
Grid	3.33	7.68	13.07	**2.03	*2.41	6.14	4.54	4.55	5.32
Naval	**2.87	*2.97	15.62	17.48	11.02	14.46	11.97	15.23	12.36
Appliance	**1.99	3.37	8.40	21.99	6.05	4.58	2.88	5.48	*2.38
Election	22.21	**8.62	23.62	21.90	*21.34	32.63	29.87	34.65	30.49
Steel	*8.81	**8.28	22.11	28.88	22.68	16.79	14.58	18.65	16.75
Facebook1	**3.03	*15.99	23.17	24.97	26.93	17.06	16.74	25.76	18.44
Facebook2	**3.74	18.90	25.33	17.15	34.26	17.50	17.54	25.96	*16.48
PM2.5	**1.05	*1.71	8.24	16.28	2.04	3.32	3.37	5.92	2.82
Bio	**1.10	3.30	7.77	24.35	2.05	2.55	2.76	*1.97	2.41
Blog	**3.34	28.07	37.39	31.42	27.33	22.23	22.57	49.47	*21.50
Consumption	2.05	1.93	**1.14	32.41	1.95	2.04	*1.15	2.11	1.51
Video	**3.85	21.47	24.68	10.45	18.92	8.39	7.59	18.11	*7.14
GPU	**1.67	*3.88	11.99	10.76	8.70	7.62	5.45	6.85	6.52
Query	*3.50	**1.02	13.10	14.96	4.44	3.70	4.38	3.63	5.56
Wave	**3.80	28.25	14.81	39.23	21.48	20.76	*11.81	27.73	17.88
Air	**1.21	3.53	7.72	7.18	2.70	2.13	1.92	*1.83	2.63
Year	*1.77	2.88	5.02	37.34	24.99	**1.57	2.11	3.38	1.84
Relative	3.86	**1.85	20.64	*2.46	16.85	4.98	3.30	8.20	3.77
Ujiindoorloc	**3.11	9.37	27.30	6.54	21.47	4.49	3.36	11.68	*3.25
Cup98	**13.87	*30.36	34.45	48.43	31.14	33.78	34.05	50.00	33.83

** is the best performance and * is the second best. EMQW wins the best on 13 out of 20 datasets, and wins the second best or above on 16 out of 20 datasets. On some datasets, the improvements over other methods are large (decrease rate over 50%), such as facebook1, facebook2, bio, blog, video, gpu, wave, and cup98.

TABLE VI
EIS RESULTS OF NINE METHODS

Dataset \ Method	EMQW	HNN	Deep Ensemble	MC Dropout	Concrete Dropout	Vanilla QR	QRW	SQR	Interval Score
Grid	20.94	**16.68	21.30	23.97	18.83	*17.38	18.52	17.71	19.72
Naval	10.26	**6.31	*8.85	27.37	1349.80	13.60	12.72	14.83	12.31
Appliance	88.67	84.52	93.84	**35.33	105.06	69.77	75.12	*67.77	77.79
Election	8.66	5.10	6.75	11.51	54749.11	2.41	*2.39	**1.64	2.49
Steel	**1.51	2.30	3.16	16.10	4.73	2.28	1.86	1.95	*1.81
Facebook1	24.73	38.24	40.72	**15.77	inf	18.88	20.50	*16.57	23.54
Facebook2	23.26	40.68	45.68	*19.13	inf	19.61	20.43	**16.62	20.27
PM2.5	54.35	65.67	71.83	**35.90	81.27	49.74	53.22	*49.33	53.27
Bio	80.51	91.93	97.79	**34.55	752080.32	*73.23	73.62	77.81	76.35
Blog	42.60	87.43	107.62	25.14	inf	*23.67	27.44	**0.71	28.49
Consumption	*90.98	95.77	98.42	**25.52	100.24	91.57	94.88	98.77	95.29
Video	60.72	86.51	89.82	**19.07	inf	53.30	52.54	*50.13	54.34
GPU	5.81	5.26	6.09	15.56	26.18	*4.09	**3.53	5.14	4.68
Query	9.44	10.29	11.96	21.85	17.37	**8.16	9.30	9.41	*8.64
Wave	12.09	**0.02	5.46	16.46	8.38	*0.43	0.50	0.47	0.49
Air	29.90	29.76	30.76	**21.71	35.59	*23.75	24.53	25.77	24.62
Year	106.32	112.28	116.45	**21.03	inf	105.59	105.05	*93.42	101.93
Relative	20.56	*10.08	37.44	18.23	179089.01	43.86	44.50	**7.16	21.75
Ujiindoorloc	22.47	69.39	122.46	37.06	inf	*17.28	24.24	**10.06	31.29
Cup98	39.57	141.73	149.39	*6.16	528.60	26.74	26.17	**0.00	26.30

** is the best performance and * is the second best. Some methods give unusual sharpness results, e.g., mc dropout or sqr gives much smaller sharpness and concrete dropout gives infinity on several datasets. Excluding these, the remaining methods give comparable sharpness results.

quantiles. We run them to predict the 99 conditional quantiles of the levels in the set Π , with each level being handled separately. A third model for comparison is the quantile regression forest proposed in [66], which is popular too.

We tune the hyper-parameters of GBDT and LightGBM on the 20% validation set: the maximum depth of a single tree, and the number of trees. For quantile regression forest, the fraction of features used in the node splitting is tuned as well. For each $\tau \in \Pi$, this tuning will be done separately. We report their results of EICE and EIS together in Table VIII. From the table, we find that for EICE, our EMQW wins the best on 16 datasets out of 20. For EIS, EMQW wins the best on 9 datasets out of

20. Conclusively, EMQW shows better performance than these traditional models do.

G. Interpretation of Learning Results

We visualize the learning process and some learning results of EMQW for better interpretation. This helps us understand how EMQW works and why it gives superior performance.

Performance Variation of EMQW With Step t : We are interested in the performance variation of EMQW on the testing set with respect to t . We plot in Fig. 2 the varying performance of the metric EICE, as the step t increases from 0 to the final T_{ada} .

TABLE VII
TICE RESULTS OF NINE METHODS

Dataset \ Method	EMQW	HNN	Deep Ensemble	MC Dropout	Concrete Dropout	Vanilla QR	QRW	SQR	Interval Score
Grid	4.72	10.88	13.14	**1.74	*2.99	7.02	5.51	5.68	5.98
Naval	*2.66	**2.46	17.89	18.53	12.56	15.11	16.23	11.93	14.79
Appliance	**2.00	*2.66	8.02	32.18	6.18	5.82	3.29	8.97	3.30
Election	*14.69	**6.66	23.08	18.49	19.18	32.72	25.69	52.57	26.10
Steel	*10.59	**6.06	23.07	20.47	19.73	13.93	13.75	15.54	11.76
Facebook1	**3.43	13.34	18.64	36.87	18.18	17.89	*8.71	39.24	15.15
Facebook2	**3.62	14.69	19.21	23.06	22.83	10.83	*10.22	39.28	10.25
PM2.5	**1.74	*1.80	10.57	24.02	3.13	4.54	3.17	6.64	3.21
Bio	**1.30	2.61	7.33	35.70	3.55	3.28	3.36	*1.72	2.34
Blog	**4.43	17.93	22.13	46.60	21.53	18.78	22.22	74.19	*10.10
Consumption	2.86	2.19	1.97	48.20	2.21	1.97	**1.15	2.94	*1.41
Video	**4.70	14.15	15.60	14.36	13.77	6.12	*5.77	26.52	6.69
GPU	**1.34	*2.71	16.04	13.05	9.26	5.75	4.51	5.09	4.70
Query	*2.41	**0.76	15.57	15.39	5.31	4.93	3.43	4.86	8.55
Wave	**4.61	20.69	16.25	25.00	21.67	15.98	*7.59	23.86	14.83
Air	1.96	3.80	8.88	12.81	3.62	*1.11	1.35	**1.08	3.08
Year	*1.95	4.04	6.78	55.60	19.27	**1.43	2.59	4.58	1.95
Relative	4.81	**1.59	20.66	*2.31	14.25	5.06	3.83	6.69	3.29
Ujiindoorloc	*3.91	5.93	22.75	7.06	17.73	6.50	4.10	19.69	**3.12
Cup98	**10.16	*18.73	20.42	72.65	20.48	26.93	26.67	75.00	25.69

They correspond to the average calibration of 90%, 80%, 70%, and 60%-intervals. ** is the best performance and * is the second best. EMQW wins the best on 10 out of 20 datasets, and wins the second best or above on 16 out of 20 datasets. On some datasets, the improvements over other methods are large (decrease rate close to or over 50%), such as facebook1, facebook2, blog, gpu, wave, and cup98.

TABLE VIII
EICE / EIS RESULTS OF EMQW AND THE THREE ENSEMBLE TREE MODELS

Dataset \ Method	EMQW	GBDT + Quantile Loss	LightGBM + Quantile Loss	Quantile Regression Forest
Grid	*3.33 / *20.94	6.34 / 39.10	7.75 / 34.07	11.29 / 72.67
Naval	*2.87 / *10.26	5.38 / 39.85	4.73 / 24.15	11.59 / 36.82
Appliance	*1.99 / 88.67	2.87 / 94.29	4.18 / *80.60	11.17 / 93.34
Election	*22.21 / 8.66	40.00 / 9.57	39.43 / 10.27	49.96 / *0.17
Steel	8.81 / *1.51	4.29 / 8.61	*4.01 / 5.95	14.55 / 5.57
Facebook1	*3.03 / 24.73	29.00 / 18.39	4.75 / *18.05	32.02 / 21.26
Facebook2	*3.74 / 23.26	28.65 / 17.81	5.31 / *17.29	31.71 / 20.12
PM2.5	*1.05 / *54.35	1.85 / 81.76	2.67 / 74.61	5.31 / 81.51
Bio	*1.10 / *80.51	2.21 / 109.07	3.02 / 97.76	4.58 / 95.48
Blog	*3.34 / 42.60	33.17 / 22.09	7.44 / *21.58	35.67 / 23.01
Consumption	2.05 / 90.98	*1.65 / 97.71	2.27 / 91.50	3.92 / *89.93
Video	*3.85 / 60.72	6.99 / 62.29	6.99 / 57.72	13.27 / *56.19
GPU	*1.67 / 5.81	3.04 / 30.59	3.03 / 22.13	6.77 / *4.77
Query	3.50 / *9.44	*2.58 / 18.94	2.98 / 14.98	15.62 / 15.13
Wave	*3.80 / 12.09	4.98 / 4.01	5.01 / *2.85	14.98 / 6.70
Air	*1.21 / *29.90	1.43 / 32.91	1.38 / 30.92	3.99 / 36.40
Year	1.77 / 106.32	*1.29 / 107.01	1.72 / *104.43	2.87 / 117.85
Relative	*3.86 / *20.56	4.06 / 36.43	4.89 / 29.60	21.42 / 22.11
Ujiindoorloc	*3.11 / *22.47	3.62 / 33.92	4.16 / 24.66	5.27 / 34.37
Cup98	*13.87 / 39.57	49.76 / *28.49	39.11 / 29.32	49.78 / 31.02

* is the best performance among the four. For eice metric, EMQW wins the best on 16 datasets out of 20. For eis metric, EMQW wins the best on 9 datasets out of 20.

Each subfigure corresponds to one dataset, and different lines in one subfigure represent different training/testing splits. In the vast majority of lines, EICE decreases as t increases with very few exceptions, indicating that the ensemble steps in EMQW work correctly as we expect and let the model gradually learn better $\mathbb{P}(\mathbf{y}|\mathbf{X} = x)$.

Densities Given by EMQW: As we have always claimed, the conditional distribution given by EMQW will achieve adaptive flexibility. We plot in Fig. 3 the densities implied by the conditional quantiles predicted by EMQW to examine this (the density is the derivative of the inverse of quantile function, so

easy to obtain numerically). One can see that EMQW can learn distributions with various shapes: exactly Gaussian, approximately Gaussian, sharp peak, asymmetry, long tail, finite bound, multi-modality, etc. This variety may even appear within one dataset, showing the complexity of real data and the ability of EMQW.

Densities Given by Other Methods: In contrast, we have claimed that many quantile-based methods with over-flexibility will give chaotic or invalid densities. We plot in Fig. 4 the densities implied by the quantiles predicted by the four quantile-based methods to verify this. Variance networks and the two

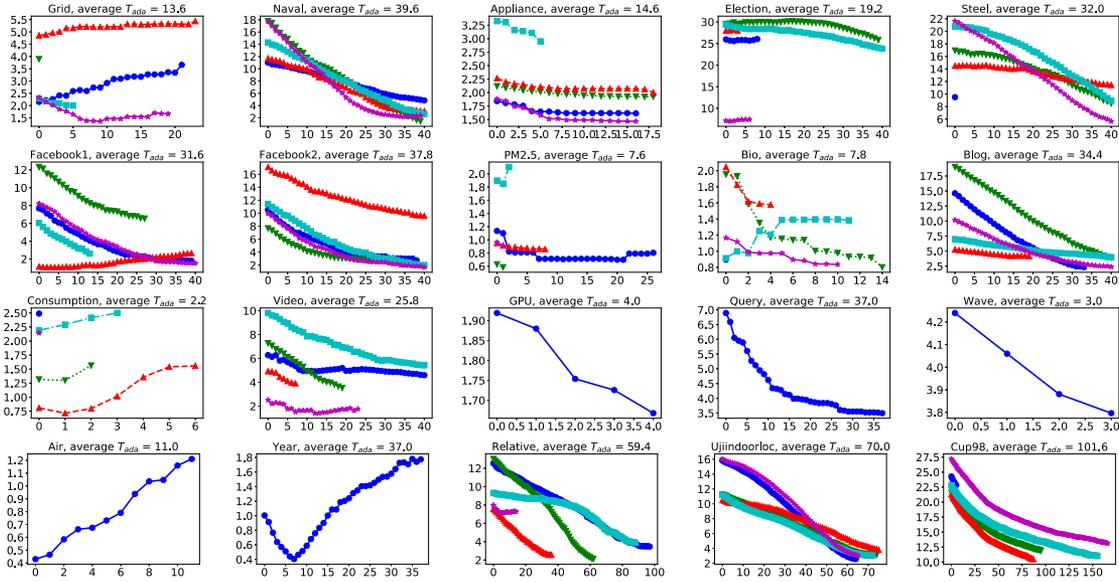


Fig. 2. The changing of EICE on the testing set as the step t increases from 0 to T_{ada} in EMQW. Each subfigure corresponds to one dataset, and the lines of different colors and styles within one subfigure represent different training/testing splits. In the vast majority of lines, EICE decreases as t increases with very few exceptions, indicating that the ensemble steps in EMQW work correctly as we expect.

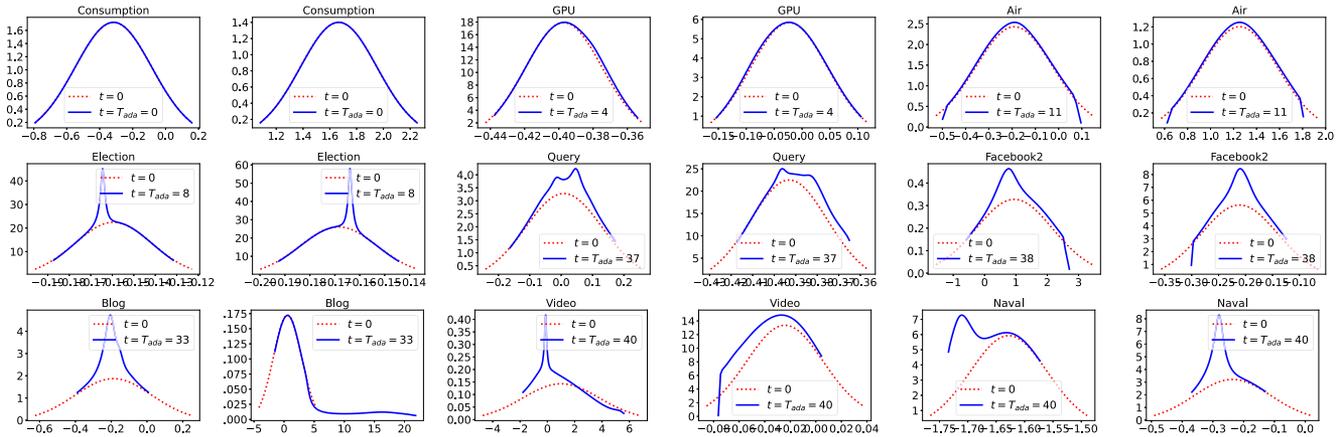


Fig. 3. Densities implied from the 99 conditional quantiles predicted by EMQW. For each dataset, we select two data points from the testing set and plot their densities of $\mathbb{P}(y|\mathbf{X} = x)$. Red lines are the results at $t = 0$ and blue ones are the results at $t = T_{ada}$. The first row of subfigures: EMQW produces exactly Gaussian, or approximately Gaussian distributions with slight differences. The second row: the densities show various peaks and asymmetric shapes. The third row: the densities show more diverse shapes, such as a long right tail, a left bound, and a multi-peak shape. The variety may occur within one dataset, revealing the complexity of real data and the ability of EMQW.

Bayesian (Dropout) methods all produce Gaussian densities, so unnecessary to plot. We select one data point from the testing set of each of the three datasets: Naval, Consumption, and Air. Its densities given by the four methods are shown in Fig. 4. From the figure, we can see they all produce rough or even invalid densities. This phenomenon is common, not just appear in these datasets.

V. DISCUSSION ON ASYMPTOTIC GUARANTEE

It is appealing to know for sure that the conditional quantiles (or conditional densities implied) approximate the true ones

asymptotically. Actually, we are looking for a function $f(\cdot) = [f_1(\cdot), \dots, f_K(\cdot)]^T$ to minimize: $\min_{f \in \mathcal{F}} \mathbb{E}[L(\mathbf{y}, f(\mathbf{X}))] = \sum_{k=1}^K \mathbb{E}[L_{\tau_k}(\mathbf{y}, f_k(\mathbf{X}))]$. Obviously, if we consider a sufficiently rich \mathcal{F} , the optimal $f_k(x)$ should be the τ_k -quantile of $\mathbb{P}(y|\mathbf{X} = x)$. In our method, \mathcal{F} consists of those functions defined via a complex ensemble of $T + 1$ neural networks. If we denote the actual τ_k -quantile of $\mathbb{P}(y|\mathbf{X} = x)$ as $q_k(x)$ and further denote $Q(x) = [q_1(x), \dots, q_K(x)]^T$, the consistency means that the estimated $Q^n(\cdot)$ given by our algorithm based on the dataset $\mathcal{D}_n = \{(\mathbf{X}_i, \mathbf{y}_i)\}_{i=1}^n$ and function class \mathcal{F}_n should satisfy: $\lim_{n \rightarrow \infty} \mathbb{E}[\int_{\mathbb{R}^d} \|Q^n(x) - Q(x)\|^2 \mu(dx)] = 0$, for all

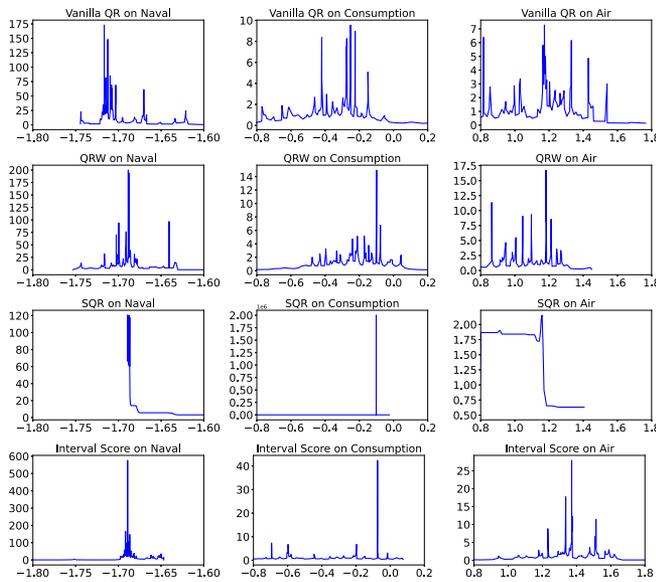


Fig. 4. Densities implied from the 99 conditional quantiles predicted by four quantile-based methods: Vanilla QR, QRW, SQR, and Interval Score. Variance networks and Dropout methods all produce Gaussian densities, so unnecessary to plot. We select one data point from the testing set of the dataset: Naval, Consumption, and Air, and plot its densities of $\mathbb{P}(y|\mathbf{X} = x)$ given by the four methods. We can see that all of them produce rough or even invalid densities.

distributions of (\mathbf{X}, y) with $\mathbb{E}[y^2] < \infty$. Here μ denotes the distribution of \mathbf{X} . This is called the weak universal consistency of Q^n (see [67] for the strong version). We assume \mathcal{F}_n consists of those functions having T_n ensemble steps, and $T_n \rightarrow \infty$ as $n \rightarrow \infty$.

The multi-quantile regression is not new in the literature. For example, [34], [46], [68], [69] all adopted similar objectives for this. However, they share some common weaknesses: i) the function class \mathcal{F} they assume is linear, linear spline, or kernel; ii) they all add explicit monotonic constraints on quantiles in the optimization; and iii) they all have no asymptotic analysis for the non-parametric estimators. Our method overcomes the first two weaknesses. For the third one, [70], [71] established the consistency of polynomial and neural-network estimators of conditional quantile, but with only one quantile level. Here we face two challenges: i) the multi-quantile regression with many levels; and ii) the function class we consider is the ensemble of many (small) neural networks, which is a novel and unique design of us.

There have been some works analyzing the asymptotic properties of neural networks (NNs). [67] gave the consistency results of infinitely wide NNs in ordinary regression problem with least square loss. Very recently, [72] showed that wide and deep NNs achieve consistency for classification and excessively deep NNs are harmful for regression. [73] analyzed the consistency of deep convolutional NNs. Back to our problem, to prove the consistency of our estimator, we may need finishing the following three steps: i) decompose $\int_{\mathbb{R}^d} \|Q^n(x) - Q(x)\|^2 \mu(dx)$ into the approximation error and estimation error; ii) prove that our function class is dense, or, our ensemble of NNs can approximate

any continuous function with a compact domain; and iii) deal with the bound on the covering numbers of our function class. Finishing each step will not be trivial and straightforward, especially, the third step will be challenging. Again it is because the function class we consider is the ensemble of many (small) NNs, which is quite different from the infinitely wide NN used in many theoretical works. Given these facts, we leave the asymptotic guarantee analysis to future work.

VI. CONCLUSION

To conclude, we propose an ensemble multi-quantiles approach that can do adaptively flexible distribution prediction. It is motivated by the observations that: i) Gaussian assumption for the conditional distribution lacks flexibility; and ii) over-flexibility may be harmful for generalization, such as the invalid densities obtained by many quantile regression methods. We design a boosting framework for multi-quantiles prediction with an initial step and the ensemble steps. Our method is also motivated by the demand for a built-in solution to the quantile-crossing issue.

The whole ensemble approach is intuitive and interpretable, in which in the initial step a relatively strong learner produces a Gaussian distribution, and in the ensemble steps T_{ada} weak learners adaptively seek a balance between the distribution structure and the flexibility. We verify these characteristics by experiments, and compare it to recent popular competing methods. Conclusively, the EMQW approach achieves start-of-the-art performance among all, under the metrics of calibration, sharpness, and tail-side calibration. Traditional ensemble tree models are also compared.

Through visualization, we find that our approach can learn distributions with various density shapes including Gaussian, approximately Gaussian, sharp peak, asymmetry, long tail, finite bound, multi-modality, etc. This variety may appear even in one dataset, confirming the complexity of real data and the necessity and merits of such an adaptively flexible ensemble approach. Although we run the experiments with tabular data, the proposed approach is general and can be applied to other formats of data such as images and texts.

REFERENCES

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [2] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2017, pp. 1321–1330.
- [3] R. Rahaman and A. Thiery, "Uncertainty quantification and deep ensembles," in *Proc. Adv. Neural Inf. Process. Syst.*, 2021, pp. 20063–20075.
- [4] E. Begoli, T. Bhattacharya, and D. Kuznezov, "The need for uncertainty quantification in machine-assisted medical decision making," *Nature Mach. Intell.*, vol. 1, no. 1, pp. 20–23, 2019.
- [5] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 6405–6416.
- [6] A. J. McNeil, R. Frey, and P. Embrechts, *Quantitative Risk Management: Concepts, Techniques and Tools-Revised Edition*. Princeton, NJ, USA: Princeton Univ. Press, 2015.
- [7] D. Duffie and J. Pan, "An overview of value at risk," *J. Derivatives*, vol. 4, no. 3, pp. 7–49, 1997.
- [8] C. Acerbi and D. Tasche, "On the coherence of expected shortfall," *J. Bank. Finance*, vol. 26, no. 7, pp. 1487–1503, 2002.

- [9] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2016, pp. 1050–1059.
- [10] W. J. Maddox, P. Izmailov, T. Garipov, D. P. Vetrov, and A. G. Wilson, "A simple baseline for Bayesian uncertainty in deep learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, Art. no. 1179.
- [11] M. Sensoy, L. Kaplan, and M. Kandemir, "Evidential deep learning to quantify classification uncertainty," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 3183–3193.
- [12] V. Kuleshov, N. Fenner, and S. Ermon, "Accurate uncertainties for deep learning using calibrated regression," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2018, pp. 2796–2804.
- [13] N. Skafté, M. Jørgensen, and S. Hauberg, "Reliable training and estimation of variance networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, Art. no. 568.
- [14] S. Zhao, T. Ma, and S. Ermon, "Individual calibration with randomized forecasting," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2020, pp. 11387–11397.
- [15] J. M. Hernández-Lobato and R. Adams, "Probabilistic backpropagation for scalable learning of Bayesian neural networks," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2015, pp. 1861–1869.
- [16] G. Shafer and V. Vovk, "A tutorial on conformal prediction," *J. Mach. Learn. Res.*, vol. 9, no. 3, pp. 371–421, 2008.
- [17] V. Balasubramanian, S.-S. Ho, and V. Vovk, *Conformal Prediction for Reliable Machine Learning: Theory, Adaptations and Applications*. Oxford, U.K.: Newnes, 2014.
- [18] H. Papadopoulos, V. Vovk, and A. Gammerman, "Regression conformal prediction with nearest neighbours," *J. Artif. Intell. Res.*, vol. 40, pp. 815–840, 2011.
- [19] J. Lei, J. Robins, and L. Wasserman, "Distribution-free prediction sets," *J. Amer. Statist. Assoc.*, vol. 108, no. 501, pp. 278–287, 2013.
- [20] Y. Romano, E. Patterson, and E. Candes, "Conformalized quantile regression," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, Art. no. 318.
- [21] P. Cui, W. Hu, and J. Zhu, "Calibrated reliable regression using maximum mean discrepancy," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, pp. 17164–17175.
- [22] T. Zhou, Y. Li, Y. Wu, and D. Carlson, "Estimating uncertainty intervals from collaborating networks," *J. Mach. Learn. Res.*, vol. 22, no. 257, pp. 1–47, 2021.
- [23] F. Wenzel, J. Snoek, D. Tran, and R. Jenatton, "Hyperparameter ensembles for robustness and uncertainty quantification," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, pp. 6514–6527.
- [24] T. Pearce, A. Brintrup, M. Zaki, and A. Neely, "High-quality prediction intervals for deep learning: A distribution-free, ensemble approach," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2018, pp. 4075–4084.
- [25] H. Song, T. Diethe, M. Kull, and P. Flach, "Distribution calibration for regression," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2019, pp. 5897–5906.
- [26] N. Tagasovska and D. Lopez-Paz, "Single-model uncertainties for deep learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, Art. no. 576.
- [27] Y. Chung, W. Neiswanger, I. Char, and J. Schneider, "Beyond pinball loss: Quantile methods for calibrated uncertainty quantification," in *Proc. Adv. Neural Inf. Process. Syst.*, 2021, pp. 10971–10984.
- [28] S. Feldman, S. Bates, and Y. Romano, "Improving conditional coverage via orthogonal quantile regression," in *Proc. Adv. Neural Inf. Process. Syst.*, 2021, pp. 2060–2071.
- [29] R. Koenker and G. Bassett Jr, "Regression quantiles," *Econometrica: J. Econometric Soc.*, vol. 46, pp. 33–50, 1978.
- [30] R. Koenker and K. F. Hallock, "Quantile regression," *J. Econ. Perspectives*, vol. 15, no. 4, pp. 143–156, 2001.
- [31] X. He, "Quantile curves without crossing," *Amer. Statistician*, vol. 51, no. 2, pp. 186–192, 1997.
- [32] I. Takeuchi, Q. V. Le, T. D. Sears, and A. J. Smola, "Nonparametric quantile estimation," *J. Mach. Learn. Res.*, vol. 7, pp. 1231–1264, 2006.
- [33] Y. Liu and Y. Wu, "Stepwise multiple quantile regression estimation using non-crossing constraints," *Statist. Interface*, vol. 2, no. 3, pp. 299–310, 2009.
- [34] H. D. Bondell, B. J. Reich, and H. Wang, "Noncrossing quantile regression curve estimation," *Biometrika*, vol. 97, no. 4, pp. 825–838, 2010.
- [35] V. Chernozhukov, I. Fernández-Val, and A. Galichon, "Quantile and probability curves without crossing," *Econometrica*, vol. 78, no. 3, pp. 1093–1125, 2010.
- [36] Y. Gal, J. Hron, and A. Kendall, "Concrete dropout," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 3584–3593.
- [37] R. Koenker and Z. Xiao, "Quantile autoregression," *J. Amer. Statist. Assoc.*, vol. 101, no. 475, pp. 980–990, 2006.
- [38] T. Gneiting and A. E. Raftery, "Strictly proper scoring rules, prediction, and estimation," *J. Amer. Statist. Assoc.*, vol. 102, no. 477, pp. 359–378, 2007.
- [39] J. Gasthaus et al., "Probabilistic forecasting with spline quantile function RNNs," in *Proc. 22nd Int. Conf. Artif. Intell. Statist.*, PMLR, 2019, pp. 1901–1910.
- [40] P. Theodossiou, "Financial data and the skewed generalized T distribution," *Manage. Sci.*, vol. 44, no. 12-part-1, pp. 1650–1661, 1998.
- [41] T. G. Bali, H. Mo, and Y. Tang, "The role of autoregressive conditional skewness and kurtosis in the estimation of conditional VaR," *J. Bank. Finance*, vol. 32, no. 2, pp. 269–282, 2008.
- [42] X. Yan, Q. Wu, and W. Zhang, "Cross-sectional learning of extremal dependence among financial assets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 3852–3862.
- [43] A. Der Kiureghian and O. Ditlevsen, "Aleatory or epistemic? Does it matter?," *Statist. Saf.*, vol. 31, no. 2, pp. 105–112, 2009.
- [44] E. Hüllermeier and W. Waegeman, "Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods," *Mach. Learn.*, vol. 110, pp. 457–506, 2021.
- [45] S. Lahlou et al., "DEUP: Direct epistemic uncertainty prediction," *Trans. Mach. Learn. Res.*, 2023.
- [46] Y. Liu and Y. Wu, "Simultaneous multiple non-crossing quantile regression estimation using kernel constraints," *J. Nonparametric Statist.*, vol. 23, no. 2, pp. 415–437, 2011.
- [47] D. Dua and C. Graff, "UCI machine learning repository," 2017. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [48] A. Coraddu, L. Oneto, A. Ghio, S. Savio, D. Anguita, and M. Figari, "Machine learning approaches for improving condition-based maintenance of naval propulsion plants," in *Proc. Inst. Mech. Engineers, Part M: J. Eng. Maritime Environ.*, vol. 230, no. 1, pp. 136–153, 2016.
- [49] L. M. Candanedo, V. Feldheim, and D. Deramaix, "Data driven prediction models of energy use of appliances in a low-energy house," *Energy Buildings*, vol. 140, pp. 81–97, 2017.
- [50] N. Moniz, "Real-time 2019 portuguese parliament election results dataset," 2019, *arXiv:1912.08922*.
- [51] K. Singh, R. K. Sandhu, and D. Kumar, "Comment volume prediction using neural networks and decision trees," in *Proc. IEEE UKSim-AMSS 17th Int. Conf. Comput. Modelling Simul.*, 2015, pp. 15–20.
- [52] K. Singh, "Facebook comment volume prediction," *Int. J. Simul. Syst. Sci. Technol.*, vol. 16, no. 5, pp. 16.1–16.9, 2015.
- [53] X. Liang et al., "Assessing Beijing's pm2.5 pollution: Severity, weather impact, apec and winter heating," in *Proc. Roy. Soc. A: Math. Phys. Eng. Sci.*, vol. 471, no. 2182, 2015, Art. no. 20150257.
- [54] K. Buza, "Feedback prediction for blogs," in *Proc. Data Anal. Mach. Learn. Knowl. Discov.*, Springer, 2013, pp. 145–152.
- [55] A. Salam and A. El Hibaoui, "Comparison of machine learning algorithms for the power consumption prediction: case study of Tetouan City-," in *Proc. 6th Int. Renewable Sustain. Energy Conf.*, 2018, pp. 1–5.
- [56] R. Ballester-Ripoll, E. G. Paredes, and R. Pajarola, "Sobol tensor trains for global sensitivity analysis," *Rel. Eng. Syst. Saf.*, vol. 183, pp. 311–322, 2019.
- [57] F. Savva, C. Anagnostopoulos, and P. Triantafillou, "Explaining aggregates for exploratory analytics," in *Proc. IEEE Int. Conf. Big Data*, 2018, pp. 478–487.
- [58] C. Anagnostopoulos, F. Savva, and P. Triantafillou, "Scalable aggregation predictive analytics: A query-driven machine learning approach," *Appl. Intell.*, vol. 48, pp. 2546–2567, 2018.
- [59] S. Zhang, B. Guo, A. Dong, J. He, Z. Xu, and S. X. Chen, "Cautionary tales on air-quality improvement in Beijing," in *Proc. Roy. Soc. A: Math. Phys. Eng. Sci.*, vol. 473, no. 2205, 2017, Art. no. 20170457.
- [60] J. Torres-Sospedra et al., "UJIIndoorLoc: A new multi-building and multi-floor database for WLAN fingerprint-based indoor localization problems," in *Proc. Int. Conf. Indoor Positioning Indoor Navigation*, 2014, pp. 261–270.
- [61] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
- [62] A. Paszke et al., "PyTorch: An imperative style, high-performance deep learning library," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 8024–8035.
- [63] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proc. 13th Int. Conf. Artif. Intell. Statist. Workshop*, 2010, pp. 249–256.

- [64] J. H. Friedman, "Greedy function approximation: A gradient boosting machine," *Ann. Statist.*, vol. 29, pp. 1189–1232, 2001.
- [65] G. Ke et al., "LightGBM: A highly efficient gradient boosting decision tree," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 3149–3157.
- [66] N. Meinshausen, "Quantile regression forests," *J. Mach. Learn. Res.*, vol. 7, pp. 983–999, 2006.
- [67] L. Györfi, M. Köhler, A. Krzyżak, and H. Walk, *A Distribution-Free Theory of Nonparametric Regression*, vol. 1. Berlin, Germany: Springer, 2002.
- [68] H. Zou and M. Yuan, "Composite quantile regression and the oracle model selection theory1," *Ann. Statist.*, vol. 36, no. 3, pp. 1108–1126, 2008.
- [69] S. J. Moon, J.-J. Jeon, J. S. H. Lee, and Y. Kim, "Learning multiple quantiles with neural networks," *J. Comput. Graphical Statist.*, vol. 30, no. 4, pp. 1238–1248, 2021.
- [70] P. Chaudhuri, "Nonparametric estimates of regression quantiles and their local Bahadur representation," *Ann. Statist.*, vol. 19, no. 2, pp. 760–777, 1991.
- [71] H. White, "Nonparametric estimation of conditional quantiles using neural networks," in *Proc. Comput. Sci. Statist. Statist. Many Parameters: Curves Images Spatial Models*, Springer, 1992, pp. 190–199.
- [72] A. Radhakrishnan, M. Belkin, and C. Uhler, "Wide and deep neural networks achieve consistency for classification," in *Proc. Nat. Acad. Sci.*, vol. 120, no. 14, 2023, Art. no. e2208779120.
- [73] S.-B. Lin, K. Wang, Y. Wang, and D.-X. Zhou, "Universal consistency of deep convolutional neural networks," *IEEE Trans. Inf. Theory*, vol. 68, no. 7, pp. 4610–4617, Jul. 2022.



Yonghua Su received the BS degree in mathematics from Wuhan University, China, in 2019. He is currently working toward the PhD degree with the Institute of Statistics and Big Data, Renmin University of China. His research interests include uncertainty quantification and generative models in machine learning.



Wenxuan Ma received the BS degree in statistics from the Minzu University of China, in 2020. She is currently working toward the PhD degree with the Institute of Statistics and Big Data, Renmin University of China. Her research interests include uncertainty quantification and tree-based models in machine learning, and their applications in finance such as empirical asset pricing.



Xing Yan received the PhD degree from the Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, in 2019. He is currently an assistant professor with the Institute of Statistics and Big Data, Renmin University of China. His research interests include generative learning, causal learning, OOD generalization, uncertainty quantification, and their applications in finance and business. He published papers in *Journal of Economic Dynamics and Control*, *NeurIPS*, *AAAI*, *ECCV*, etc.