On the Impact of Weight Quantization on Deep Neural Network Uncertainty

Shuang Liang[1,2], **Xun Lu**[2], **Zi-Ang Liu**[3], **Ming-Liang Wang**[3], **Yan Lyu**[3*], **Shao-Qun Zhang**[1,2*]

[1]National Key Laboratory for Novel Software Technology, Nanjing University, China
[2]School of Intelligent Science and Technology, Nanjing University, China
[3]China Mobile Zijin Innovation Institute, China
lvyansgs@js.chinamobile.com, zhangsq@lamda.nju.edu.cn

Abstract

Weight quantization (WQ) is a key technique for lightweight Deep Neural Network (DNN) computations. While existing algorithms often pursue memory compression and inference acceleration with accuracy comparable to full-precision models, the effect of WQ on DNN uncertainty remains largely unexplored. In this paper, we quantify the impact of WQ on DNN uncertainty through the novel Exact Moment Propagation (EMP) uncertainty estimator. It is observed that WQ significantly increases DNN uncertainty. Based on the EMP estimator, we propose the MOMent Alignment (MOMA) to reduce WQ-induced uncertainty and preserve the accuracy of weight-quantized DNNs. Empirical results across various DNN architectures and datasets validate the effectiveness of both EMP and MOMA methods.

1 Introduction

In recent years, Weight Quantization (WQ) has attracted increasing interest in lightweight computations of Deep Neural Networks (DNNs), particularly in resource-constrained environments (Gope, Dasika, and Mattina 2019; Shen et al. 2024). Seminal studies show that WQ algorithms achieve at least 16x model compression (Li et al. 2016) and 3x training acceleration (Courbariaux, Bengio, and David 2015) while maintaining comparable accuracy to full-precision counterparts (Zhu et al. 2017; Li et al. 2022); however, the extent to which WQ techniques affect the uncertainty associated with specific outcomes remains largely unexplored.

This work focuses on quantifying the impact of WQ on DNN uncertainty. In the past decades, significant efforts have been made to quantify uncertainty in DNNs without WQ, such as Deep Ensembles (Lakshminarayanan, Pritzel, and Blundell 2017) and Moment Propagation (MP) (Wu et al. 2019). In addition, research has also examined uncertainty quantification in weight-quantized networks, such as BLRNet (Peters and Welling 2018) and BayesBiNN (Meng, Bachmann, and Khan 2020). However, a unified measure to quantify model uncertainty before and after WQ is lacking. Without this measure, one cannot fairly quantify the impact of WQ on DNN uncertainty.

*Corresponding authors.

In this paper, we propose the Exact Moment Propagation (EMP), a theoretical uncertainty quantification method that is independent of weight precision. Based on the EMP, we observe that quantized models exhibit greater uncertainty than their full-precision counterparts. Inspired by research on uncertainty-aware training for neural networks (Tabarisaadi et al. 2022; Einbinder et al. 2022), we propose the MOMent Alignment (MOMA), an optimization method compatible with various weight quantization training algorithms, to maintain accuracy and reduce uncertainty in quantized DNNs. This work fills the gap in quantifying and optimizing the impact of WQ on DNN uncertainty.

Our main contributions are illustrated in Figure 1(a) and further summarized as follows.

- We propose the EMP for DNN uncertainty quantification that is independent of weight precision in Algorithm 1. Empirical evidence in Subsection 5.1 shows that EMP achieves near-ideal uncertainty estimation performance, surpassing the MP, and that quantized DNNs exhibit greater uncertainty than their full-precision versions.
- We propose the MOMA for uncertainty optimization in Algorithm 2. Experiments in Subsection 5.2 show that MOMA greatly reduces the uncertainty of quantized models across different datasets.

The rest of this paper is organized as follows. Section 2 reviews related work. Section 3 and Section 4 formally introduce the EMP and MOMA methods, respectively. Section 5 presents experiments on real-world datasets to validate the effectiveness of the EMP and MOMA methods. Section 6 concludes this work.

2 Related Work

Weight Quantization. WQ aims to encode model weights into a format of k bits, where $k = 1$ and $k = 1.58$ often lead to extreme compression. For $k = 1$, the weights of DNNs are compressed into binary values, such as $\{0, 1\}$ or $\{-1, 1\}$. Seminal works include BinaryConnect (Courbariaux, Bengio, and David 2015), Binarized Neural Networks (Hubara et al. 2016), and XNOR-Net (Rastegari et al. 2016). For $k = 1.58$, the weights are compressed into ternary values $\{-1, 0, 1\}$. Early milestone works include Ternary Weight Networks (Li et al. 2016), Trained

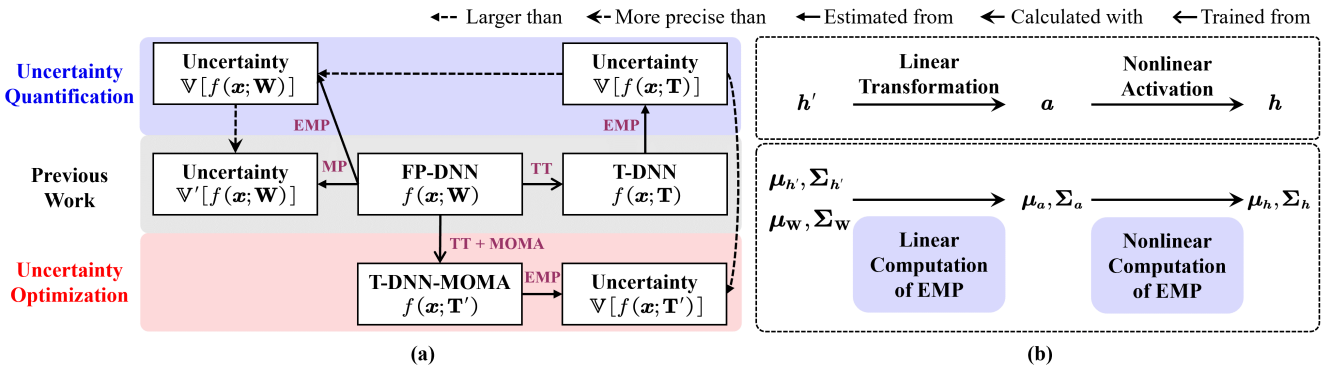

Figure 1: (a) Overview of our work, where TT is short for Ternary Training, and (b) the workflow of the EMP.

Ternary Quantization (Zhu et al. 2017), and Ternary Connect (Lin et al. 2016). In recent years, ternary WQ has been successfully applied to transformer architectures (Vaswani et al. 2017), such as TernaryBERT (Zhang et al. 2020) and Q-ViT (Li et al. 2022). TernaryCLIP (Zhang et al. 2025) achieves extremely low-cost and high-efficiency computations in image-text contrastive modeling.

Uncertainty Quantification. To quantify the uncertainty in DNNs without WQ, Deep Ensembles (Lakshminarayanan, Pritzel, and Blundell 2017) trains independent networks in parallel, estimating predictive uncertainty via negative log likelihood. PL-DNN (Bibi, Alfadly, and Ghanem 2018) and MP (Wu et al. 2019) propagate input mean and covariance through the network to derive predictive uncertainty from output covariance. For uncertainty quantification in weight-quantized DNNs, BLRNet (Peters and Welling 2018) and BayesBiNN (Meng, Bachmann, and Khan 2020) adopt the Bayesian Neural Networks paradigm (Jospin et al. 2022), sampling the posterior distribution of weights to estimate predictive uncertainty through output covariance. Despite these advances, a unified uncertainty quantification method independent of weight precision is still lacking, precluding fair comparisons of DNN uncertainty before and after WQ.

3 Uncertainty Quantification

In this section, we propose the Exact Moment Propagation (EMP) method for quantifying the uncertainty of DNNs. In contrast to previous studies, which were designed solely for full-precision models and provided either exact output covariance under strong assumptions (Bibi, Alfadly, and Ghanem 2018) or approximated output covariance (Wu et al. 2019), the proposed EMP requires only mild assumptions, improves the precision of uncertainty estimation, and extends its applicability to quantized models. The EMP method consists of the linear and nonlinear computations, whose workflow is shown in Figure 1(b).

3.1 Problem Formulation

This subsection introduces the problem formulation of WQ and uncertainty quantification of DNNs.

Ternary Weight Quantization. Throughout this work, we focus on the ternary quantization and begin with an L-layer DNN $f(\boldsymbol{x};\mathbf{W})$, of which the forward propagation follows

$$\boldsymbol{h}^{(0)} = \boldsymbol{x}\ ,\ \boldsymbol{a}^{(l)} = \mathbf{W}^{(l)}\boldsymbol{h}^{(l-1)} + \boldsymbol{b}^{(l)}\ ,\ \boldsymbol{h}^{(l)} = \tau\left(\boldsymbol{a}^{(l)}\right)\ ,$$

for $l \in [L]$, where $\mathbf{W}^{(l)}$ and $\boldsymbol{b}^{(l)}$ are the weight matrix and bias vector of the l-th layer, respectively, τ is the nonlinear activation function, $[L]$ represents the set $\{1, 2, \cdots, L\}$, and $\mathbf{W}$ denotes the set of weight matrices, i.e., $\mathbf{W} = \{\mathbf{W}^{(l)}\}_{l=1}^{L}$. For notational simplicity, $\mathbf{W}$ is referred to as either a weight matrix or a set, depending on the context. We define the output of the DNN as $\hat{\boldsymbol{y}} = \boldsymbol{a}^{(L)}$. In this paper, we specify τ as the ReLU activation. For notational clarity, we will omit the explicit layer index l in the following discussion and use primed symbols to represent variables from the $(l-1)$-th layer, e.g., $\boldsymbol{h}' = \boldsymbol{h}^{(l-1)}$.

The objective of ternary quantization is to establish a ternary-weight DNN, provided a pretrained one. In the remainder of this paper, we denote $f(\boldsymbol{x};\mathbf{W})$ and $f(\boldsymbol{x};\mathbf{T})$ as Full-Precision DNNs (FP-DNNs) and Ternary DNNs (T-DNNs), respectively. Here, $\mathbf{T} = \{\mathbf{T}^{(l)}\}_{l=1}^{L}$, each element of which belongs to $\{-1, 0, 1\}$. Notice that $\mathbf{T}$ is referred to as either a weight matrix or a set, depending on the context.

Uncertainty Quantification. Similar to Bayesian-based uncertainty quantification methods, we also treat weights as random variables and make the following assumptions.

Assumption 1. *We assume that the weights within the same layer of FP-DNNs follow a Gaussian distribution, that is,* $w^{(l)} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mu^{(l)}, (\sigma^{(l)})^2)$ *for* $l \in [L]$, *where* $w^{(l)}$ *denotes an arbitrary element of matrix* $\mathbf{W}^{(l)}$.

Assumption 2. *We assume that the weights within the same layer of T-DNNs follow a ternary-valued distribution, that is,* $t^{(l)} \overset{\text{i.i.d.}}{\sim} \mathcal{T}(p^{(l)}, 1-p^{(l)}-q^{(l)}, q^{(l)})$ *for* $l \in [L]$, *where* $t^{(l)}$ *denotes an arbitrary element of matrix* $\mathbf{T}^{(l)}$ *and* $X \sim \mathcal{T}(p, 1-p-q, q)$ *denotes that the random variable* X *follows a discrete distribution with* $P(X=-1)=p$, $P(X=0)=1-p-q$, *and* $P(X=1)=q$ *for* $p, q \in [0, 1]$.

Seminal studies (He and Fan 2019; Alawad and Ishak 2024; Zhang, Wang, and Fan 2024; Tian et al. 2025) made mild as-

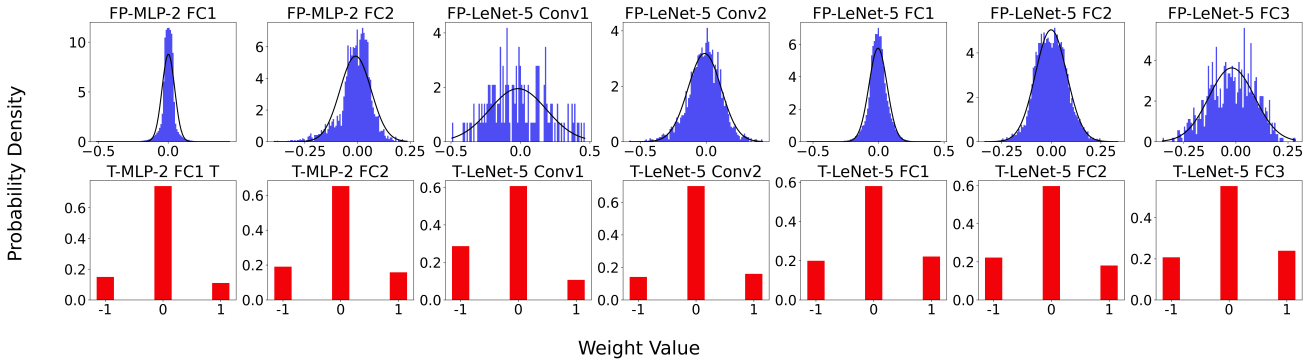

Figure 2: Empirical weight distributions of MLP-2 and LeNet-5 layer by layer.

sumptions similar to Assumption 1 for describing the weight distributions of FP-DNNs. We also conduct empirical investigations on the mildness of Assumptions 1 and 2. We first train a two-layer MLP (Rumelhart, Hinton, and Williams 1986) and a LeNet-5 (LeCun et al. 1998) with the SGD optimizer (Sutskever et al. 2013) on MNIST (LeCun et al. 1998), denoted as FP-MLP-2 and FP-LeNet-5, respectively. We then train their ternary-weight counterparts, i.e., T-MLP-2 and T-LeNet-5. Figure 2 shows the empirical weight distributions within the same layer for the concerned models, where the blue and red regions represent those of full-precision and ternary-value, respectively. It is observed that each blue region approximately follows a Gaussian distribution, hence supporting Assumption 1. In addition, it is obvious that each red region follows a ternary-valued distribution, hence supporting Assumption 2. Following this logic, we can estimate the parameters of weight distributions in Assumptions 1 and 2 from the empirical distributions.

3.2 Exact Moment Propagation

In this subsection, we introduce the EMP method. Figure 1(b) illustrates the EMP through a single layer. During feed-forward propagation, the concerned moments undergo two types of conversions. The linear transformation from $\boldsymbol{h}'$ to $\boldsymbol{a}$ induces the conversion from $(\boldsymbol{\mu}_{\boldsymbol{h}'}, \boldsymbol{\Sigma}_{\boldsymbol{h}'})$ to $(\boldsymbol{\mu}_{\boldsymbol{a}}, \boldsymbol{\Sigma}_{\boldsymbol{a}})$, which is formally expressed by the linear computation of EMP. Similarly, the nonlinear activation from $\boldsymbol{a}$ to $\boldsymbol{h}$ induces the conversion from $(\boldsymbol{\mu}_{\boldsymbol{a}}, \boldsymbol{\Sigma}_{\boldsymbol{a}})$ to $(\boldsymbol{\mu}_{\boldsymbol{h}}, \boldsymbol{\Sigma}_{\boldsymbol{h}})$, which is formally expressed by the nonlinear computation of EMP. The nonlinear computation of EMP calculates the covariance matrix $\boldsymbol{\Sigma}_{\boldsymbol{h}}$ element by element, where each element is obtained by solving an integral. The challenge lies in solving the integrals for the non-diagonal elements, as the function being integrated becomes more complex than that in the diagonal case (Frey and Hinton 1999). Previous studies have approximated these integrals by either constraining the integration intervals (Bibi, Alfadly, and Ghanem 2018) or using asymptotic methods (Wu et al. 2019). In contrast, we solve these integrals directly by applying a recursive approach based on integration by parts, leading to an exact solution for $\boldsymbol{\Sigma}_{\boldsymbol{h}}$. The following subsections elaborate on the completion of linear computation and nonlinear computation of the EMP.

Linear Computation of EMP. The goal of the linear computation of EMP is to provide exact expressions to calculate $(\boldsymbol{\mu}_{\boldsymbol{a}}, \boldsymbol{\Sigma}_{\boldsymbol{a}})$ using $(\boldsymbol{\mu}_{\boldsymbol{h}'}, \boldsymbol{\Sigma}_{\boldsymbol{h}'})$ and $(\boldsymbol{\mu}_{\mathbf{W}}, \boldsymbol{\Sigma}_{\mathbf{W}})$. We achieve this goal with the following theorem.

Theorem 3. *Let $\boldsymbol{h}'$ be a random vector characterized by moments $\boldsymbol{\mu}_{\boldsymbol{h}'}$ and $\boldsymbol{\Sigma}_{\boldsymbol{h}'}$, b_i denote a constant, and $\mathbf{W}_i$ be a random vector characterized by moments $\boldsymbol{\mu}_{\mathbf{W}_i}$ and $\boldsymbol{\Sigma}_{\mathbf{W}_i}$ for $i \in \{1, 2\}$. Define the scalar random variable $a_i = \mathbf{W}_i^\top \boldsymbol{h}' + b_i$, where $i \in \{1, 2\}$. Then, the following holds.*

$$\mathbb{E}[a_i] = \boldsymbol{\mu}_{\mathbf{W}_i}^\top \boldsymbol{\mu}_{\boldsymbol{h}'} + b_i \,,$$

$$\mathbb{V}(a_i) = \boldsymbol{\mu}_{\boldsymbol{h}'}^\top \boldsymbol{\Sigma}_{\mathbf{W}_i} \boldsymbol{\mu}_{\boldsymbol{h}'} + \operatorname{tr}(\boldsymbol{\Sigma}_{\mathbf{W}_i} \boldsymbol{\Sigma}_{\boldsymbol{h}'}) + \boldsymbol{\mu}_{\mathbf{W}_i}^\top \boldsymbol{\Sigma}_{\boldsymbol{h}'} \boldsymbol{\mu}_{\mathbf{W}_i} \,,$$

$$\operatorname{Cov}(a_1, a_2) = \operatorname{tr}(\operatorname{Cov}(\mathbf{W}_1, \mathbf{W}_2) \boldsymbol{\Sigma}_{\boldsymbol{h}'}) + \boldsymbol{\mu}_{\mathbf{W}_1}^\top \boldsymbol{\Sigma}_{\boldsymbol{h}'} \boldsymbol{\mu}_{\mathbf{W}_2} + \boldsymbol{\mu}_{\boldsymbol{h}'}^\top \operatorname{Cov}(\mathbf{W}_1, \mathbf{W}_2) \boldsymbol{\mu}_{\boldsymbol{h}'} \,.$$

Theorem 3 provides analytic expressions of the expectation and covariance of pre-activation variables. Notice that Theorem 3 is compatible with FP-DNNs and T-DNNs, leading to the following corollaries.

Corollary 4. *Under the conditions given in Theorem 3 and Assumption 1, the following holds.*

$$\mathbb{E}[a_i] = \mu \mathbf{1}^\top \boldsymbol{\mu}_{\boldsymbol{h}'} + b_i \,,$$

$$\mathbb{V}(a_i) = \sigma^2 \left(\|\boldsymbol{\mu}_{\boldsymbol{h}'}\|_2^2 + \operatorname{tr}(\boldsymbol{\Sigma}_{\boldsymbol{h}'}) \right) + \mu^2 \mathbf{1}^\top \boldsymbol{\Sigma}_{\boldsymbol{h}'} \mathbf{1} \,,$$

$$\operatorname{Cov}(a_1, a_2) = \mu^2 \mathbf{1}^\top \boldsymbol{\Sigma}_{\boldsymbol{h}'} \mathbf{1} \,,$$

where $\mathbf{1}$ denotes the vector with all components equal to 1 and $\|\cdot\|_2$ denotes the L_2 norm.

Corollary 5. *Under the conditions in Theorem 3 and Assumption 2, the following holds.*

$$\mathbb{E}[a_i] = (q - p) \mathbf{1}^\top \boldsymbol{\mu}_{\boldsymbol{h}'} + b_i \,,$$

$$\mathbb{V}(a_i) = \left[(q + p) - (q - p)^2 \right] \left(\|\boldsymbol{\mu}_{\boldsymbol{h}'}\|_2^2 + \operatorname{tr}(\boldsymbol{\Sigma}_{\boldsymbol{h}'}) \right) + (q - p)^2 \mathbf{1}^\top \boldsymbol{\Sigma}_{\boldsymbol{h}'} \mathbf{1} \,,$$

$$\operatorname{Cov}(a_1, a_2) = (q - p)^2 \mathbf{1}^\top \boldsymbol{\Sigma}_{\boldsymbol{h}'} \mathbf{1} \,.$$

Corollary 4 and Corollary 5 provide analytic expressions of the expectation and covariance of the pre-activation variable, avoiding the computation of the term $\operatorname{Cov}(\mathbf{W}_1, \mathbf{W}_2)$

in Theorem 3, thereby enabling a more efficient linear computation of EMP than that of Theorem 3. The full proof of Theorem 3 can be found in Appendix C.

Nonlinear Computation of EMP. The goal of the nonlinear computation of EMP is to provide exact expressions to calculate $(\boldsymbol{\mu}_{\boldsymbol{h}}, \boldsymbol{\Sigma}_{\boldsymbol{h}})$ using $(\boldsymbol{\mu}_{\boldsymbol{a}}, \boldsymbol{\Sigma}_{\boldsymbol{a}})$. The core is to characterize the pre-activation variable $\boldsymbol{a}$ according to the Law of the Unconscious Statistician (Feller 1991). In general, we make the following assumptions.

Assumption 6. *Under Assumption 1, we assume that the pre-activation variables of FP-DNNs follow a Gaussian distribution, that is,* $\boldsymbol{a}^{(l)} \sim \mathcal{N}(\boldsymbol{\mu}_{\boldsymbol{a}^{(l)}}, \boldsymbol{\Sigma}_{\boldsymbol{a}^{(l)}})$ *for* $l \in [L]$.

Assumption 7. *Under Assumption 2, we assume that the pre-activation variables of T-DNNs follow a Gaussian distribution, that is,* $\boldsymbol{a}^{(l)} \sim \mathcal{N}(\boldsymbol{\mu}_{\boldsymbol{a}^{(l)}}, \boldsymbol{\Sigma}_{\boldsymbol{a}^{(l)}})$ *for* $l \in [L]$.

We validate the mildness of Assumptions 6 and 7, and discuss the role and interdependencies of all assumptions made throughout this paper in Appendix A. Next, it suffices to analytically express $(\boldsymbol{\mu}_{\boldsymbol{h}}, \boldsymbol{\Sigma}_{\boldsymbol{h}})$ under the Gaussian assumption on $\boldsymbol{a}$. We separately denote the Probability Density Function (PDF) and the Cumulative Distribution Function (CDF) of the standard univariate Gaussian distribution as

$$\phi(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right) \ , \ \Phi(x) = \int_{-\infty}^{x} \phi(u) \, \mathrm{d}u \ .$$

In addition, we separately denote the PDF and the CDF of the standard bivariate Gaussian distribution with correlation coefficient $\rho \in (-1, 1)$ as

$$\begin{cases} \phi_2(x_1, x_2; \rho) = \dfrac{1}{2\pi\sqrt{1-\rho^2}} \exp\left(\dfrac{x_1^2 - 2\rho x_1 x_2 + x_2^2}{-2(1-\rho^2)}\right) \ , \\ \Phi_2(x_1, x_2; \rho) = \displaystyle\int_{-\infty}^{x_1} \int_{-\infty}^{x_2} \phi_2(u_1, u_2; \rho) \ \mathrm{d}u_1 \, \mathrm{d}u_2 \ . \end{cases}$$

With the notations above, we present the following theorem.

Theorem 8. *Given* $(a_1, a_2)^\top \sim \mathcal{N}_2(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ *with*

$$\boldsymbol{\mu} = (\mu_1, \mu_2)^\top \ , \ \boldsymbol{\Sigma} = \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix} \ ,$$

and $h_i = \tau(a_i)$ *for* $i \in \{1, 2\}$ *where* $\tau(\cdot)$ *denotes the ReLU activation, then the following holds.*

$$\begin{aligned} &\mathbb{E}[h_i] = \sigma_i \left(r_i \Phi(r_i) + \phi(r_i)\right) \ , \\ &\mathbb{V}(h_i) = \sigma_i^2 \left[\left(1 + r_i^2\right) \Phi(r_i) + r_i \phi(r_i)\right] - \mathbb{E}[h_i]^2 \ , \\ &\mathrm{Cov}(h_1, h_2) = \sigma_1\sigma_2 \big[r_2 \phi(r_1) \, \Phi(c_2) + r_1 \phi(r_2) \, \Phi(c_1) \\ &\quad + (r_1 r_2 + \rho) \, \Phi_2(r_1, r_2; \rho) + \sqrt{1-\rho^2} \phi(r_2) \, \phi(c_1) \big] \\ &\quad - \mathbb{E}[Z_1] \, \mathbb{E}[Z_2] \ , \end{aligned}$$

in which $r_i = \mu_i / \sigma_i$, $c_i = (r_i - \rho r_{3-i}) / \sqrt{1-\rho^2}$ *for* $i \in \{1, 2\}$, $\phi(\cdot)$ *and* $\Phi(\cdot)$ *denote the PDF and the CDF of the standard univariate normal distribution, respectively, and* $\Phi_2(\cdot, \cdot; \rho)$ *denotes the CDF of standard bivariate normal distribution with correlation coefficient* ρ.

Theorem 8 provides analytic expressions of expectation and covariance of $\tau(\boldsymbol{a})$, completing the nonlinear computation of EMP. Notice that Theorem 8 is straightforwardly compatible with both FP-DNNs and T-DNNs. To the best of our knowledge, Theorem 8 provides the first exact solution to the covariance $\boldsymbol{\Sigma}_{\boldsymbol{h}}$ of a Gaussian vector passed through the ReLU activation under mild assumptions; in contrast, the previous methods provided either exact $\boldsymbol{\Sigma}_{\boldsymbol{h}}$ under strong assumptions (Bibi, Alfadly, and Ghanem 2018) or approximated $\boldsymbol{\Sigma}_{\boldsymbol{h}}$ with $\boldsymbol{h} = \tau(\boldsymbol{a})$ (Wu et al. 2019). The full proof of Theorem 8 can be found in Appendix D.

To summarize the computations above, we can complete the EMP method for uncertainty quantification. Algorithm 1 lists the procedure of the EMP method for quantifying the uncertainty of T-DNNs, where Step 3 and Step 4 represent the linear computation and the nonlinear computation of the EMP, respectively. Algorithm 1 also adapts to FP-DNNs with two modifications: replacing the input $\{p^{(l)}, \ q^{(l)}\}_{l=1}^{L}$ with $\{\mu^{(l)}, \ \sigma^{(l)}\}_{l=1}^{L}$ and substituting Corollary 5 in Step 3 with Corollary 4.

Algorithm 1: EMP for T-DNNs

Input: Sample distribution $\boldsymbol{\mu}_{\boldsymbol{x}}, \boldsymbol{\Sigma}_{\boldsymbol{x}}$, weight distribution parameters $\{p^{(l)}, \ q^{(l)}\}_{l=1}^{L}$

Output: Output covariance $\boldsymbol{\Sigma}_{\hat{\boldsymbol{y}}}$

Procedure:

1: Set $\boldsymbol{\mu}_{\boldsymbol{h}}^{(0)} = \boldsymbol{\mu}_{\boldsymbol{x}}, \boldsymbol{\Sigma}_{\boldsymbol{h}}^{(0)} = \boldsymbol{\Sigma}_{\boldsymbol{x}}$
2: **for** $1 \leq l \leq L$ **do**
3: $\quad (\boldsymbol{\mu}_{\boldsymbol{a}}^{(l)}, \boldsymbol{\Sigma}_{\boldsymbol{a}}^{(l)}) \Leftarrow (\boldsymbol{\mu}_{\boldsymbol{h}}^{(l-1)}, \boldsymbol{\Sigma}_{\boldsymbol{h}}^{(l-1)})$ and $(p^{(l)}, \ q^{(l)})$ by Corollary 5
4: $\quad (\boldsymbol{\mu}_{\boldsymbol{h}}^{(l)}, \boldsymbol{\Sigma}_{\boldsymbol{h}}^{(l)}) \Leftarrow (\boldsymbol{\mu}_{\boldsymbol{a}}^{(l)}, \boldsymbol{\Sigma}_{\boldsymbol{a}}^{(l)})$ by Theorem 8
5: **end for**
6: $\boldsymbol{\Sigma}_{\hat{\boldsymbol{y}}} = \boldsymbol{\Sigma}_{\boldsymbol{a}}^{(L)}$

4 Uncertainty Optimization

In this section, we propose an optimization method termed MOMent Alignment (MOMA) for improving the accuracy and reducing the uncertainty of T-DNNs. The MOMA is compatible with various quantization training algorithms for DNNs. Here, we also adopt ternary quantization as an example and draw the key ideas from Li et al. (2016) and Alemdar et al. (2017) to achieve ternary training. We detail the optimization procedure in Subsection 4.1, and then formally propose the corresponding MOMA in Subsection 4.2.

4.1 Ternary Training of T-DNNs

Following the supervised learning paradigm, the optimization problem of training a T-DNN f can be formulated as

$$\min_{\mathbf{T}} \mathcal{L}(\boldsymbol{y}, f(\boldsymbol{x}; \mathbf{T})) \ , \tag{1}$$

where $(\boldsymbol{x}, \boldsymbol{y})$ denotes the pair of input and target variables, and $\mathcal{L}(\cdot, \cdot)$ denotes the loss function. Let $s \in \mathbb{N}$ indicate the training iteration. The weight-updating procedure is listed as

$$\begin{cases} \mathbf{T}_{s+1} = \mathrm{sgn}_{\gamma_s}(\mathbf{W}_{s+1}) \ , \\ \mathbf{W}_{s+1} = \mathbf{W}_s + \eta \dfrac{\partial \mathcal{L}(\boldsymbol{y}, f(\boldsymbol{x}; \mathbf{T}_s))}{\partial \mathbf{W}_s} \ , \end{cases} \tag{2}$$

where η denotes the learning rate and $\text{sgn}_{\gamma_s}(\cdot)$ is the partition function. Here, we employ the partition function used in Li et al. (2016), that is,

$$t = \text{sgn}_{\gamma_s}(w) = \begin{cases} 1\,, & w > \gamma_s\,, \\ 0\,, & -\gamma_s \le w \le \gamma_s\,, \\ -1\,, & w < -\gamma_s\,, \end{cases} \tag{3}$$

where t and w denote an arbitrary element of the weight matrices $\mathbf{T}_s$ and $\mathbf{W}_s$, respectively, and $\gamma_s \in \mathbb{R}^+$ is a partition factor defined as

$$\gamma_s = \frac{1}{N}\left\|\mathbf{W}_s\right\|_1\,. \tag{4}$$

The gradient-related term $\partial\mathcal{L}(\boldsymbol{y}, f(\boldsymbol{x};\mathbf{T}_s))/\partial\mathbf{W}_s$ in Eq. (2) is computed via surrogate gradients (Bengio, Léonard, and Courville 2013) as

$$\begin{cases} \dfrac{\partial\mathcal{L}\left(\boldsymbol{y}, f\left(\boldsymbol{x};\mathbf{T}_s\right)\right)}{\partial\mathbf{W}_s} = \dfrac{\partial\mathcal{L}\left(\boldsymbol{y}, f\left(\boldsymbol{x};\mathbf{T}_s\right)\right)}{\partial\mathbf{T}_s}\dfrac{\partial\mathbf{T}_s}{\partial\mathbf{W}_s}\,, \\ \dfrac{\partial\mathbf{T}_s}{\partial\mathbf{W}_s} = \dfrac{\partial\text{sgn}_{\gamma}\left(\mathbf{W}_s\right)}{\partial\mathbf{W}_s} = 1\,. \end{cases} \tag{5}$$

In summary, one can obtain a T-DNN by initializing $\mathbf{W}_0$ with the pretrained full-precision weights, and following the optimization procedure listed in Eq. (2), Eq. (4), and Eq. (5).

4.2 Moment Alignment

This subsection proposes the MOMA optimization for improving the accuracy and reducing the uncertainty of T-DNNs, ternarized from a pretrained FP-DNN. The key idea of MOMA is to formulate an optimization problem by aligning the first and second central moments of T-DNN weights to zero. The following theorem supports our line of thought.

Theorem 9. *Under Assumption 2 and Assumption 7, if* $q - p = 0$ *and* $p + q - (q-p)^2 = 0$ *hold for the ternary-valued distribution parameters* $(p,\ q)$ *of each layer, then the output covariance satisfies* $\boldsymbol{\Sigma}_{\hat{\boldsymbol{y}}} = \mathbf{0}$.

Theorem 9 demonstrates that there exists an optimal solution, that is, the uncertainty of T-DNNs can be reduced to zero if one aligns the concerned moments to zero, which provides a theoretical guarantee for MOMA. The full proof of Theorem 9 can be found in Appendix E. Inspired by Theorem 9, we formulate the MOMA optimization as

$$\begin{cases} \tilde{\gamma}_s = \arg\min\limits_{\gamma}\ \left\{\left|g_3\left(\gamma, \mathbf{W}_s\right)\right| + \lambda\left|g_4\left(\gamma, \mathbf{W}_s\right)\right|\right\}\,, \\ \text{s.t.}\quad \gamma \in \left((1-c)\,\gamma_s, (1+c)\,\gamma_s\right)\,. \end{cases} \tag{6}$$

Algorithm 2: Ternary Training with MOMA for T-DNNs

Input: Pretrained full-precision weights $\mathbf{W}_0$, the maximum number of training epochs S_0
Output: Ternary weight $\mathbf{T}$
Procedure:

1: **for** $0 \le s \le S_0$ **do**
2: $\quad\mathbf{W}_{s+1} = \mathbf{W}_s + \eta\,\partial\mathcal{L}(\boldsymbol{y}, f(\boldsymbol{x};\mathbf{T}_s))/\partial\mathbf{T}_s$ in Eq. (2)
3: $\quad\gamma_s = \frac{1}{N}\|\mathbf{W}_s\|_1$ in Eq. (4)
4: $\quad\tilde{\gamma}_s \leftarrow$ solving Eq. (6)
5: $\quad\mathbf{T}_{s+1} = \text{sgn}_{\tilde{\gamma}_s}(\mathbf{W}_{s+1})$ in Eq. (3).
6: **end for**

where $c \in (0,1)$, $\lambda \in \mathbb{R}^+$, g_3 is an estimator for $q - p$, and g_4 is an estimator for $p + q - (q-p)^2$. The search interval is designed to be a neighborhood of the conventional partition factor of Eq. (4) to maintain accuracy, while the objective function is designed to align the concerned moments to zero for uncertainty reduction. The estimators in Eq. (6) are defined as follows.

$$\begin{cases} g_1\left(\gamma, \mathbf{W}_s\right) = \sum_{i,j}\chi_{(-\infty,-\gamma)}\left(\mathbf{W}_{ij}\right)/N\,, \\ g_2\left(\gamma, \mathbf{W}_s\right) = \sum_{i,j}\chi_{(\gamma,+\infty)}\left(\mathbf{W}_{ij}\right)/N\,, \\ g_3\left(\gamma, \mathbf{W}_s\right) = g_2\left(\gamma, \mathbf{W}_s\right) - g_1\left(\gamma, \mathbf{W}_s\right)\,, \\ g_4\left(r, \mathbf{W}_s\right) = g_1\left(\gamma, \mathbf{W}_s\right) + g_2\left(\gamma, \mathbf{W}_s\right) - \left(g_3\left(\gamma, \mathbf{W}_s\right)\right)^2\,, \end{cases}$$

where g_1 is an estimator for p, g_2 is an estimator for q, $\mathbf{W}_{ij}$ denotes the (i,j)-th element of $\mathbf{W}$, and $\chi_A(\cdot)$ denotes the indicator function of a concerned interval A that satisfies $\chi_A(x) = 1$ if $x \in A$ and $\chi_A(x) = 0$ if $x \notin A$.

Algorithm 2 details the calculation procedure of the MOMA optimization for T-DNN training. In contrast to conventional T-DNN training algorithms, MOMA involves an additional search for the partition factor, as detailed in Step 5. Hence, our proposed MOMA method can be incorporated into other quantization training algorithms that use a threshold-based partition function like Eq. (3).

5 Experiments

In this section, we conduct experiments to verify the effectiveness of the proposed EMP and MOMA methods. The experiments are performed to answer the following questions.

Q1. Is the proposed EMP method comparable to the classical uncertainty estimation methods?

Q2. To what extent, if at all, does weight ternarization increase the uncertainty of DNNs?

Q3. Whether and to what extent does the proposed MOMA method reduce the uncertainty of T-DNNs?

All experiments were conducted on MNIST (LeCun et al. 1998), CIFAR-10, and CIFAR-100 (Krizhevsky 2009) datasets. The concerned architectures include MLP 2[1] (Rumelhart, Hinton, and Williams 1986), LeNet-5 (LeCun et al. 1998), VGG-13 (Simonyan 2014), and ResNet-18 (He et al. 2016). We denote the full-precision DNNs as FP-DNNs, such as FP-MLP-2. The ternary-weight DNNs trained using the procedures outlined in Subsection 4.1 are referred to as T-DNNs, such as T-MLP-2. The ternary-weight DNNs trained with Algorithm 2 are denoted as T-DNNs-MOMA, such as T-MLP-2-MOMA. The proposed methods and contenders are implemented in `PyTorch` and `SciPy`. All experiments were conducted on a `13th gen Intel(R) Core(TM) i9-13900KF` CPU and an `NVIDIA RTX 4090` GPU. Experiments were repeated 5 times with different initialization seeds, reporting the mean α and standard deviation β as α_β.

[1]MLP-2 denotes a two-layer fully-connected architecture.

Datasets	o	FP-LeNet-5			T-LeNet-5		
		EMP	MP	PL-DNN	EMP	MP	PL-DNN
MNIST	0	1.0007 ± 0.0101	1.1809 ± 0.0217	1.2772 ± 0.0153	1.0023 ± 0.0109	–	–
	1	0.9999 ± 0.0103	1.1793 ± 0.0238	1.2759 ± 0.0139	1.0023 ± 0.0121	–	–
	2	0.9990 ± 0.0118	1.1794 ± 0.0239	1.2748 ± 0.0170	1.0021 ± 0.0117	–	–
	3	1.0008 ± 0.0104	1.1808 ± 0.0215	1.2770 ± 0.0153	1.0022 ± 0.0108	–	–
	4	1.0011 ± 0.0096	1.1807 ± 0.0239	1.2775 ± 0.0136	1.0023 ± 0.0107	–	–
	5	1.0010 ± 0.0106	1.1815 ± 0.0237	1.2773 ± 0.0149	1.0022 ± 0.0109	–	–
	6	1.0011 ± 0.0096	1.1818 ± 0.0246	1.2774 ± 0.0130	1.0023 ± 0.0105	–	–
	7	0.9996 ± 0.0099	1.1803 ± 0.0243	1.2755 ± 0.0143	1.0021 ± 0.0101	–	–
	8	1.0014 ± 0.0098	1.1814 ± 0.0218	1.2422 ± 0.0140	1.0021 ± 0.0107	–	–
	9	0.9999 ± 0.0100	1.1801 ± 0.0225	1.2759 ± 0.0142	1.0019 ± 0.0106	–	–
CIFAR-10	0	1.0098 ± 0.0099	1.5029 ± 0.0177	1.6167 ± 0.0334	1.0141 ± 0.0113	–	–
	1	1.0092 ± 0.0105	1.5020 ± 0.0207	1.6156 ± 0.0299	1.0136 ± 0.0115	–	–
	2	1.0090 ± 0.0127	1.5017 ± 0.0218	1.6154 ± 0.0358	1.0141 ± 0.0110	–	–
	3	1.0100 ± 0.0110	1.5031 ± 0.0189	1.6169 ± 0.0338	1.0141 ± 0.0115	–	–
	4	1.0106 ± 0.0098	1.5041 ± 0.0179	1.6179 ± 0.0307	1.0137 ± 0.0116	–	–
	5	1.0099 ± 0.0109	1.5030 ± 0.0200	1.6167 ± 0.0311	1.0139 ± 0.0118	–	–
	6	1.0103 ± 0.0094	1.5036 ± 0.0185	1.6174 ± 0.0282	1.0139 ± 0.0109	–	–
	7	1.0090 ± 0.0099	1.5017 ± 0.0184	1.6153 ± 0.0307	1.0139 ± 0.0120	–	–
	8	1.0103 ± 0.0093	1.5036 ± 0.0179	1.6173 ± 0.0312	1.0140 ± 0.0118	–	–
	9	1.0095 ± 0.0100	1.5025 ± 0.0179	1.6162 ± 0.0314	1.0141 ± 0.0109	–	–

Table 1: Performance of the EMP and contenders for LeNet-5 on MNIST and CIFAR-10. Results for other architectures and datasets are detailed in Appendix F.

5.1 Verifications on EMP

This subsection demonstrates the effectiveness of the proposed EMP method. Let $\boldsymbol{x} \in D$ denote a test sample from the test set D. In Algorithm 1, we set $\boldsymbol{\mu}_{\boldsymbol{x}} = \boldsymbol{x}$ and $\boldsymbol{\Sigma}_{\boldsymbol{x}} = \mathbf{0}$.

To compare the performance of various uncertainty estimators, we follow the evaluation metric in Bibi, Alfadly, and Ghanem (2018), where they treat the classical Monte Carlo (MC) estimator (Jospin et al. 2022) as the ground truth, with other estimators being evaluated based on their closeness to it. Formally, we define the variance ratio as $r_o(\boldsymbol{x}) = \mathbb{V}_o^{\mathrm{MC}}(\boldsymbol{x}) \;/\; \mathbb{V}_o^{\mathrm{E}}(\boldsymbol{x})$, where $\mathbb{V}_o^{\mathrm{MC}}(\boldsymbol{x})$ and $\mathbb{V}_o^{\mathrm{E}}(\boldsymbol{x})$ denote the o-th element of the output variance vector obtained by the MC method and by an estimator, respectively, for $o \in \{0, 1, \cdots, 9\}$. A ratio $r_o(\boldsymbol{x})$ closer to 1 indicates higher precision of an estimator on $\boldsymbol{x}$. We then compute the mean and standard deviation of $r_o(\boldsymbol{x})$ over D as

$$\mathbb{E}\left[r_o\left(\boldsymbol{x}\right)\right] = \sum\nolimits_{\boldsymbol{x}\in D} r_o\left(\boldsymbol{x}\right) / \left|D\right| \;,$$
$$\sqrt{\mathbb{V}\left[r_o\left(\boldsymbol{x}\right)\right]} = \sqrt{\sum\nolimits_{\boldsymbol{x}\in D} \left(r_o\left(\boldsymbol{x}\right) - \mathbb{E}\left[r_o\left(\boldsymbol{x}\right)\right]\right)^2 / \left|D\right|} \;.$$

Then we use $\mathbb{E}[r_o(\boldsymbol{x})] \pm \sqrt{\mathbb{V}[r_o(\boldsymbol{x})]}$ as a metric, which is the same as that in Bibi, Alfadly, and Ghanem (2018), to measure the uncertainty estimation performance of an estimator. A mean $\mathbb{E}[r_o(\boldsymbol{x})]$ closer to 1 indicates higher precision on average, while a standard deviation $\sqrt{\mathbb{V}[r_o(\boldsymbol{x})]}$ closer to 0 indicates greater robustness on average.

Table 1 shows the uncertainty estimation performance of EMP and its contenders, including PL-DNN (Bibi, Alfadly, and Ghanem 2018) and MP (Wu et al. 2019) for LeNet-5 on MNIST and CIFAR-10. It is observed that the EMP method achieves near-ideal performance and significantly outperforms its contenders. Experimental results for other architectures and datasets, detailed in Appendix F, further support this observation. These findings demonstrate that the proposed EMP estimator is comparable to the MC estimator and more precise than its contenders, thus answering Q1.

5.2 Verifications on MOMA

This experiment demonstrates the effectiveness of the MOMA method. We first analyze the uncertainty difference between DNNs with and without weight ternarization to take a closer look at Q2. We begin this analysis by defining two variance ratios as

$$\tilde{r}_o\left(\boldsymbol{x}\right) = \mathbb{V}_o^{\mathrm{E}}(\boldsymbol{x}; \text{T-DNN}) \;/\; \mathbb{V}_o^{\mathrm{E}}(\boldsymbol{x}; \text{FP-DNN}) \;,$$
$$\tilde{r}'_o\left(\boldsymbol{x}\right) = \mathbb{V}_o^{\mathrm{E}}(\boldsymbol{x}; \text{T-DNN-MOMA}) \;/\; \mathbb{V}_o^{\mathrm{E}}(\boldsymbol{x}; \text{FP-DNN}) \;,$$

where $\mathbb{V}_o^{\mathrm{E}}(\boldsymbol{x}; \text{Model})$ represents the variance, i.e., $\mathbb{V}_o^{\mathrm{E}}(\boldsymbol{x})$, computed using the corresponding model. A ratio greater than 1 indicates that weight ternarization increases DNN uncertainty by the corresponding ratio.

Figure 3 shows the concerned uncertainty difference on MLP-2 by providing the histogram of the two ratios. Due to space limitations, we only plot the results for $o \in \{0, 1, \cdots, 5\}$ and the ratio range of $[0, 3000]$, containing 99% of the empirical points for the variance ratios. Full results are given in Appendix G. The blue and red regions represent the histograms of $\tilde{r}_o(\boldsymbol{x})$ and $\tilde{r}'_o(\boldsymbol{x})$, respectively. It is observed that $\tilde{r}_o(\boldsymbol{x}) > 1.4$ and $\tilde{r}'_o(\boldsymbol{x}) > 1.4$ for all

Models	Accuracy	Uncertainty	Reduction
FP-MLP-2	$0.9761_{0.0002}$	$1.4697_{0.0082} \times 10^{3}$	$68.08_{0.1389}\%$
T-MLP-2	$0.9752_{0.0005}$	$1.2572_{0.0236} \times 10^{5}$	
T-MLP-2-MOMA	$0.9751_{0.0006}$	$4.0129_{0.0039} \times 10^{4}$	
FP-LeNet-5	$0.9911_{0.0001}$	$1.7728_{0.0232} \times 10^{3}$	$98.55_{0.0025}\%$
T-LeNet-5	$0.9903_{0.0003}$	$2.8113_{0.0490} \times 10^{5}$	
T-LeNet-5-MOMA	$0.9885_{0.0006}$	$4.0906_{0.1433} \times 10^{3}$	
FP-VGG-13	$0.9985_{0.0001}$	$4.6029_{0.2302} \times 10^{-1}$	$20.86_{2.5366}\%$
T-VGG-13	$0.9954_{0.0005}$	$1.4500_{0.0210} \times 10^{1}$	
T-VGG-13-MOMA	$0.9957_{0.0005}$	$1.1463_{0.0198} \times 10^{1}$	
FP-ResNet-18	$0.9963_{0.0001}$	$6.9757_{0.3743} \times 10^{-3}$	$44.27_{4.4212}\%$
T-ResNet-18	$0.9938_{0.0002}$	$4.2808_{0.1772} \times 10^{-1}$	
T-ResNet-18-MOMA	$0.9951_{0.0002}$	$2.4067_{0.0501} \times 10^{-1}$	

Table 2: Accuracy and uncertainty of concerned models on MNIST. Results for other datasets are detailed in Appendix G.

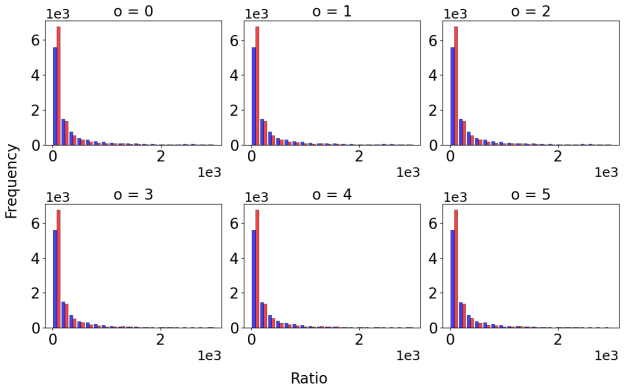

Figure 3: Histograms of the concerned variance ratios.

$\boldsymbol{x} \in D$, indicating that DNN uncertainty increases after weight ternarization, regardless of whether MOMA is used. This observation answers Q2. Furthermore, the red bars are taller at lower x-values and shorter at higher x-values, indicating that the proposed MOMA method effectively reduces the concerned variance ratio.

To further measure the uncertainty reduction performance of the proposed MOMA method, we define the uncertainty of a model on a dataset D as

$$\text{Uncertainty} = \frac{1}{|D|} \sum_{\boldsymbol{x} \in D} \text{tr}\left(\boldsymbol{\Sigma}_{\hat{\boldsymbol{y}}(\boldsymbol{x})}\right).$$

Table 2 shows the uncertainty reduction performance of MOMA across various DNN architectures on MNIST by comparing the uncertainty of the T-DNNs and T-DNNs-MOMA. It is observed that the MOMA greatly reduces the uncertainty of T-DNNs on MNIST. Experimental results for CIFAR-10 and CIFAR-100, provided in Appendix G, further support this observation. This finding demonstrates that the proposed MOMA can effectively reduce the uncertainty of T-DNNs, thus answering Q3.

6 Conclusions

In this paper, we proposed the EMP, a theoretical uncertainty quantification method that works independently of weight precision. Experimental results showed that T-DNNs exhibit significantly greater uncertainty than FP-DNNs across varying ratios. Inspired by the EMP estimator, we proposed MOMA, an optimization method for reducing the uncertainty of T-DNNs. Empirical results validated the effectiveness of both EMP and MOMA.

Acknowledgements

This research was supported by the Jiangsu Provincial Natural Science Foundation Youth Project (BK20230782) and Nanjing University-China Mobile Communications Group Co., Ltd. Joint Institute.

References

Alawad, M.; and Ishak, M. 2024. Probabilistic Bayesian Neural Networks for Efficient Inference. In *Proceedings of the 34th Great Lakes Symposium on VLSI*, 724–729.

Alemdar, H.; Leroy, V.; Prost-Boucle, A.; and Pétrot, F. 2017. Ternary neural networks for resource-efficient AI applications. In *Proceedings of the 30th International Joint Conference on Neural Networks*, 2547–2554.

Bengio, Y.; Léonard, N.; and Courville, A. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*.

Bibi, A.; Alfadly, M.; and Ghanem, B. 2018. Analytic expressions for probabilistic moments of PL-DNN with Gaussian input. In *Proceedings of the 31st IEEE Conference on Computer Vision and Pattern Recognition*, 9099–9107.

Courbariaux, M.; Bengio, Y.; and David, J.-P. 2015. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in Neural Information Processing Systems 28*, 3123–3131.

Einbinder, B.-S.; Romano, Y.; Sesia, M.; and Zhou, Y. 2022. Training uncertainty-aware classifiers with conformalized

deep learning. In *Advances in Neural Information Processing Systems 35*, 22380–22395.

Feller, W. 1991. *An introduction to probability theory and its applications*. John Wiley & Sons.

Frey, B.; and Hinton, G. 1999. Variational learning in nonlinear Gaussian belief networks. *Neural Computation*, 11(1): 193–213.

Gope, D.; Dasika, G.; and Mattina, M. 2019. Ternary hybrid neural-tree networks for highly constrained iot applications. *Proceedings of Machine Learning and Systems*, 1: 190–200.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the 29th IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.

He, Z.; and Fan, D. 2019. Simultaneously optimizing weight and quantizer of ternary neural network using truncated gaussian approximation. In *Proceedings of the 32nd IEEE Conference on Computer Vision and Pattern Recognition*, 11438–11446.

Hubara, I.; Courbariaux, M.; Soudry, D.; El-Yaniv, R.; and Bengio, Y. 2016. Binarized neural networks. *Advances in Neural Information Processing Systems 29*, 4107–4115.

Jospin, L.-V.; Laga, H.; Boussaid, F.; Buntine, W.; and Bennamoun, M. 2022. Hands-on Bayesian neural networks—A tutorial for deep learning users. *IEEE Computational Intelligence Magazine*, 17(2): 29–48.

Krizhevsky, A. 2009. Learning multiple layers of features from tiny images. *Master's thesis*.

Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Proceedings of the 31st Advances in Neural Information Processing Systems 30*, 6405–6416.

LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.

Li, F.; Liu, B.; Wang, X.; Zhang, B.; and Yan, J. 2016. Ternary weight networks. *arXiv preprint arXiv:1605.04711*.

Li, Y.; Xu, S.; Zhang, B.; Cao, X.; Gao, P.; and Guo, G. 2022. Q-vit: Accurate and fully quantized low-bit vision transformer. In *Advances in neural information processing systems 35*, 34451–34463.

Lin, Z.; Courbariaux, M.; Memisevic, R.; and Bengio, Y. 2016. Neural networks with few multiplications. In *Proceedings of the 4th International Conference on Learning Representations*.

Meng, X.; Bachmann, R.; and Khan, M.-E. 2020. Training binary neural networks using the Bayesian learning rule. In *Proceedings of the 37th International Conference on Machine Learning*, 6852–6861.

Peters, J.; and Welling, M. 2018. Probabilistic binary neural networks. *arXiv preprint arXiv:1809.03368*.

Rastegari, M.; Ordonez, V.; Redmon, J.; and Farhadi, A. 2016. Xnor-net: Imagenet classification using binary convolutional neural networks. In *Proceedings of the 14th European Conference on Computer Vision*, 525–542.

Rumelhart, D.; Hinton, G.; and Williams, R. 1986. Learning representations by back-propagating errors. *Nature*, 323(6088): 533–536.

Shen, X.; Dong, P.; Lu, L.; Kong, Z.; Li, Z.; Lin, M.; Wu, C.; and Wang, Y. 2024. Agile-Quant: Activation-Guided Quantization for Faster Inference of LLMs on the Edge. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, 18944–18951.

Simonyan, K. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.

Sutskever, I.; Martens, J.; Dahl, G.; and Hinton, G. 2013. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th International Conference on Machine Learning*, 1139–1147.

Tabarisaadi, P.; Khosravi, A.; Nahavandi, S.; Shafie-Khah, M.; and Catalão, J. 2022. An optimized uncertainty-aware training framework for neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 35(5): 6928–6935.

Tian, Y.-M.; Liang, S.; Zhang, S.-Q.; and Fan, F.-L. 2025. Depth-induced NTK: Bridging Over-parameterized Neural Networks and Deep Neural Kernels. *arXiv preprint arXiv:2511.05585*.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems 30*, 5998–6008.

Wu, A.; Nowozin, S.; Meeds, E.; Turner, R.; Hernandez-Lobato, J.; and Gaunt, A. 2019. Deterministic variational inference for robust Bayesian neural networks. In *Proceedings of the 7th International Conference on Learning Representations*.

Zhang, S.-H.; Tang, W.-C.; Wu, C.; Hu, P.; Li, N.; Zhang, L.-J.; Zhang, Q.; and Zhang, S.-Q. 2025. TernaryCLIP: Efficiently Compressing Vision-Language Models with Ternary Weights and Distilled Knowledge. *arXiv:2510.21879*.

Zhang, S.-Q.; Wang, F.; and Fan, F.-L. 2024. Neural Network Gaussian Processes by Increasing Depth. *IEEE Transactions on Neural Networks and Learning Systems*, 35(2): 2881–2886.

Zhang, W.; Hou, L.; Yin, Y.; Shang, L.; Chen, X.; Jiang, X.; and Liu, Q. 2020. Ternarybert: Distillation-aware ultra-low bit bert. In *Proceedings of the 25th Empirical Methods in Natural Language Processing*, 509–521.

Zhu, C.; Han, S.; Mao, H.; and Dally, W. 2017. Trained ternary quantization. In *Proceedings of the 5th International Conference on Learning Representations*, 1–10.

Technical Appendix

This Technical Appendix provides the supplementary materials for our work "On the Impact of Weight Quantization on Deep Neural Network Uncertainty", constructed according to the corresponding sections therein.

A About the Assumptions

This appendix validates Assumptions 6 and 7, and discusses the role and interdependencies of all assumptions made throughout this paper.

Seminal studies (Bibi, Alfadly, and Ghanem 2018; Wu et al. 2019) made similar assumptions to Assumption 6 for uncertainty quantification. To further verify the mildness of Assumption 6 and Assumption 7, we conduct the following experiment. We sample weights according to Assumption 1 or Assumption 2, then propagate the networks 10,000 times on a fixed MNIST instance to observe the empirical distributions of hidden layer and pre-activation variable for an MLP-2.

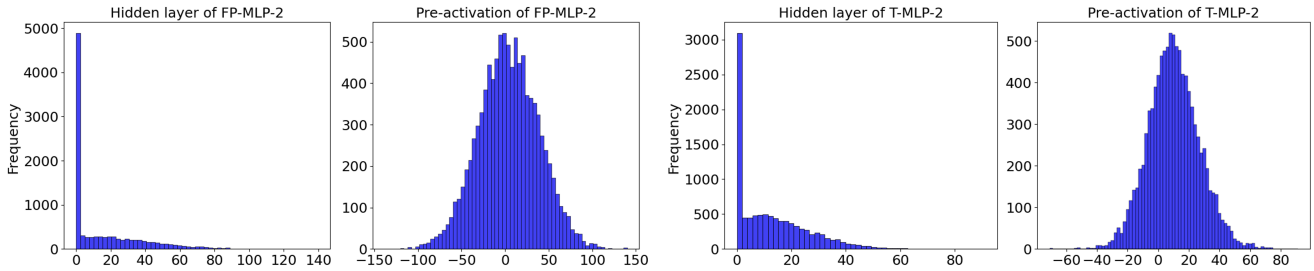

Figure 4: Empirical distributions of the hidden layer and pre-activation variables for FP-MLP-2 and T-MLP-2.

Figure 4 displays the empirical distributions of the pre-activation variable a and the hidden layer h, where the weight vector $\mathbf{W}$ satisfies Assumption 1 for FP-MLP-2 and Assumption 2 for T-MLP-2. In both cases, the empirical distribution of a closely approximates a Gaussian distribution even though h does not, indicating that Assumption 6 and Assumption 7 empirically follow from Assumption 1 and Assumption 2, respectively.

Discussion. The assumptions of weights are listed in Assumptions 1 and 2, where the weights within the same layer are i.i.d. following a Gaussian distribution for FP-DNNs and a ternary-valued distribution for T-DNNs. These assumptions are key to deriving Corollary 4, Corollary 5, and Theorem 9, providing theoretical guarantees for the effectiveness of EMP and MOMA, such as closed-form solutions for linear EMP computations and moment alignment. Besides, Assumptions 6 and 7 empirically follow from Assumptions 1 and 2. Assumptions 6 and 7 are key to deriving Theorem 8, providing theoretical guarantees for the effectiveness of EMP, such as closed-form solutions for nonlinear EMP computations.

Notice that EMP and MOMA are implemented by Algorithm 1 and Algorithm 2, respectively, which make no assumptions about the input distribution in practice. We set the input distribution as a single-point distribution in our experiments, that is, $\boldsymbol{\mu}_{\boldsymbol{x}} = \boldsymbol{x}$ and $\boldsymbol{\Sigma}_{\boldsymbol{x}} = \mathbf{0}$ for an input $\boldsymbol{x}$.

B About the limitation and complexity

B.1 EMP

About the complexity and acceleration. Uncertainty quantification has suffered from high computational demands. We address this by proposing two weight distribution assumptions: a Gaussian distribution for weights in the same layer of FP-DNNs in Assumption 1 and a ternary-valued distribution for T-DNNs in Assumption 2. These assumptions reduce variance-related uncertainty quantification complexity. While traditional BNNs require storage of $2m^2L$ distribution parameters for an MLP with depth L and width m, EMP requires only $2L$ parameters, substantially enhancing efficiency through this significant parameter reduction.

About the limitation. The EMP estimator is specifically designed for DNNs with ReLU activations.

B.2 MOMA

About the complexity. The key idea of MOMA is transforming Eq. (4) into the optimization problem Eq. (6). Therefore, the additional computational complexity introduced by MOMA comes solely from solving the optimization problem Eq. (6). This optimization problem can be efficiently solved using zero-order optimization algorithms such as grid search.

About the limitation. The MOMA is specifically designed for threshold-based quantized training algorithms, where quantized weights are obtained through the partition function Eq. (3) or its variants.

C Full Proof of Theorem 3

The key idea to calculate $\mathbb{E}[a_i]$ is to use the law of total expectation as

$$\begin{aligned}\mathbb{E}\left[a_i\right] &= \mathbb{E}_{\boldsymbol{h}'}\left[\mathbb{E}_{\mathbf{W}}\left[a_i|\boldsymbol{h}'\right]\right] \\ &= \mathbb{E}_{\boldsymbol{h}'}\left[\mathbb{E}_{\mathbf{W}}\left[{\mathbf{W}_i}^\top \boldsymbol{h}' + b_i|\boldsymbol{h}'\right]\right] \\ &= \mathbb{E}_{\boldsymbol{h}'}\left[\mathbb{E}_{\mathbf{W}}\left[{\mathbf{W}_i}^\top\right]\boldsymbol{h}' + b_i\right] \\ &= \mathbb{E}_{\boldsymbol{h}'}\left[\boldsymbol{\mu}_{\mathbf{W}_i}^\top \boldsymbol{h}' + b_i\right] \\ &= \boldsymbol{\mu}_{\mathbf{W}_i}^\top \boldsymbol{\mu}_{\boldsymbol{h}'} + b_i\,.\end{aligned} \tag{7}$$

The key idea to calculate $\mathbb{V}[a_i]$ is to utilize the law of total variance as

$$\begin{aligned}\mathbb{V}\left(a_i\right) &= \mathbb{E}_{\boldsymbol{h}'}\left[\mathbb{V}_{\mathbf{W}}\left(a_i|\boldsymbol{h}'\right)\right] + \mathbb{V}_{\boldsymbol{h}'}\left(\mathbb{E}_{\mathbf{W}}\left[a_i|\boldsymbol{h}'\right]\right) \\ &= \mathbb{E}_{\boldsymbol{h}'}\left[\mathbb{V}_{\mathbf{W}}\left({\mathbf{W}_i}^\top \boldsymbol{h}' + b_i|\boldsymbol{h}'\right)\right] + \mathbb{V}_{\boldsymbol{h}'}\left(\mathbb{E}_{\mathbf{W}}\left[{\mathbf{W}_i}^\top \boldsymbol{h}' + b_i|\boldsymbol{h}'\right]\right) \\ &= \mathbb{E}_{\boldsymbol{h}'}\left[\left(\boldsymbol{h}'\right)^\top \boldsymbol{\Sigma}_{\mathbf{W}_i}\boldsymbol{h}'\right] + \mathbb{V}_{\boldsymbol{h}'}\left(\boldsymbol{\mu}_{\mathbf{W}_i}^\top \boldsymbol{h}' + b_i\right) \\ &= \boldsymbol{\mu}_{\boldsymbol{h}'}^\top \boldsymbol{\Sigma}_{\mathbf{W}_i}\boldsymbol{\mu}_{\boldsymbol{h}'} + \mathrm{tr}\left(\boldsymbol{\Sigma}_{\mathbf{W}_i}\boldsymbol{\Sigma}_{\boldsymbol{h}'}\right) + \boldsymbol{\mu}_{\mathbf{W}_i}^\top \boldsymbol{\Sigma}_{\boldsymbol{h}'}\boldsymbol{\mu}_{\mathbf{W}_i}\,,\end{aligned} \tag{8}$$

The key idea to calculate $\mathrm{Cov}(a_1, a_2)$ is to use the law of total covariance as

$$\begin{aligned}\mathrm{Cov}\left(a_1, a_2\right) &= \mathbb{E}_{\boldsymbol{x}}\left[\mathrm{Cov}\left(a_1, a_2|\boldsymbol{h}'\right)\right] + \mathrm{Cov}\left(\mathbb{E}_{\mathbf{W}}\left[y_1|\boldsymbol{h}'\right], \mathbb{E}_{\mathbf{W}}\left[y_2|\boldsymbol{h}'\right]\right) \\ &= \mathbb{E}_{\boldsymbol{h}'}\left[\mathrm{Cov}\left(\mathbf{W}_1^\top \boldsymbol{h}' + b_1, \mathbf{W}_2^\top \boldsymbol{h}' + b_2|\boldsymbol{h}'\right)\right] + \mathrm{Cov}\left(\mathbb{E}_{\mathbf{W}}\left[\mathbf{W}_1^\top \boldsymbol{h}' + b_1|\boldsymbol{h}'\right], \mathbb{E}_{\mathbf{W}}\left[\mathbf{W}_2^\top \boldsymbol{h}' + b_2|\boldsymbol{h}'\right]\right) \\ &= \mathbb{E}_{\boldsymbol{h}'}\left[\mathrm{Cov}\left(\mathbf{W}_1^\top \boldsymbol{h}', \mathbf{W}_2^\top \boldsymbol{h}'|\boldsymbol{h}'\right)\right] + \mathrm{Cov}\left(\mathbb{E}_{\mathbf{W}}\left[\mathbf{W}_1^\top\right]\boldsymbol{h}' + b_1, \mathbb{E}_{\mathbf{W}}\left[\mathbf{W}_2^\top\right]\boldsymbol{h}' + b_2\right) \\ &= \mathbb{E}_{\boldsymbol{h}'}\left[\left(\boldsymbol{h}'\right)^\top \mathrm{Cov}\left(\mathbf{W}_1, \mathbf{W}_2\right)\boldsymbol{h}'\right] + \mathrm{Cov}\left(\boldsymbol{\mu}_{\mathbf{W}_1}^\top \boldsymbol{h}', \boldsymbol{\mu}_{\mathbf{W}_2}^\top \boldsymbol{h}'\right) \\ &= \boldsymbol{\mu}_{\boldsymbol{h}'}^\top \mathrm{Cov}\left(\mathbf{W}_1, \mathbf{W}_2\right)\boldsymbol{\mu}_{\boldsymbol{h}'} + \mathrm{tr}\left(\mathrm{Cov}\left(\mathbf{W}_1, \mathbf{W}_2\right)\boldsymbol{\Sigma}_{\boldsymbol{h}'}\right) + \boldsymbol{\mu}_{\mathbf{W}_1}^\top \boldsymbol{\Sigma}_{\boldsymbol{h}'}\boldsymbol{\mu}_{\mathbf{W}_2}\,.\end{aligned} \tag{9}$$

This completes the proof. □

D Full Proof of Theorem 8

We start this proof with some useful notations. We denote the Probability Density Function (PDF) of a univariate Gaussian distribution equipped with μ and σ as

$$\phi\left(x; \mu, \sigma^2\right) = \frac{1}{\sqrt{2\pi}\sigma}\exp\left(-\frac{\left(x-\mu\right)^2}{2\sigma^2}\right)\,.$$

Particularly, when $\mu = 0$ and $\sigma = 1$, this univariate Gaussian distribution is called the standard univariate Gaussian distribution, whose PDF is denoted as

$$\phi\left(x\right) = \frac{1}{\sqrt{2\pi}}\exp\left(-\frac{x^2}{2}\right)\,.$$

Further, we denote the Cumulative Distribution Function (CDF) of the standard univariate Gaussian distribution as

$$\Phi\left(x\right) = \int_{-\infty}^{x}\phi\left(u\right)\mathrm{d}u\,.$$

For a bivariate Gaussian distribution, we denote the PDF of a bivariate normal distribution equipped with

$$\boldsymbol{\mu} = \left(\mu_1, \mu_2\right)^\top \quad \text{and} \quad \boldsymbol{\Sigma} = \begin{pmatrix}\sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2\end{pmatrix}$$

as

$$\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right) = \frac{1}{2\pi\left|\boldsymbol{\Sigma}\right|^{1/2}}\exp\left(-\frac{1}{2}\left(\boldsymbol{x}-\boldsymbol{\mu}\right)^\top \boldsymbol{\Sigma}^{-1}\left(\boldsymbol{x}-\boldsymbol{\mu}\right)\right)\,,$$

where $\boldsymbol{x} = \left(x_1, x_2\right)^\top$. Especially, when

$$\boldsymbol{\mu} = \left(0, 0\right)^\top \quad \text{and} \quad \boldsymbol{\Sigma} = \begin{pmatrix}1 & \rho \\ \rho & 1\end{pmatrix}\,,$$

this univariate Gaussian distribution is called the standard bivariate Gaussian distribution with correlation coefficient ρ, whose PDF is denoted as follows

$$\phi_2\left(x_1, x_2; \rho\right) = \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left(-\frac{1}{2\left(1-\rho^2\right)}\left(x_1^2 - 2\rho x_1 x_2 + x_2^2\right)\right) .$$

Further, we denote the CDF of the standard bivariate Gaussian distribution with correlation coefficient ρ as

$$\Phi_2\left(x_1, x_2; \rho\right) = \int_{-\infty}^{x_1}\int_{-\infty}^{x_2} \phi_2\left(u_1, u_2; \rho\right) \mathrm{d}u_1 \mathrm{d}u_2 .$$

With the notations above, we then introduce some useful equations as

$$\begin{aligned}
\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)|_{x_1=a} &= \phi\left(a; \mu_1, \sigma_1^2\right) \phi\left(x_2; \mu_2 - \rho\frac{\sigma_2}{\sigma_1}\left(\mu_1 - a\right), \left(1-\rho^2\right)\sigma_2^2\right) , \\
\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)|_{x_2=a} &= \phi\left(a; \mu_2, \sigma_2^2\right) \phi\left(x_1; \mu_1 - \rho\frac{\sigma_1}{\sigma_2}\left(\mu_2 - a\right), \left(1-\rho^2\right)\sigma_1^2\right) .
\end{aligned}$$

Hence, the following holds.

$$\begin{aligned}
&\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)|_{x_1=+\infty} \doteq \lim_{a\to+\infty} \phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)|_{x_1=a} = 0, \\
&\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)|_{x_2=+\infty} \doteq \lim_{a\to+\infty} \phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)|_{x_2=a} = 0 , \\
&x_1\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)|_{x_1=0}^{x_1=+\infty} \doteq \left[x_1\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)\right]|_{x_1=+\infty} - \left[x_1\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)\right]|_{x_1=0} = 0 , \\
&\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)|_{x_1=0} = \frac{1}{\sigma_1}\phi\left(\frac{\mu_1}{\sigma_1}\right) \phi\left(x_2; \mu_2 - \rho\frac{\sigma_2}{\sigma_1}\mu_1, \left(1-\rho^2\right)\sigma_2^2\right) , \\
&\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)|_{x_2=0} = \frac{1}{\sigma_2}\phi\left(\frac{\mu_2}{\sigma_2}\right) \phi\left(x_1; \mu_1 - \rho\frac{\sigma_1}{\sigma_2}\mu_2, \left(1-\rho^2\right)\sigma_1^2\right) .
\end{aligned} \tag{10}$$

To calculate $\mathbb{E}\left[h_i\right]$ and $\mathbb{V}\left(h_i\right)$, we follow the key idea proposed by Frey and Hinton (1999). Accordingly, it suffices to give an analytic expression on $\mathbb{E}(h_1, h_2)$, that is,

$$\begin{aligned}
\mathbb{E}(h_1, h_2) &= \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} \tau\left(x_1\right)\tau\left(x_2\right)\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)\mathrm{d}x_1\mathrm{d}x_2 \\
&= \int_{0}^{+\infty}\int_{0}^{+\infty} x_1 x_2 \phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)\mathrm{d}x_1\mathrm{d}x_2 .
\end{aligned} \tag{11}$$

The key idea is to solve the integral in Eq. (11) recursively, that is, we first build a recursive formula between the integral in Eq. (11) and $\int_0^{+\infty}\int_0^{+\infty} x_1\phi_2(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})\mathrm{d}x_1\mathrm{d}x_2$, and then solve the latter one. To build this recursive formula, we first note that

$$\frac{\partial\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)}{\partial\boldsymbol{x}} = -\boldsymbol{\Sigma}^{-1}\left(\boldsymbol{x} - \boldsymbol{\mu}\right)\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right) . \tag{12}$$

We then multiply both sides of Eq. (12) by $\boldsymbol{\Sigma}$, explicitly write the components of the vectors $\partial\phi_2(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})/\partial\boldsymbol{x}$ and $(\boldsymbol{x} - \boldsymbol{\mu})$, and obtain

$$\boldsymbol{\Sigma}\begin{pmatrix} \partial\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)/\partial x_1 \\ \partial\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)/\partial x_2 \end{pmatrix} = -\begin{pmatrix} x_1\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right) - \mu_1\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right) \\ x_2\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right) - \mu_2\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right) \end{pmatrix} . \tag{13}$$

We further unfold the left-hand side of Eq. (13) and obtain

$$\begin{pmatrix} \sigma_1^2\partial\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)/\partial x_1 + \rho\sigma_1\sigma_2\partial\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)/\partial x_2 \\ \rho\sigma_1\sigma_2\partial\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)/\partial x_1 + \sigma_2^2\partial\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)/\partial x_2 \end{pmatrix} = -\begin{pmatrix} x_1\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right) - \mu_1\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right) \\ x_2\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right) - \mu_2\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right) \end{pmatrix} . \tag{14}$$

Observing the first and the second components of the vectors on both sides of Eq. (14), we have

$$x_1\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right) = \mu_1\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right) - \sigma_1^2\frac{\partial\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)}{\partial x_1} - \rho\sigma_1\sigma_2\frac{\partial\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)}{\partial x_2} , \tag{15}$$

and

$$x_2\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right) = \mu_2\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right) - \rho\sigma_1\sigma_2\frac{\partial\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)}{\partial x_1} - \sigma_2^2\frac{\partial\phi_2\left(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)}{\partial x_2} . \tag{16}$$

Multiplying x_1 on both sides of Eq. (16), we have

$$x_1 x_2 \phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right) = x_1\mu_2\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right) - \rho\sigma_1\sigma_2 x_1 \frac{\partial\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)}{\partial x_1} - \sigma_2^2 x_1 \frac{\partial\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)}{\partial x_2} \ . \tag{17}$$

Integrating both sides of Eq. (17) over $(0,+\infty)\times(0,+\infty)$, we have the concerned recursive formula as follows.

$$\begin{aligned}
&\int_0^\infty\int_0^\infty x_1 x_2 \phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)\mathrm{d}x_1\mathrm{d}x_2 = \mu_2\int_0^\infty\int_0^\infty x_1\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)\mathrm{d}x_1\mathrm{d}x_2 - \rho\sigma_1\sigma_2\int_0^\infty\int_0^\infty x_1\frac{\partial\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)}{\partial x_1}\mathrm{d}x_1\mathrm{d}x_2 \\
&\quad - \sigma_2^2\int_0^\infty\int_0^\infty x_1\frac{\partial\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)}{\partial x_2}\mathrm{d}x_1\mathrm{d}x_2 \\
&= \mu_2\int_0^\infty\int_0^\infty x_1\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)\mathrm{d}x_1\mathrm{d}x_2 - \rho\sigma_1\sigma_2\int_0^\infty\left(x_1\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)|_{x_1=0}^{x_1=\infty} - \int_0^\infty\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)\mathrm{d}x_1\right)\mathrm{d}x_2 \\
&\quad - \sigma_2^2\int_0^\infty x_1\left(\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)|_{x_2=0}^{x_2=\infty}\right)\mathrm{d}x_1 \\
&= \mu_2\int_0^\infty\int_0^\infty x_1\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)\mathrm{d}x_1\mathrm{d}x_2 + \rho\sigma_1\sigma_2\int_0^\infty\int_0^\infty\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)\mathrm{d}x_1\mathrm{d}x_2 \\
&\quad + \sigma_2\phi\left(\frac{\mu_2}{\sigma_2}\right)\int_0^\infty x_1\phi\left(x_1;\mu_1-\rho\frac{\sigma_1}{\sigma_2}\mu_2,\left(1-\rho^2\right)\sigma_1^2\right)\mathrm{d}x_1 \\
&= \mu_2\int_0^\infty\int_0^\infty x_1\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)\mathrm{d}x_1\mathrm{d}x_2 + \rho\sigma_1\sigma_2\Phi_2\left(\frac{\mu_1}{\sigma_1},\frac{\mu_2}{\sigma_2};\rho\right) + \sigma_1\sigma_2\sqrt{1-\rho^2}\phi\left(\frac{\mu_2}{\sigma_2}\right)\phi\left(c_1\right) \\
&\quad + \sigma_1\sigma_2\phi\left(\frac{\mu_2}{\sigma_2}\right)\left(\frac{\mu_1}{\sigma_1}-\rho\frac{\mu_2}{\sigma_2}\right)\Phi\left(c_1\right) \ ,
\end{aligned} \tag{18}$$

where $c_i = (r_i - \rho r_{3-i})/\sqrt{1-\rho^2}$ for $i\in\{1,2\}$, the second equality uses integration by parts, and the third equality utilizes Eq. (10). Next, it suffices to solve $\int_0^{+\infty}\int_0^{+\infty} x_1\phi_2(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma})\mathrm{d}x_1\mathrm{d}x_2$. Integrating both sides of Eq. (15) over $(0,+\infty)\times(0,+\infty)$, one has

$$\begin{aligned}
&\int_0^\infty\int_0^\infty x_1\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)\mathrm{d}x_1\mathrm{d}x_2 = \int_0^\infty\int_0^\infty \mu_1\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)\mathrm{d}x_1\mathrm{d}x_2 - \int_0^\infty\int_0^\infty \sigma_1^2\frac{\partial\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)}{\partial x_1}\mathrm{d}x_1\mathrm{d}x_2 \\
&\quad - \int_0^\infty\int_0^\infty \rho\sigma_1\sigma_2\frac{\partial\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)}{\partial x_2}\mathrm{d}x_1\mathrm{d}x_2 \\
&= \mu_1\Phi_2\left(\frac{\mu_1}{\sigma_1},\frac{\mu_2}{\sigma_2};\rho\right) - \sigma_1^2\int_0^\infty\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)|_{x_1=0}^{x_1=\infty}\,\mathrm{d}x_2 - \rho\sigma_1\sigma_2\int_0^\infty\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)|_{x_2=0}^{x_2=\infty}\,\mathrm{d}x_1 \\
&= \mu_1\Phi_2\left(\frac{\mu_1}{\sigma_1},\frac{\mu_2}{\sigma_2};\rho\right) + \sigma_1\phi\left(\frac{\mu_1}{\sigma_1}\right)\int_0^\infty\phi\left(x_2;\mu_2-\rho\frac{\sigma_2}{\sigma_1}\mu_1,\left(1-\rho^2\right)\sigma_2^2\right)\mathrm{d}x_2 \\
&\quad + \rho\sigma_1\phi\left(\frac{\mu_2}{\sigma_2}\right)\int_0^\infty\phi\left(x_1;\mu_1-\rho\frac{\sigma_1}{\sigma_2}\mu_2,\left(1-\rho^2\right)\sigma_1^2\right)\mathrm{d}x_1 \\
&= \mu_1\Phi_2\left(\frac{\mu_1}{\sigma_1},\frac{\mu_2}{\sigma_2};\rho\right) + \sigma_1\phi\left(\frac{\mu_1}{\sigma_1}\right)\Phi\left(c_2\right) + \rho\sigma_1\phi\left(\frac{\mu_2}{\sigma_2}\right)\Phi\left(c_1\right) \ ,
\end{aligned} \tag{19}$$

where $r_i = \mu_i/\sigma_i$ for $i\in\{1,2\}$, the second equality utilizes integration by parts, and the third equality uses Eq. (10). Substituting Eq. (19) into Eq. (18), one has

$$\begin{aligned}
&\int_0^\infty\int_0^\infty x_1 x_2\phi_2\left(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}\right)\mathrm{d}x_1\mathrm{d}x_2 = \\
&\qquad \sigma_1\sigma_2\left[r_2\phi\left(r_1\right)\Phi\left(c_2\right) + r_1\phi\left(r_2\right)\Phi\left(c_1\right) + \left(r_1 r_2+\rho\right)\Phi_2\left(r_1,r_2;\rho\right) + \sqrt{1-\rho^2}\phi\left(r_2\right)\phi\left(c_1\right)\right] .
\end{aligned}$$

This completes the proof. □

E Full Proof of Theorem 9

According to Corollary 5, if

$$\begin{cases} q^{(l)} - p^{(l)} = 0 \ , \\ p^{(l)} + q^{(l)} - \left(q^{(l)} - p^{(l)}\right)^2 = 0 \ , \end{cases}$$

for $l \in [L]$, then one has

$$\mathbb{E}\left[\boldsymbol{a}^{(l)}\right] = \mathbf{0}\,, \quad \mathbb{V}\left[\boldsymbol{a}^{(l)}\right] = \mathbf{0}\,, \quad \mathrm{Cov}\left(\boldsymbol{a}^{(l)}, \boldsymbol{a}^{(l)}\right) = \mathbf{0}\,,$$

where

$$\boldsymbol{a}^{(l)} = \mathbf{W}^{(l)}\boldsymbol{h}^{(l-1)} + \boldsymbol{b}^{(l)}\,.$$

It is obvious that

$$\boldsymbol{\Sigma}_{\hat{\boldsymbol{y}}} = \mathrm{Cov}\left(\boldsymbol{a}^{(L)}, \boldsymbol{a}^{(L)}\right) = \mathbf{0}\,,$$

which completes the proof. □

This theorem admits an alternative proof. It suffices to note that

$$\begin{cases} q^{(l)} - p^{(l)} = 0\,, \\ p^{(l)} + q^{(l)} - \left(q^{(l)} - p^{(l)}\right)^2 = 0\,, \end{cases}$$

are equivalent to

$$p^{(l)} = q^{(l)} = 0 \quad \text{for} \quad l \in [L]\,.$$

Thus, distribution $\mathcal{T}(p^{(l)}, 1 - p^{(l)} - q^{(l)}, q^{(l)})$ is considered deterministic, leading to $\boldsymbol{\Sigma}_{\hat{\boldsymbol{y}}} = \mathbf{0}$ and thus completing the proof. □

F Experiments Details of Subsection 5.1

This appendix provides the experimental details for Subsection 5.1.

Table 3, Table 4, and Table 5 show the uncertainty estimation performance of EMP and its contenders, including PL-DNN (Bibi, Alfadly, and Ghanem 2018) and MP (Wu et al. 2019) across datasets for MLP-2, LeNet-5, and VGG-13, respectively. We plot the results of only the first 10 components of CIFAR-100 due to space constraints. It is observed that the EMP method achieves near-ideal performance and significantly outperforms its contenders. This observation supports the effectiveness of the proposed EMP estimator.

Datasets	o	FP-MLP-2			T-MLP-2		
		EMP	MP	PL-DNN	EMP	MP	PL-DNN
MNIST	0	$1.0001_{0.0002} \pm 0.0101_{0.0008}$	$1.0818_{0.0226} \pm 0.2029_{0.0401}$	$0.9083_{0.0042} \pm 0.1957_{0.0250}$	$0.9954_{0.0004} \pm 0.0103_{0.0010}$	–	–
	1	$1.0000_{0.0001} \pm 0.0100_{0.0008}$	$1.0818_{0.0226} \pm 0.2032_{0.0400}$	$0.9083_{0.0043} \pm 0.1957_{0.0250}$	$0.9953_{0.0003} \pm 0.0104_{0.0010}$	–	–
	2	$0.9999_{0.0001} \pm 0.0100_{0.0008}$	$1.0818_{0.0226} \pm 0.2031_{0.0401}$	$0.9086_{0.0044} \pm 0.1958_{0.0250}$	$0.9953_{0.0003} \pm 0.0104_{0.0009}$	–	–
	3	$0.9998_{0.0002} \pm 0.0100_{0.0007}$	$1.0818_{0.0226} \pm 0.2029_{0.0401}$	$0.9085_{0.0043} \pm 0.1958_{0.0250}$	$0.9955_{0.0002} \pm 0.0105_{0.0009}$	–	–
	4	$1.0000_{0.0002} \pm 0.0099_{0.0008}$	$1.0817_{0.0225} \pm 0.2030_{0.0401}$	$0.9084_{0.0044} \pm 0.1957_{0.0250}$	$0.9956_{0.0004} \pm 0.0103_{0.0010}$	–	–
	5	$1.0000_{0.0001} \pm 0.0100_{0.0008}$	$1.0818_{0.0225} \pm 0.2031_{0.0401}$	$0.9084_{0.0044} \pm 0.1957_{0.0250}$	$0.9954_{0.0004} \pm 0.0104_{0.0010}$	–	–
	6	$1.0000_{0.0000} \pm 0.0099_{0.0007}$	$1.0817_{0.0226} \pm 0.2030_{0.0399}$	$0.9084_{0.0043} \pm 0.1959_{0.0250}$	$0.9952_{0.0005} \pm 0.0104_{0.0009}$	–	–
	7	$1.0001_{0.0001} \pm 0.0100_{0.0007}$	$1.0818_{0.0226} \pm 0.2032_{0.0400}$	$0.9084_{0.0043} \pm 0.1957_{0.0250}$	$0.9952_{0.0005} \pm 0.0104_{0.0009}$	–	–
	8	$0.9999_{0.0002} \pm 0.0101_{0.0007}$	$1.0819_{0.0224} \pm 0.2030_{0.0402}$	$0.9083_{0.0044} \pm 0.1956_{0.0250}$	$0.9951_{0.0004} \pm 0.0106_{0.0009}$	–	–
	9	$1.0000_{0.0001} \pm 0.0100_{0.0007}$	$1.0819_{0.0226} \pm 0.2031_{0.0400}$	$0.9084_{0.0043} \pm 0.1959_{0.0249}$	$0.9954_{0.0004} \pm 0.0104_{0.0010}$	–	–
CIFAR-10	0	$1.0000_{0.0003} \pm 0.0077_{0.0009}$	$1.1916_{0.0281} \pm 0.2538_{0.0167}$	$0.7905_{0.0258} \pm 0.3391_{0.0212}$	$0.9947_{0.0007} \pm 0.0108_{0.0011}$	–	–
	1	$0.9999_{0.0003} \pm 0.0078_{0.0011}$	$1.1912_{0.0275} \pm 0.2535_{0.0165}$	$0.7903_{0.0263} \pm 0.3389_{0.0212}$	$0.9947_{0.0006} \pm 0.0106_{0.0013}$	–	–
	2	$1.0001_{0.0003} \pm 0.0077_{0.0011}$	$1.1912_{0.0276} \pm 0.2537_{0.0165}$	$0.7905_{0.0264} \pm 0.3390_{0.0213}$	$0.9948_{0.0006} \pm 0.0109_{0.0012}$	–	–
	3	$0.9999_{0.0002} \pm 0.0077_{0.0011}$	$1.1916_{0.0277} \pm 0.2537_{0.0165}$	$0.7906_{0.0261} \pm 0.3391_{0.0212}$	$0.9945_{0.0005} \pm 0.0108_{0.0013}$	–	–
	4	$0.9998_{0.0002} \pm 0.0078_{0.0012}$	$1.1914_{0.0279} \pm 0.2538_{0.0164}$	$0.7905_{0.0263} \pm 0.3392_{0.0211}$	$0.9947_{0.0007} \pm 0.0106_{0.0012}$	–	–
	5	$1.0001_{0.0003} \pm 0.0076_{0.0010}$	$1.1913_{0.0276} \pm 0.2539_{0.0169}$	$0.7905_{0.0262} \pm 0.3390_{0.0213}$	$0.9948_{0.0007} \pm 0.0107_{0.0011}$	–	–
	6	$1.0000_{0.0003} \pm 0.0078_{0.0008}$	$1.1915_{0.0279} \pm 0.2538_{0.0166}$	$0.7905_{0.0263} \pm 0.3392_{0.0210}$	$0.9946_{0.0008} \pm 0.0107_{0.0015}$	–	–
	7	$1.0003_{0.0002} \pm 0.0077_{0.0011}$	$1.1916_{0.0275} \pm 0.2539_{0.0165}$	$0.7907_{0.0260} \pm 0.3392_{0.0211}$	$0.9946_{0.0008} \pm 0.0108_{0.0015}$	–	–
	8	$0.9999_{0.0003} \pm 0.0079_{0.0012}$	$1.1913_{0.0277} \pm 0.2536_{0.0165}$	$0.7906_{0.0263} \pm 0.3392_{0.0212}$	$0.9946_{0.0007} \pm 0.0107_{0.0011}$	–	–
	9	$0.9999_{0.0001} \pm 0.0078_{0.0011}$	$1.1914_{0.0275} \pm 0.2537_{0.0166}$	$0.7904_{0.0263} \pm 0.3391_{0.0211}$	$0.9948_{0.0007} \pm 0.0105_{0.0012}$	–	–
CIFAR-100	0	$1.0003_{0.0008} \pm 0.0074_{0.0008}$	$1.2432_{0.0339} \pm 0.2675_{0.0154}$	$0.6636_{0.0217} \pm 0.4170_{0.0146}$	$0.9942_{0.0012} \pm 0.0104_{0.0014}$	–	–
	1	$1.0004_{0.0007} \pm 0.0070_{0.0006}$	$1.2430_{0.0341} \pm 0.2673_{0.0153}$	$0.6630_{0.0214} \pm 0.4170_{0.0146}$	$0.9946_{0.0016} \pm 0.0105_{0.0012}$	–	–
	2	$1.0001_{0.0006} \pm 0.0072_{0.0008}$	$1.2424_{0.0334} \pm 0.2661_{0.0145}$	$0.6634_{0.0215} \pm 0.4170_{0.0145}$	$0.9946_{0.0015} \pm 0.0105_{0.0011}$	–	–
	3	$0.9997_{0.0007} \pm 0.0066_{0.0011}$	$1.2428_{0.0339} \pm 0.2673_{0.0150}$	$0.6631_{0.0210} \pm 0.4174_{0.0146}$	$0.9948_{0.0015} \pm 0.0108_{0.0011}$	–	–
	4	$1.0003_{0.0010} \pm 0.0070_{0.0012}$	$1.2436_{0.0337} \pm 0.2674_{0.0154}$	$0.6635_{0.0212} \pm 0.4170_{0.0150}$	$0.9946_{0.0015} \pm 0.0108_{0.0015}$	–	–
	5	$1.0001_{0.0004} \pm 0.0069_{0.0008}$	$1.2433_{0.0343} \pm 0.2673_{0.0150}$	$0.6632_{0.0207} \pm 0.4169_{0.0144}$	$0.9942_{0.0011} \pm 0.0107_{0.0012}$	–	–
	6	$1.0000_{0.0003} \pm 0.0069_{0.0011}$	$1.2436_{0.0338} \pm 0.2676_{0.0149}$	$0.6630_{0.0211} \pm 0.4168_{0.0149}$	$0.9943_{0.0016} \pm 0.0105_{0.0012}$	–	–
	7	$1.0002_{0.0006} \pm 0.0073_{0.0008}$	$1.2429_{0.0337} \pm 0.2668_{0.0150}$	$0.6633_{0.0217} \pm 0.4171_{0.0148}$	$0.9943_{0.0013} \pm 0.0107_{0.0014}$	–	–
	8	$1.0000_{0.0006} \pm 0.0066_{0.0006}$	$1.2431_{0.0343} \pm 0.2675_{0.0154}$	$0.6632_{0.0214} \pm 0.4169_{0.0148}$	$0.9943_{0.0017} \pm 0.0107_{0.0012}$	–	–
	9	$1.0000_{0.0006} \pm 0.0071_{0.0012}$	$1.2431_{0.0339} \pm 0.2675_{0.0158}$	$0.6632_{0.0212} \pm 0.4168_{0.0146}$	$0.9948_{0.0013} \pm 0.0106_{0.0015}$	–	–

Table 3: Performance of the EMP and contenders for MLP-2 across datasets.

Datasets	o	FP-MLP-2			T-MLP-2		
		EMP	MP	PL-DNN	EMP	MP	PL-DNN
MNIST	0	$1.0007_{0.0006} \pm 0.0101_{0.0002}$	$1.1809_{0.0301} \pm 0.0217_{0.0069}$	$1.2772_{0.0497} \pm 0.0153_{0.0011}$	$1.0023_{0.0112} \pm 0.0109_{0.0017}$	–	–
	1	$0.9999_{0.0002} \pm 0.0103_{0.0000}$	$1.1793_{0.0304} \pm 0.0238_{0.0058}$	$1.2759_{0.0499} \pm 0.0139_{0.0008}$	$1.0023_{0.0111} \pm 0.0121_{0.0012}$	–	–
	2	$0.9990_{0.0005} \pm 0.0118_{0.0007}$	$1.1794_{0.0297} \pm 0.0239_{0.0060}$	$1.2748_{0.0501} \pm 0.0170_{0.0015}$	$1.0021_{0.0110} \pm 0.0117_{0.0011}$	–	–
	3	$1.0008_{0.0004} \pm 0.0104_{0.0001}$	$1.1808_{0.0301} \pm 0.0215_{0.0065}$	$1.2770_{0.0498} \pm 0.0153_{0.0011}$	$1.0022_{0.0110} \pm 0.0108_{0.0016}$	–	–
	4	$1.0011_{0.0006} \pm 0.0096_{0.0004}$	$1.1807_{0.0304} \pm 0.0239_{0.0067}$	$1.2775_{0.0497} \pm 0.0136_{0.0011}$	$1.0023_{0.0109} \pm 0.0107_{0.0016}$	–	–
	5	$1.0010_{0.0006} \pm 0.0106_{0.0001}$	$1.1815_{0.0300} \pm 0.0237_{0.0059}$	$1.2773_{0.0498} \pm 0.0149_{0.0010}$	$1.0022_{0.0109} \pm 0.0109_{0.0013}$	–	–
	6	$1.0011_{0.0007} \pm 0.0096_{0.0004}$	$1.1818_{0.0299} \pm 0.0246_{0.0062}$	$1.2774_{0.0496} \pm 0.0130_{0.0011}$	$1.0023_{0.0112} \pm 0.0105_{0.0013}$	–	–
	7	$0.9996_{0.0002} \pm 0.0099_{0.0002}$	$1.1803_{0.0296} \pm 0.0243_{0.0067}$	$1.2755_{0.0500} \pm 0.0143_{0.0009}$	$1.0021_{0.0109} \pm 0.0101_{0.0011}$	–	–
	8	$1.0014_{0.0007} \pm 0.0098_{0.0003}$	$1.1814_{0.0301} \pm 0.0218_{0.0066}$	$1.2422_{0.0714} \pm 0.0140_{0.0010}$	$1.0021_{0.0111} \pm 0.0107_{0.0012}$	–	–
	9	$0.9999_{0.0004} \pm 0.0100_{0.0001}$	$1.1801_{0.0298} \pm 0.0225_{0.0063}$	$1.2759_{0.0497} \pm 0.0142_{0.0009}$	$1.0019_{0.0111} \pm 0.0106_{0.0013}$	–	–
CIFAR-10	0	$1.0098_{0.0017} \pm 0.0099_{0.0001}$	$1.5029_{0.0445} \pm 0.0177_{0.0010}$	$1.6167_{0.0559} \pm 0.0334_{0.0022}$	$1.0141_{0.0133} \pm 0.0113_{0.0017}$	–	–
	1	$1.0092_{0.0017} \pm 0.0105_{0.0001}$	$1.5020_{0.0446} \pm 0.0207_{0.0011}$	$1.6156_{0.0559} \pm 0.0299_{0.0022}$	$1.0136_{0.0133} \pm 0.0115_{0.0018}$	–	–
	2	$1.0090_{0.0018} \pm 0.0127_{0.0002}$	$1.5017_{0.0447} \pm 0.0218_{0.0009}$	$1.6154_{0.0560} \pm 0.0358_{0.0022}$	$1.0141_{0.0135} \pm 0.0110_{0.0017}$	–	–
	3	$1.0100_{0.0017} \pm 0.0110_{0.0002}$	$1.5031_{0.0446} \pm 0.0189_{0.0010}$	$1.6169_{0.0558} \pm 0.0338_{0.0023}$	$1.0141_{0.0133} \pm 0.0115_{0.0021}$	–	–
	4	$1.0106_{0.0017} \pm 0.0098_{0.0004}$	$1.5041_{0.0446} \pm 0.0179_{0.0011}$	$1.6179_{0.0557} \pm 0.0307_{0.0023}$	$1.0137_{0.0136} \pm 0.0116_{0.0015}$	–	–
	5	$1.0099_{0.0017} \pm 0.0109_{0.0001}$	$1.5030_{0.0445} \pm 0.0200_{0.0011}$	$1.6167_{0.0556} \pm 0.0311_{0.0022}$	$1.0139_{0.0133} \pm 0.0118_{0.0022}$	–	–
	6	$1.0103_{0.0017} \pm 0.0094_{0.0002}$	$1.5036_{0.0445} \pm 0.0185_{0.0011}$	$1.6174_{0.0556} \pm 0.0282_{0.0022}$	$1.0139_{0.0134} \pm 0.0109_{0.0018}$	–	–
	7	$1.0090_{0.0017} \pm 0.0099_{0.0003}$	$1.5017_{0.0446} \pm 0.0184_{0.0010}$	$1.6153_{0.0559} \pm 0.0307_{0.0022}$	$1.0139_{0.0134} \pm 0.0120_{0.0014}$	–	–
	8	$1.0103_{0.0017} \pm 0.0093_{0.0001}$	$1.5036_{0.0445} \pm 0.0179_{0.0010}$	$1.6173_{0.0558} \pm 0.0312_{0.0022}$	$1.0140_{0.0133} \pm 0.0118_{0.0015}$	–	–
	9	$1.0095_{0.0017} \pm 0.0100_{0.0001}$	$1.5025_{0.0446} \pm 0.0179_{0.0010}$	$1.6162_{0.0559} \pm 0.0314_{0.0023}$	$1.0141_{0.0135} \pm 0.0109_{0.0020}$	–	–
CIFAR-100	0	$1.0133_{0.0006} \pm 0.0128_{0.0002}$	$0.6602_{0.0171} \pm 0.0959_{0.0034}$	$0.5738_{0.0204} \pm 0.1323_{0.0040}$	$1.0135_{0.0131} \pm 0.0112_{0.0016}$	–	–
	1	$1.0134_{0.0006} \pm 0.0134_{0.0001}$	$0.6603_{0.0171} \pm 0.0961_{0.0033}$	$0.5740_{0.0205} \pm 0.1324_{0.0039}$	$1.0132_{0.0128} \pm 0.0117_{0.0020}$	–	–
	2	$1.0135_{0.0006} \pm 0.0131_{0.0002}$	$0.6602_{0.0171} \pm 0.0949_{0.0034}$	$0.5738_{0.0205} \pm 0.1314_{0.0040}$	$1.0133_{0.0128} \pm 0.0114_{0.0018}$	–	–
	3	$1.0122_{0.0006} \pm 0.0106_{0.0001}$	$0.6598_{0.0170} \pm 0.0974_{0.0033}$	$0.5736_{0.0204} \pm 0.1336_{0.0039}$	$1.0136_{0.0129} \pm 0.0114_{0.0017}$	–	–
	4	$1.0134_{0.0007} \pm 0.0121_{0.0002}$	$0.6603_{0.0171} \pm 0.0960_{0.0034}$	$0.5740_{0.0205} \pm 0.1324_{0.0040}$	$1.0132_{0.0129} \pm 0.0112_{0.0012}$	–	–
	5	$1.0126_{0.0006} \pm 0.0120_{0.0002}$	$0.6598_{0.0171} \pm 0.0959_{0.0033}$	$0.5736_{0.0204} \pm 0.1323_{0.0040}$	$1.0132_{0.0128} \pm 0.0118_{0.0015}$	–	–
	6	$1.0125_{0.0007} \pm 0.0135_{0.0002}$	$0.6596_{0.0172} \pm 0.0951_{0.0034}$	$0.5732_{0.0205} \pm 0.1316_{0.0040}$	$1.0134_{0.0130} \pm 0.0117_{0.0019}$	–	–
	7	$1.0126_{0.0006} \pm 0.0121_{0.0002}$	$0.6600_{0.0171} \pm 0.0970_{0.0034}$	$0.5737_{0.0204} \pm 0.1333_{0.0040}$	$1.0134_{0.0128} \pm 0.0118_{0.0019}$	–	–
	8	$1.0133_{0.0006} \pm 0.0116_{0.0001}$	$0.6603_{0.0171} \pm 0.0964_{0.0033}$	$0.5739_{0.0205} \pm 0.1328_{0.0039}$	$1.0138_{0.0128} \pm 0.0110_{0.0019}$	–	–
	9	$1.0131_{0.0006} \pm 0.0118_{0.0001}$	$0.6602_{0.0171} \pm 0.0963_{0.0033}$	$0.5739_{0.0205} \pm 0.1327_{0.0040}$	$1.0133_{0.0127} \pm 0.0112_{0.0017}$	–	–

Table 4: Performance of the EMP and contenders for LeNet-5 across datasets.

Datasets	o	FP-MLP-2			T-MLP-2		
		EMP	MP	PL-DNN	EMP	MP	PL-DNN
MNIST	0	$0.8923_{0.0521} \pm 0.0089_{0.0017}$	$0.5461_{0.1998} \pm 0.0736_{0.0052}$	$0.5420_{0.2202} \pm 0.0741_{0.0054}$	$0.8875_{0.0633} \pm 0.0174_{0.0013}$	–	–
	1	$0.8934_{0.0520} \pm 0.0089_{0.0016}$	$0.5476_{0.1993} \pm 0.0688_{0.0048}$	$0.5435_{0.2297} \pm 0.0694_{0.0049}$	$0.8873_{0.0635} \pm 0.0177_{0.0015}$	–	–
	2	$0.8937_{0.0523} \pm 0.0091_{0.0018}$	$0.5483_{0.1995} \pm 0.0712_{0.0050}$	$0.5442_{0.2299} \pm 0.0718_{0.0051}$	$0.8869_{0.0632} \pm 0.0169_{0.0014}$	–	–
	3	$0.8935_{0.0522} \pm 0.0093_{0.0017}$	$0.5489_{0.1996} \pm 0.0722_{0.0053}$	$0.5448_{0.2201} \pm 0.0727_{0.0054}$	$0.8873_{0.0634} \pm 0.0174_{0.0012}$	–	–
	4	$0.8935_{0.0521} \pm 0.0089_{0.0016}$	$0.5490_{0.1997} \pm 0.0727_{0.0055}$	$0.5449_{0.2200} \pm 0.0732_{0.0056}$	$0.8875_{0.0631} \pm 0.0172_{0.0014}$	–	–
	5	$0.8935_{0.0520} \pm 0.0089_{0.0017}$	$0.5492_{0.1994} \pm 0.0704_{0.0051}$	$0.5451_{0.2298} \pm 0.0710_{0.0052}$	$0.8868_{0.0635} \pm 0.0173_{0.0013}$	–	–
	6	$0.8940_{0.0524} \pm 0.0091_{0.0018}$	$0.5495_{0.1992} \pm 0.0702_{0.0049}$	$0.5454_{0.2296} \pm 0.0707_{0.0050}$	$0.8874_{0.0633} \pm 0.0176_{0.0015}$	–	–
	7	$0.8934_{0.0522} \pm 0.0090_{0.0017}$	$0.5483_{0.1995} \pm 0.0713_{0.0054}$	$0.5442_{0.2299} \pm 0.0718_{0.0055}$	$0.8867_{0.0634} \pm 0.0177_{0.0013}$	–	–
	8	$0.8934_{0.0519} \pm 0.0087_{0.0016}$	$0.5504_{0.1999} \pm 0.0734_{0.0057}$	$0.5462_{0.2203} \pm 0.0739_{0.0058}$	$0.8874_{0.0632} \pm 0.0173_{0.0014}$	–	–
	9	$0.8937_{0.0521} \pm 0.0089_{0.0017}$	$0.5505_{0.1993} \pm 0.0706_{0.0052}$	$0.5463_{0.2297} \pm 0.0712_{0.0053}$	$0.8876_{0.0631} \pm 0.0172_{0.0012}$	–	–
CIFAR-10	0	$0.8363_{0.0625} \pm 0.0153_{0.0021}$	$0.3579_{0.1534} \pm 0.1387_{0.0075}$	$0.3567_{0.1535} \pm 0.1389_{0.0076}$	$0.8441_{0.0757} \pm 0.0382_{0.0028}$	–	–
	1	$0.8358_{0.0624} \pm 0.0152_{0.0020}$	$0.3577_{0.1532} \pm 0.1380_{0.0073}$	$0.3565_{0.1533} \pm 0.1381_{0.0074}$	$0.8434_{0.0755} \pm 0.0376_{0.0026}$	–	–
	2	$0.8361_{0.0623} \pm 0.0146_{0.0019}$	$0.3596_{0.1539} \pm 0.1446_{0.0078}$	$0.3584_{0.1540} \pm 0.1447_{0.0079}$	$0.8442_{0.0758} \pm 0.0385_{0.0029}$	–	–
	3	$0.8361_{0.0626} \pm 0.0158_{0.0022}$	$0.3569_{0.1535} \pm 0.1392_{0.0076}$	$0.3557_{0.1536} \pm 0.1393_{0.0077}$	$0.8442_{0.0757} \pm 0.0385_{0.0028}$	–	–
	4	$0.8363_{0.0626} \pm 0.0158_{0.0021}$	$0.3562_{0.1534} \pm 0.1389_{0.0075}$	$0.3550_{0.1535} \pm 0.1390_{0.0076}$	$0.8441_{0.0756} \pm 0.0383_{0.0027}$	–	–
	5	$0.8359_{0.0624} \pm 0.0148_{0.0020}$	$0.3560_{0.1531} \pm 0.1377_{0.0072}$	$0.3548_{0.1532} \pm 0.1378_{0.0073}$	$0.8441_{0.0756} \pm 0.0381_{0.0026}$	–	–
	6	$0.8360_{0.0624} \pm 0.0149_{0.0019}$	$0.3571_{0.1538} \pm 0.1418_{0.0077}$	$0.3559_{0.1539} \pm 0.1419_{0.0078}$	$0.8444_{0.0754} \pm 0.0375_{0.0025}$	–	–
	7	$0.8348_{0.0625} \pm 0.0152_{0.0021}$	$0.3575_{0.1534} \pm 0.1390_{0.0076}$	$0.3563_{0.1535} \pm 0.1391_{0.0077}$	$0.8435_{0.0757} \pm 0.0382_{0.0028}$	–	–
	8	$0.8356_{0.0623} \pm 0.0147_{0.0019}$	$0.3578_{0.1536} \pm 0.1401_{0.0078}$	$0.3566_{0.1537} \pm 0.1403_{0.0079}$	$0.8433_{0.0757} \pm 0.0383_{0.0027}$	–	–
	9	$0.8362_{0.0625} \pm 0.0151_{0.0020}$	$0.3565_{0.1530} \pm 0.1366_{0.0071}$	$0.3553_{0.1531} \pm 0.1367_{0.0072}$	$0.8443_{0.0757} \pm 0.0384_{0.0029}$	–	–
CIFAR-100	0	$0.8258_{0.0636} \pm 0.0246_{0.0028}$	$0.3926_{0.1439} \pm 0.1445_{0.0082}$	$0.3914_{0.1580} \pm 0.1447_{0.0083}$	$0.8553_{0.0748} \pm 0.0327_{0.0023}$	–	–
	1	$0.8256_{0.0636} \pm 0.0246_{0.0027}$	$0.3922_{0.1438} \pm 0.1439_{0.0081}$	$0.3909_{0.1589} \pm 0.1440_{0.0082}$	$0.8544_{0.0746} \pm 0.0321_{0.0022}$	–	–
	2	$0.8256_{0.0638} \pm 0.0253_{0.0029}$	$0.3943_{0.1444} \pm 0.1508_{0.0085}$	$0.3931_{0.1585} \pm 0.1510_{0.0086}$	$0.8554_{0.0748} \pm 0.0328_{0.0024}$	–	–
	3	$0.8251_{0.0637} \pm 0.0247_{0.0028}$	$0.3914_{0.1440} \pm 0.1451_{0.0083}$	$0.3902_{0.1581} \pm 0.1453_{0.0084}$	$0.8552_{0.0746} \pm 0.0321_{0.0021}$	–	–
	4	$0.8251_{0.0637} \pm 0.0247_{0.0028}$	$0.3906_{0.1439} \pm 0.1448_{0.0082}$	$0.3894_{0.1580} \pm 0.1450_{0.0083}$	$0.8556_{0.0749} \pm 0.0328_{0.0024}$	–	–
	5	$0.8249_{0.0638} \pm 0.0252_{0.0029}$	$0.3905_{0.1438} \pm 0.1434_{0.0080}$	$0.3892_{0.1589} \pm 0.1436_{0.0081}$	$0.8548_{0.0747} \pm 0.0326_{0.0023}$	–	–
	6	$0.8258_{0.0638} \pm 0.0251_{0.0029}$	$0.3915_{0.1442} \pm 0.1478_{0.0084}$	$0.3902_{0.1583} \pm 0.1479_{0.0085}$	$0.8546_{0.0746} \pm 0.0322_{0.0022}$	–	–
	7	$0.8250_{0.0637} \pm 0.0248_{0.0028}$	$0.3920_{0.1439} \pm 0.1449_{0.0082}$	$0.3908_{0.1580} \pm 0.1451_{0.0083}$	$0.8552_{0.0748} \pm 0.0326_{0.0023}$	–	–
	8	$0.8251_{0.0637} \pm 0.0247_{0.0028}$	$0.3924_{0.1441} \pm 0.1460_{0.0083}$	$0.3912_{0.1582} \pm 0.1462_{0.0084}$	$0.8553_{0.0748} \pm 0.0327_{0.0023}$	–	–
	9	$0.8253_{0.0638} \pm 0.0253_{0.0029}$	$0.3910_{0.1437} \pm 0.1424_{0.0079}$	$0.3898_{0.1588} \pm 0.1426_{0.0080}$	$0.8554_{0.0749} \pm 0.0328_{0.0024}$	–	–

Table 5: Performance of the EMP and contenders for VGG-13 across datasets.

G Experiments Details of Subsection 5.2

This appendix provides the experimental details for Subsection 5.2.

The hyperparameters of Algorithm 2 are chosen as $c = 0.3$ and $\lambda = 10$ for the MLP-2 architecture, and set as $c = 0.9$ and $\lambda = 1$ for other architectures.

Table 6 shows the uncertainty reduction performance of MOMA by comparing the uncertainty of the T-DNNs and T-DNNs-MOMA. It is observed that the MOMA greatly reduces the uncertainty of T-DNNs across various datasets and DNN architectures. This observation supports the effectiveness of the proposed MOMA method.

Datasets	Models	Accuracy	Uncertainty	Reduction
MNIST	FP-MLP-2	$0.9761_{0.0002}$	$1.4697_{0.0082} \times 10^{3}$	$68.08_{0.1389}\%$
	T-MLP-2	$0.9752_{0.0005}$	$1.2572_{0.0236} \times 10^{5}$	
	T-MLP-2-MOMA	$0.9751_{0.0006}$	$4.0129_{0.0039} \times 10^{4}$	
	FP-LeNet-5	$0.9911_{0.0001}$	$1.7728_{0.0232} \times 10^{3}$	$98.55_{0.0025}\%$
	T-LeNet-5	$0.9903_{0.0003}$	$2.8113_{0.0490} \times 10^{5}$	
	T-LeNet-5-MOMA	$0.9885_{0.0006}$	$4.0906_{0.1433} \times 10^{3}$	
	FP-VGG-13	$0.9985_{0.0001}$	$4.6029_{0.2302} \times 10^{-1}$	$20.86_{2.5366}\%$
	T-VGG-13	$0.9954_{0.0005}$	$1.4500_{0.0210} \times 10^{1}$	
	T-VGG-13-MOMA	$0.9957_{0.0005}$	$1.1463_{0.0198} \times 10^{1}$	
	FP-ResNet-18	$0.9963_{0.0001}$	$6.9757_{0.3743} \times 10^{-3}$	$44.27_{4.4212}\%$
	T-ResNet-18	$0.9938_{0.0002}$	$4.2808_{0.1772} \times 10^{-1}$	
	T-ResNet-18-MOMA	$0.9951_{0.0002}$	$2.4067_{0.0501} \times 10^{-1}$	
CIFAR-10	FP-MLP-2	$0.4806_{0.0010}$	$5.8953_{0.1782} \times 10^{3}$	$86.20_{1.0919}\%$
	T-MLP-2	$0.1066_{0.0058}$	$5.6167_{0.3858} \times 10^{6}$	
	T-MLP-2-MOMA	$0.1055_{0.0059}$	$7.7525_{0.3044} \times 10^{5}$	
	FP-LeNet-5	$0.6651_{0.0003}$	$2.0605_{0.2572} \times 10^{-1}$	$98.53_{1.5265}\%$
	T-LeNet-5	$0.6140_{0.0062}$	$4.3177_{2.0793} \times 10^{2}$	
	T-LeNet-5-MOMA	$0.6138_{0.0066}$	$6.3389_{5.8414} \times 10^{0}$	
	FP-VGG-13	$0.9263_{0.0002}$	$3.1242_{0.3023} \times 10^{-1}$	$31.67_{4.0737}\%$
	T-VGG-13	$0.9212_{0.0082}$	$9.9900_{0.4232} \times 10^{1}$	
	T-VGG-13-MOMA	$0.9096_{0.0083}$	$6.7903_{0.3624} \times 10^{1}$	
	FP-ResNet-18	$0.9395_{0.0002}$	$4.4530_{0.3822} \times 10^{-2}$	$18.27_{0.9785}\%$
	T-ResNet-18	$0.9174_{0.0076}$	$2.9628_{0.0252} \times 10^{-1}$	
	T-ResNet-18-MOMA	$0.9233_{0.0075}$	$2.4215_{0.0249} \times 10^{-1}$	
CIFAR-100	FP-MLP-2	$0.0794_{0.0015}$	$1.0644_{0.0472} \times 10^{5}$	$69.90_{0.7244}\%$
	T-MLP-2	$0.0624_{0.0066}$	$9.9100_{0.1074} \times 10^{6}$	
	T-MLP-2-MOMA	$0.0725_{0.0022}$	$2.9831_{0.0641} \times 10^{6}$	
	FP-LeNet-5	$0.2950_{0.0024}$	$3.4251_{0.0548} \times 10^{3}$	$30.93_{0.5204}\%$
	T-LeNet-5	$0.0102_{0.0001}$	$1.0975_{1.7877} \times 10^{5}$	
	T-LeNet-5-MOMA	$0.0101_{0.0001}$	$7.5830_{0.4088} \times 10^{4}$	
	FP-VGG-13	$0.7026_{0.0009}$	$3.6694_{0.0772} \times 10^{-1}$	$12.51_{5.4372}\%$
	T-VGG-13	$0.6806_{0.0045}$	$2.1427_{0.1035} \times 10^{1}$	
	T-VGG-13-MOMA	$0.6705_{0.0042}$	$1.8700_{0.0981} \times 10^{1}$	
	FP-ResNet-18	$0.7354_{0.0007}$	$1.6989_{0.2109} \times 10^{-2}$	$16.14_{5.6127}\%$
	T-ResNet-18	$0.6672_{0.0038}$	$1.5869_{0.0811} \times 10^{-1}$	
	T-ResNet-18-MOMA	$0.6880_{0.0040}$	$1.3263_{0.0762} \times 10^{-1}$	

Table 6: Accuracy and uncertainty of concerned models on MNIST, CIFAR-10, and CIFAR-100.

Figure 5, Figure 6, Figure 7, and Figure 8 are results for the MLP-2 on MNIST across different ratio ranges, and they serve two purposes. Firstly, they show the uncertainty reduction performance of MOMA by visualizing the uncertainty ratio of samples across a dataset. It is observed that the red bars are taller at lower x-values and shorter at higher x-values. This observation supports the effectiveness of the proposed MOMA method. Secondly, they show that weight ternarization enlarges DNN uncertainty, regardless of whether the MOMA is used during WQ.

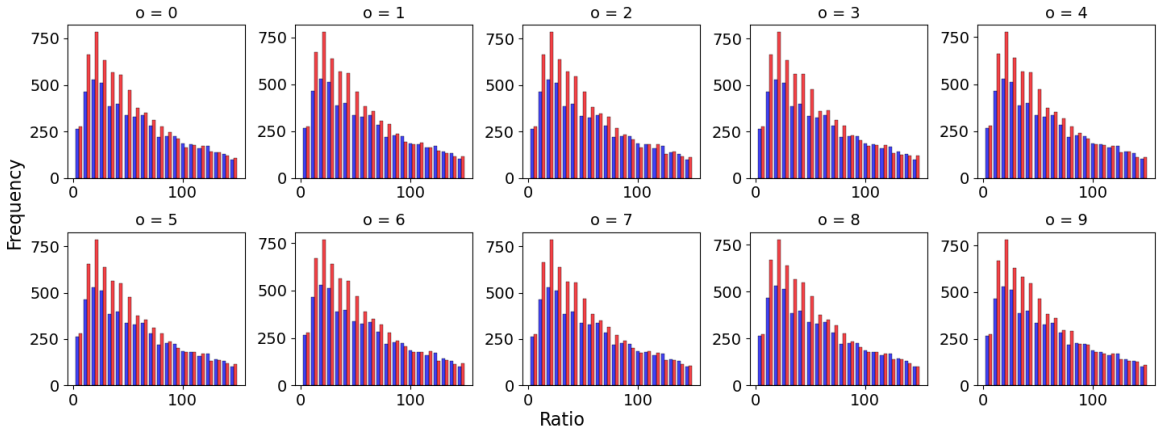

Figure 5: Histograms of $\tilde{r}_o(\boldsymbol{x})$ (blue bars) and $\tilde{r}_o^{'}(\boldsymbol{x})$ (red bars) with range 1.4 to 150.

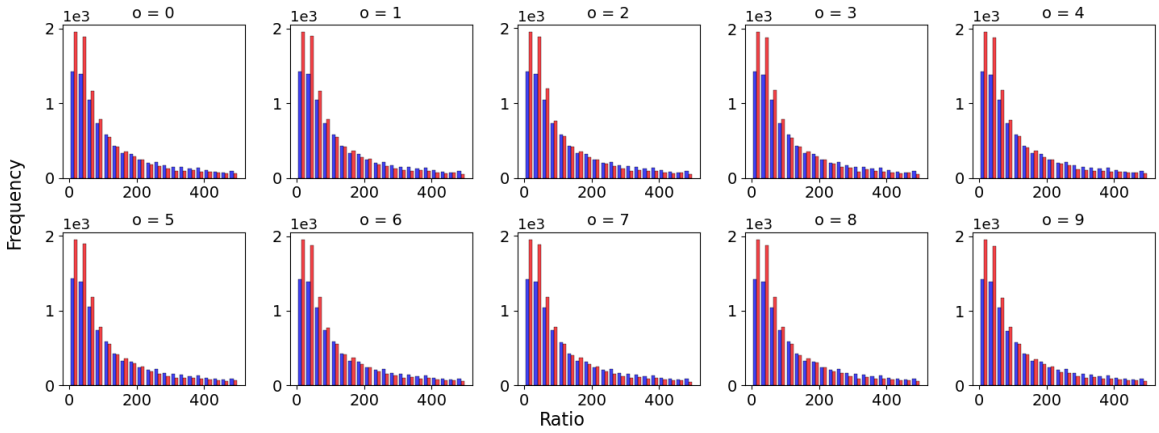

Figure 6: Histograms of $\tilde{r}_o(\boldsymbol{x})$ (blue bars) and $\tilde{r}_o^{'}(\boldsymbol{x})$ (red bars) with range 1.4 to 500.

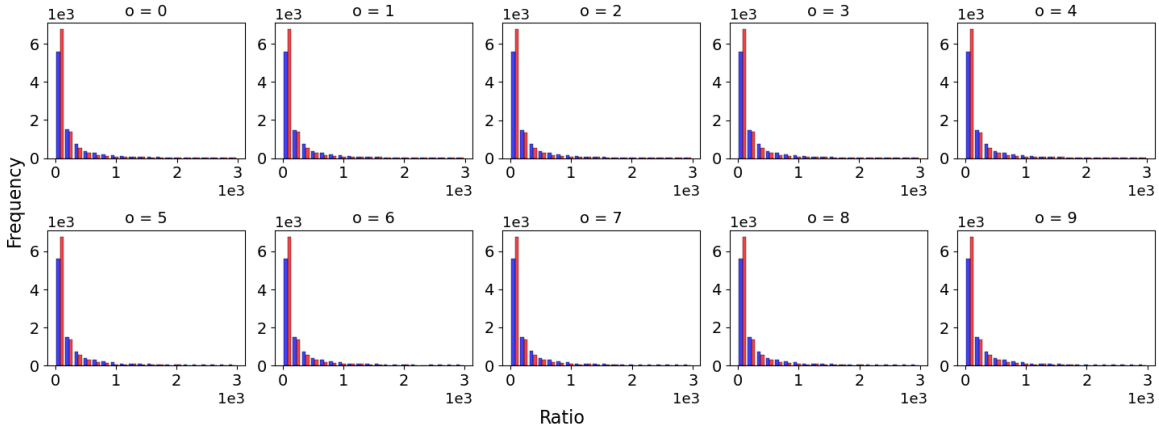

Figure 7: Histograms of $\tilde{r}_o(\boldsymbol{x})$ (blue bars) and $\tilde{r}_o^{'}(\boldsymbol{x})$ (red bars) with range 1.4 to 3000.

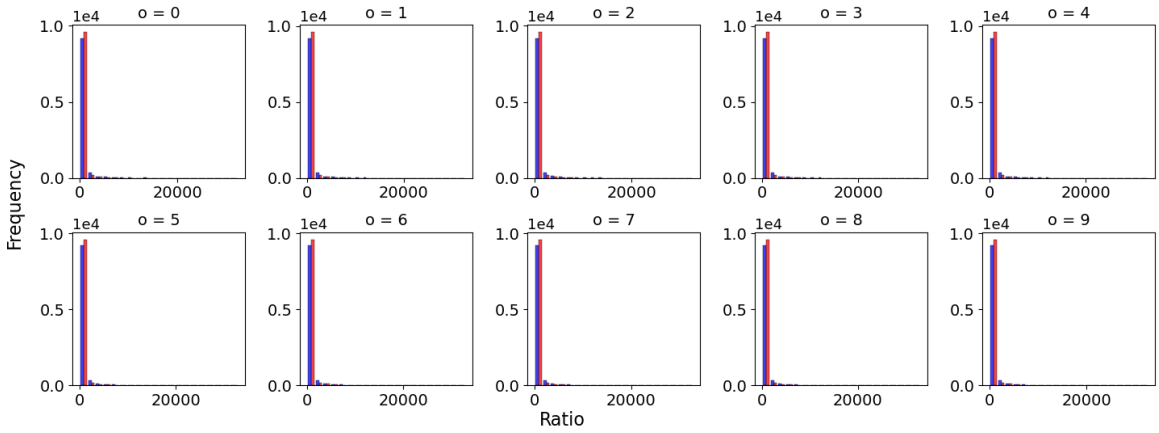

Figure 8: Full Histograms of $\tilde{r}_o(\boldsymbol{x})$ (blue bars) and $\tilde{r}_o^{'}(\boldsymbol{x})$ (red bars) with range 1.4 to 33000.