

QUTE: Quantifying Uncertainty in TinyML models with Early-exit-assisted ensembles for model-monitoring

Nikhil P Ghanathe¹ Steven J E Wilton¹

Abstract

Uncertainty quantification (UQ) provides a resource-efficient solution for on-device monitoring of tinyML models deployed without access to true labels. However, existing UQ methods impose significant memory and compute demands, making them impractical for ultra-low-power, KB-sized TinyML devices. Prior work has attempted to reduce overhead by using early-exit ensembles to quantify uncertainty in a single forward pass, but these approaches still carry prohibitive costs. To address this, we propose QUTE, a novel resource-efficient early-exit-assisted ensemble architecture optimized for tinyML models. QUTE introduces additional output blocks at the final exit of the base network, distilling early-exit knowledge into these blocks to form a diverse yet lightweight ensemble. We show that QUTE delivers superior uncertainty quality on tiny models, achieving comparable performance on larger models with **59% smaller model sizes** than the closest prior work. When deployed on a micro-controller, QUTE demonstrates a **31% reduction in latency** on average. In addition, we show that QUTE excels at detecting accuracy-drop events, outperforming all prior works.

1. Introduction

Recent advancements in embedded systems and machine learning (ML) have produced a new class of edge devices containing powerful ML models. These milliwatt-scale KB-sized devices, often termed *TinyML* devices, have low compute and memory requirements. They are often deployed in remote environments with no availability of true labels. This makes them susceptible to both out-of-distribution (OOD) and covariate-shifted data caused by environmental and sen-

^{*}Equal contribution ¹Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, Canada. Correspondence to: Nikhil P Ghanathe <nikhilghanathe@ece.ubc.ca>.

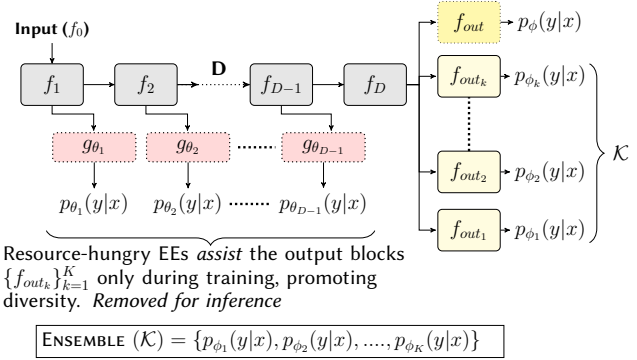


Figure 1: QUTE architecture. $\{f_{out_k}\}_{k=1}^K$ represents the 'K' additional output blocks at the final exit, which are *assisted* by 'K' early-exit blocks $\{g_{\theta_k}\}_{k=1}^K$ *only during training* to promote diversity (see Figure 2). For inference, all (*early-exits & f_{out}*) are removed

sor variations that manifest often unpredictably in the field. The ability of tinyML models to accurately measure the uncertainty of their predictions is crucial for two reasons. First, these devices generate data that is frequently used in critical decision-making. In the context of an autonomous vehicle (AV), uncertain predictions could result in the downstream system making cautious driving decisions (Tang et al., 2022). Second, tinyML devices often operate in remote environments (Vargas et al., 2021). If the inputs to the model changes (perhaps due to spatter, fog, frost, noise, motion blur, etc.), being aware of the model's unreliability may prompt an engineer to take remedial action. Many prior methods for drift/corruption detection are either statistical tests/methods that operate directly on input data (Bifet & Gavalda, 2007) with huge memory requirements, or require true labels for error-rate calculation (Gama et al., 2004; Baier et al., 2022). In contrast, uncertainty quantification (UQ) methods aim to estimate a model's *prediction confidence* without requiring true labels, providing an efficient solution for model monitoring, particularly in resource and power-constrained environments.

In ML models, there are two main types of uncertainty (Kendall & Gal, 2017): *epistemic* and *aleatoric*. Epistemic uncertainty is *reducible* and stems from limited data or knowledge, while aleatoric uncertainty is *irreducible*,

arising from inherent randomness in data or model limitations. In field-deployed tinyML models, both types of uncertainty can emerge due to semantic (OOD) and non-semantic changes in inputs. The non-semantic category, which we term *corrupted-in-distribution* (CID) data, occurs when sensor malfunctions or environmental factors result in corrupted versions of expected data (e.g., fogged camera lens). CID data introduces both aleatoric (e.g., due to random noise or weather) and epistemic (e.g., from blur or digital corruptions) uncertainties, often simultaneously, making it crucial to capture these corruption-induced uncertainties to ensure model *reliability*. Unfortunately, modern neural networks are poor in estimating uncertainty of their predictions (Ovadia et al., 2019). Many prior works have proposed uncertainty-aware networks: ensemble networks (Lakshminarayanan et al., 2017) have been found to be extremely effective in capturing both aleatoric and epistemic uncertainties. Alternatively, early-exit networks have been converted into ensembles (Antoran et al., 2020; Qendro et al., 2021; Ferienc & Rodrigues, 2023). However, both these approaches incur high memory/compute overheads and are untenable for tinyML.

In this paper, we propose QUTE (Figure 1), a novel and resource-efficient early-exit-assisted ensemble architecture that enables high-quality uncertainty quantification in tinyML models in the context of both in-distribution (ID) and corrupted-in-distribution (CID) data. As shown in Figure 1, we append additional lightweight classification heads to an existing base neural network to create ensemble members, and crucially, these classification heads are assisted by the early-exits in a manner similar to the Early-View assistance technique (Ghanathe & Wilton, 2023) to promote diversity. Post-training, we eliminate the resource-hungry early-exits while retaining only the economical additional classification heads. Our approach has significantly less memory and compute overhead compared to prior works (59% smaller models and $3.2\times$ fewer FLOPS compared to the most relevant prior work (Qendro et al., 2021)). Furthermore, QUTE performs better than prior methods in estimating uncertainty caused due to CID error sources, and on-par with prior methods for uncertainty due to OOD. We further show that higher uncertainty is correlated with a drop in accuracy. We evaluate QUTE’s ability to detect such accuracy drop events caused by CID against prior methods, and show that QUTE outperforms all prior methods. To the best of our knowledge, this is the first early-exit ensemble architecture for uncertainty quantification optimized for tinyML models.

This paper is organized as follows. Section 2 presents related work. The context and problem formulation are in Section 3. Our approach is described in Section 4. The experimental methodology and evaluation results are in Sections 5 and 6 respectively. Section 7 concludes the paper.

2. Related Work

Uncertainty quantification Bayesian neural networks (BNN) are well-suited to quantify uncertainty of a model (Blundell et al., 2015; Hernández-Lobato & Adams, 2015; Teye et al., 2018), but are parameter-inefficient and incur a high resource/compute overhead. Monte Carlo Dropout (MCD) (Gal & Ghahramani, 2016) creates implicit ensembles by enabling dropout during multiple inference passes. Recently, ensemble networks like Deep-Ensembles (Lakshminarayanan et al., 2017) have been shown to produce good uncertainty estimates (Zaidi & Zela, 2020; Wenzel et al., 2020; Rahaman et al., 2021). However, this too requires multiple inferences of individual networks, and is impractical in terms of memory for tinyML. Other prior works have proposed multi-input and multi-output networks that combine multiple independent networks into one (Havasi et al., 2021; Ferienc & Rodrigues, 2023), but they do not scale well and remain impractical for tinyML. (Ahmed et al., 2024) ensembles only normalization layers, but requires specialized hardware. Alternatively, prior works (Qendro et al., 2021; Ferienc & Rodrigues, 2023) have leveraged early-exit networks to create implicit ensembles. The closest work to QUTE is EE-ensemble (Qendro et al., 2021), which uses outputs of early-exits as ensemble members. However, the early-exits appended with extra learning layers incur a prohibitive cost (Section 3). In contrast, ensemble distillation methods (Havasi et al., 2021; Malinin et al., 2019; Tran et al., 2020) filter the knowledge of all ensemble members into a conventional neural network (NN). Hydra (Tran et al., 2020), which also uses a single multi-headed network is closest to our work.

Single-pass deterministic network Several single-pass (non-Bayesian) methods for UQ (Van Amersfoort et al., 2020; Mukhoti et al., 2023; Sensoy et al., 2018; Deng et al., 2023) offer lower memory footprints compared to ensemble methods. However, they remain impractical for tinyML due to specialized output layers or architectural constraints. DUQ (Van Amersfoort et al., 2020) adds a specialized layer post-softmax, drastically increasing resource use (e.g., $10\times$ more parameters for ResNet on CIFAR-10). DDU (Mukhoti et al., 2023) simplifies this but relies on residual connections for feature space regularization, limiting its applicability. Priornets (Malinin & Gales, 2018) require OOD data, which is often unrealistic. Postnets (Charpentier et al., 2020), though OOD-free, focus primarily on OOD detection, often sacrificing accuracy below the base network’s level. Other methods like (Meronen et al., 2023) that address overconfidence in early-exit networks require excessive computational resources (e.g., laplace approximations).

Model monitoring There has been a plethora of work to detect OOD samples, which represents a semantic shift

Model	In-distribution data				Corrupted-in-distribution data			
	Acc (↑)	ECE (↓)	BS (↓)	NLL (↓)	Acc (↑)	ECE (↓)	BS (↓)	NLL (↓)
- 1stack	0.72	0.0343	0.0381	0.8174	0.204	0.2796	0.0993	3.216
- 2stack	0.78	0.0315	0.0305	0.6441	0.207	0.2518	0.0969	3.158
- 3stack	0.87	0.0373	0.01907	0.4063	0.27	0.2534	0.0953	3.201
- 4stack	0.89	0.0562	0.01615	0.4019	0.25	0.4255	0.1087	3.677

Table 1: Comparing calibration Metrics of Resnet with {1,2,3,4} stack models on CIFAR10 ID and CID. *Best results marked in bold.* For ECE, BS, NLL *lower is better*

from in-distribution (Yang et al., 2021; 2022; Zhang et al., 2023; Liu et al., 2021; Hendrycks & Gimpel, 2018; Dinari & Freifeld, 2022). However, there have been fewer works that deal with detecting corrupted-ID, which represent a non-semantic shift w.r.t. ID, and are much harder to detect (Liang et al., 2020). Some prior works utilize corrupted-ID/covariate-shifted ID to better generalize on OOD (Bai et al., 2023; Katz-Samuels et al., 2022), which is difficult to obtain. The most relevant work that leverages *only* training data is generalized-ODIN (G-ODIN) (Hsu et al., 2020). It adds 1) a preprocessing layer that perturbs the input image and 2) decomposes confidence score for better OOD detection. However, the preprocessing might prove costly/impractical on tinyML devices. Recent studies (Xia & Bouganis, 2023; Jaeger et al., 2023) propose *failure detection* with a *reject* option, rejecting high-uncertainty samples from both ID (potentially incorrect predictions) and OOD, unlike traditional OOD detection which only separates ID from OOD.

Early-exit networks Early-exit networks add intermediate exits along the length of the base network thereby, providing avenues for reduction in average inference time (Teerapittayanon et al., 2016; Kaya et al., 2019; Huang et al., 2017; Bonato & Bouganis, 2021; Ghanathe & Wilton, 2023; Jazbec et al., 2024). Ghanathe & Wilton (2023) introduces an early-exit architecture optimized for tinyML models. In addition, it introduces the *early-view assistance* method, primarily to mitigate network overthinking in neural networks to boost accuracy. In contrast, we utilize the underlying principle and develop a modified version of this method in our work to principally create a diverse ensemble (Section 4).

3. Background and Problem formulation

Consider an *in-distribution* dataset $S_{ID} = \{x_n, y_n\}_{n=1}^N$ of size N where, x_n and y_n are the n^{th} input sample and its corresponding true label respectively. A discriminative model $\mathcal{M}_\Theta(x)$ learns parameters Θ on S_{ID} and outputs a class posterior probability vector $p_\Theta(y|x)$, which yields a predicted label \hat{y} . For a classification problem, $\hat{y} \in \{1, 2, ..L\}$, where L is the number of classes. Uncertainty quantification methods endeavor to improve uncertainty estimation quality of $\mathcal{M}_\Theta(x)$ on S_{ID} . Specifically, we want to calibrate the model such that its *predictive confidence* ($\mathcal{C}_{\mathcal{M}_\Theta}$) is in sync with its accuracy ($\mathcal{A}_{\mathcal{M}_\Theta}$). A well-calibrated model will see a commensurate drop in $\mathcal{C}_{\mathcal{M}_\Theta}$ as $\mathcal{A}_{\mathcal{M}_\Theta}$ drops. The predictive

confidence is given by, $\mathcal{C}_{\mathcal{M}_\Theta}(x) = \max_{l \in \{1, 2, \dots, L\}} p_\Theta(y = l|x)$

Confidence-calibration is a well studied problem, and it helps achieve synergy between the predicted probabilities and ground truth correctness likelihood (Guo et al., 2017). However, when \mathcal{M}_Θ is deployed in a real-world scenario, it may encounter either 1) out-of-distribution data (S_{OOD}) or 2) a corrupted/covariate-shifted version (S_{CID}), due to environmental/sensor variations (E.g., frost, noise, fog, rain etc.). Unfortunately, we find that despite good uncertainty estimation quality on S_{ID} , many uncertainty-aware networks see a drop in quality and remain overconfident, particularly for severely-corrupted S_{CID} (Hsu et al., 2020). Interestingly, we find that the model becomes *less overconfident* to S_{CID} as its *size shrinks*. This phenomenon is illustrated in Table 1. We report the calibration metrics of Resnet (He et al., 2016) on CIFAR10, with four model sizes ranging from 1 to 4 residual stacks on both S_{ID} and S_{CID} . For calibration metrics (ECE, BS, NLL), *lower is better*. More details on calibration metrics and CIFAR10-CID can be found in Section 5 and Appendix C. As seen, the 4stack model with the highest model capacity is clearly the best on S_{ID} . However, on S_{CID} , it performs the worst, remaining overconfident in the presence of corruptions, even outperformed by 1stack. In contrast, 2stack and 3stack model are much better on S_{CID} . This demonstrates that smaller models are *less overconfident* for corruption in inputs, leading to better calibration on S_{CID} .

The capabilities of smaller models can be harnessed for better uncertainty estimation on S_{CID} through early-exit networks, as illustrated by prior research like EE-ensemble (Qendro et al., 2021), which combines the predictions from multiple early-exits and final exit to form an ensemble. For example, in our evaluations, we find that on TinyImagenet-ID (Le & Yang, 2015) with MobilenetV2 (Howard et al., 2017), the negative log-likelihood (NLL) of EE-ensemble is 26% lower than the popular Deep-Ensemble (Lakshminarayanan et al., 2017), and 34% lower on TinyImagenet-Corrupted (Hendrycks & Dietterich, 2019) (lower is better). This showcases the crucial role of early-exits in achieving higher quality uncertainty overall.

However, we find that EE-ensemble is resource-intensive because it adds additional learning layers at the early-exits to accommodate their varying learning capacities. For example, adding two early-exits *without additional learning layers* (say EE-0 and EE-1) after 1st and 2nd residual stack of a 3-residual stack Resnet for CIFAR10 results in individual accuracies of 0.602 and 0.707 at EE-0 and EE-1 respectively. This is significantly worse than the individual accuracy of the final exit, 0.84, which in turn leads to poor ensemble behavior. To address this, EE-ensemble uses early-exits *with additional densely-connected layers*, which incurs a high memory overhead affecting practicality in

tinyML ($\text{Overhead}_{\text{EE-ensemble}} = 397\text{K}$ params compared to $\text{Overhead}_{\text{QUTE}} = 8.2\text{K}$ params as shown in Section 4 & 6). We present an alternative strategy to leverage the knowledge of early-exits in ensemble-creation that is extremely resource-efficient thereby enabling superior uncertainty estimates on both S_{ID} and S_{CID} , which allows us to reliably detect accuracy-drops in the model due to S_{CID} .

4. QUTE

A neural network (NN) like the one shown in Figure 1, with depth D is composed of several blocks of linear/non-linear operations (e.g., convolution, pooling) that are stacked. The NN consists of $D - 1$ intermediate blocks $\{f_i(\cdot)\}_{i=1}^{D-1}$; $f(\cdot) = f_1(\cdot) \circ f_2(\cdot) \circ f_3(\cdot) \dots \circ f_{D-1}(\cdot) \circ f_D(\cdot)$ and an output block $f_{out}(\cdot)$. $f_0(\cdot)$ is the input block. In the base network, the network processes the input $f_0(\cdot)$ through each block $f_i(\cdot)$ until layer D , ultimately producing prediction $p_{\Theta}(y|x)$ through output block $f_{out}(\cdot)$.

Given a base NN (the grey blocks in Figure 1), we first create K additional classification heads/output blocks $\{f_{out_k}\}_{k=1}^K$ at the final exit (after f_D) as shown in Figure 1. These output blocks constitute the ensemble \mathcal{K} . Next, we add K early-exit blocks, $\{g_{\theta_k}\}_{k=1}^K$ along the length of the base network, which are used *only during training* for knowledge distillation. A more intuitive view of QUTE is presented in the Appendix. In Figure 1, K is set to $D - 1$. In practice, K is a hyperparameter that depends on computation/resource budget and the required ensemble size $|\mathcal{K}|$, where the ensemble size is bounded above by the depth of the network. Thus, $\{f_{out_k}\}_{k=1}^K$ produces K predictions, where each f_{out_k} is assisted in producing predictions by its corresponding early-exit g_{θ_k} using *early-view assistance* method.

$$\mathcal{K} = \{p_{\phi_1}(y|x), p_{\phi_2}(y|x), p_{\phi_3}(y|x), \dots, p_{\phi_K}(y|x)\} \quad (1)$$

Figure 2 illustrates the architecture of *early-view assistance*, where a single final exit block f_{out_k} and its *assisting* early-exit g_{θ_k} is shown. As shown, we create an additional learning layer $h_{\phi_k}(\cdot)$ at the k^{th} output block. An identical learning layer $h_{\theta_k}(\cdot)$ is also added at the k^{th} early-exit. h_{ϕ_k} with parameters ϕ_k is responsible for assimilating knowledge from the corresponding h_{θ_k} with parameters θ_k . $\sigma(\cdot)$ is the output activation (E.g., dense+softmax). h_{ϕ_k} and h_{θ_k} are a single depthwise-convolution layer in our evaluations owing to its low computation and memory demand. We ensure equal dimensionality between h_{ϕ_k} and h_{θ_k} by adding a pointwise-convolution layer before h_{θ_k} at all early-exits (not shown in figure for brevity). Fundamentally, our aim is to imbue the diverse predictive behaviors of early-exits from different depths into the output blocks $\{f_{out_k}\}_{k=1}^K$.

Training Unlike EE-ensemble, all output blocks of QUTE have similar learning capacities since the input to all $\{h_{\phi_k}\}_{k=1}^K$ is f_D of the base network. The output of h_{ϕ_k} is




Figure 2: Early-view-assistance architecture. One *assisting* early-exit g_{θ_k} and corresponding *early-view-assisted* exit f_{out_k} is shown. Early-exits weights θ_k are transferred/copied to h_{ϕ_k} before each train batch, i.e., $\phi_k = \theta_k$

termed as *early-view* since it is assisted by the early-exit. After attaching early-exits (EE) and early-view-assisted (EV) exits to the untrained base network, all weights are learned simultaneously during training. Unlike some prior works (Lakshminarayanan et al., 2017) that use extra data (e.g., adversarial samples), we only use the available training data. We use random data augmentation (E.g., random rotation/flip/crop), a well-known method to improve model robustness to corruptions to ensure fairness in accuracy-drop detection experiments (Section 6.2). The EEs are trained with a *weighted-loss* function, similar to previous studies (Kaya et al., 2019). Further, we assign a *higher weighting* factor, w_{EV_k} for losses at EV-exits to minimize the effect of an overconfident f_{out} (original output block). In addition, we increase each w_{EV_k} by δ such that $w_{EV_k} = w_{EV_{k-1}} + \delta$ to vary the influence each EV-exit has on the weights of the shared base network, thereby promoting diversity. We empirically determine that $w_{EV_0} = 3$ and $\delta = 0.5$ yields the best results (see Appendix A.1 for details). During training,

- 1) Before each train batch, we copy weights of k^{th} EE-exit (h_{θ_k}) into k^{th} EV-exit (h_{ϕ_k}) i.e., $\phi_k = \theta_k$. Thus, each EV-exit iterates through a train batch with weights of the EE-exit. Weight transfer is limited solely from h_{θ_k} to h_{ϕ_k} , with no other layers involved.

- 2) After a train batch, the loss computed at an EV-exit is w.r.t. the weights copied from EE, after which weights of h_{ϕ_k} are updated during backpropagation. However, before the next train batch begins, we overwrite the weights of h_{ϕ_k} by copying from EE. Therefore, the weights learnt by layers before/after h_{ϕ_k} always align with EE weights. Thus, this does not affect convergence (see Figure 5 and Section 6.5).

$$\mathcal{L} = \sum_{k=1}^K \mathcal{L}_{EE_k} + \sum_{k=1}^K \mathcal{L}_{EV_k} + \mathcal{L}_{EF} \quad (2)$$

$$\mathcal{L}_{EE_k} = \tau_k \cdot \mathcal{L}_{CE} \quad \mathcal{L}_{EV_k} = w_{EV_k} \cdot \mathcal{L}_{CE}$$

The training objective reduces to minimizing Eqn 2, where \mathcal{L} is the total loss of QUTE, which is the sum of the losses of EE, EV and final exits. \mathcal{L}_{CE} is the cross-entropy loss, and $\mathcal{L}_{EF} = \mathcal{L}_{CE}$, which is loss of original final exit. \mathcal{L}_{EE_k} and \mathcal{L}_{EV_k} are losses at the k^{th} EE-exit and EV-exit respectively.

τ_k is the weighting applied to EE loss (Kaya et al., 2019). w_{EV_k} is the weighting applied to the k^{th} EV-exit.

The weights of base network (only grey blocks in Figure 1) are frozen for the last 10% of epochs while continuing weight transfer, enabling isolated training of EV-exits to foster increased diversity. In this way, the EE-knowledge is transferred to the final exit(s) to obtain better uncertainty estimation on S_{ID}/S_{CID} . In contrast, without EV-assistance, the various final exits risk learning the same predictive distribution, leading to poor ensemble behavior. We compare the effectiveness of EV-assistance in Section 6.5.

Inference During inference, the resource-heavy EE-blocks are removed from the model, leaving only the base network and the lighter output blocks $\{f_{out_k}(\cdot)\}_{k=1}^K$. We also eliminate the original output block f_{out} due to its overconfident behavior, which compromises calibration quality. The final prediction ($p_{\Theta}(y|x)$) is obtained by calculating the mean of all $|\mathcal{K}|$ prediction vectors of \mathcal{K} .

$$p_{\Theta}(y|x) = \frac{1}{|\mathcal{K}|} \left(\sum_{k=1}^K p_{\phi_k}(y|x) \right) \quad (3)$$

5. Evaluation Methodology

For our microcontroller (MCU) evaluations, we utilize two types of boards.

1) **Big-MCU**: The STM32 Nucleo-144 development board (STM32F767ZI) (STMicroelectronics, 2019) with 2MB flash and 512KB SRAM, clocked at 216MHz with a typical power consumption of 285mW. This board accommodates high-resource baselines that require substantial memory and processing power.

2) **Small-MCU**: The STM32 Nucleo-32 development board (STM32L432KC) (STMicroelectronics, 2018) with 256KB flash and 64KB RAM, clocked at 80MHz with a typical power consumption of 25mW. This ultra-low-power device better reflects our goal of deployment on smaller, energy-efficient systems.

We evaluate QUTE in three settings: 1) Accuracy-drop/CID detection, which monitors system performance over time to detect gradual shifts and temporal degradation patterns (Section 6.2), 2) Failure detection, which targets instance-level identification of both misclassifications and OOD samples (Section 6.3), and 3) uncertainty estimation quality (Section 6.4). Since the focus of this work is to detect failures reliably in-the-field, accuracy-drop and failure detection experiments evaluate this capability.

Datasets and Models We evaluate QUTE on one audio classification and three image classification tasks of differing complexities: 1) MNIST (LeCun et al., 1998) on a 4-layer CNN, 2) SpeechCommands (Warden, 2018) on a 4-layer depthwise-separable model (DSCNN) (Zhang et al.,

2017) for keyword spotting task, 3) CIFAR10 (Krizhevsky, 2009) on Resnet-8 (from MLPerf tiny benchmark suite (Banbury et al., 2021)) and 4) TinyImagenet (Le & Yang, 2015) on MobilenetV2 (Howard et al., 2017). While TinyImageNet/MobileNetV2 isn't a typical tinyML dataset/model, we include it to show QUTE's broader applicability.

CID datasets: For CID datasets, we use corrupted versions of ID i.e., MNIST-C (15 corruptions) (Mu & Gilmer, 2019), CIFAR10-C (19 corruptions) and TinyImagenet-C (15 corruptions) (Hendrycks & Dietterich, 2019), with various corruptions such as frost, noise, fog, blur, pixelate etc. For corrupting SpeechCmd, we use the audiomentations library (Jordal, 2024). We apply 11 types of corruptions such as noise, air absorption, time masking/stretching to obtain corrupted-SpeechCmd (denoted SpeechCmd-C).

OOD datasets: SpeechCommands has utterances of 35 words. Like Zhang et al. (2017), we train DSCNN to recognize ten words and use the rest as OOD (denoted SpeechCmd-OOD). Furthermore, we use FashionMNIST (Xiao et al., 2017) and SVHN (Netzer et al., 2011) as OOD for MNIST and CIFAR10 respectively. More details on models and datasets used can be found in Appendix A.1.

Placement of Early-exits The number of early-exits (K) is dictated by the resource budget. There are only a handful locations in a tinyML model where we can insert the early-exits. Our strategy is to insert them at equally-spaced locations along the base network. For MNIST and SpeechCmd, which both use 4-layer models, we insert two EEs after the 1st and 2nd layers. For Resnet-8 with 3 residual stacks, we insert EEs after 1st and 2nd residual stacks. We insert 5 EEs for MobilenetV2 at equally-spaced locations starting from the input. A detailed analysis of the effect of number of early-exits on UQ is provided in Appendix B.1.

Accuracy-drop detection Unlike with OOD data, correct predictions may still occur on CID inputs. Therefore, we focus on detecting accuracy-drop events caused by CID inputs. We present a realistic mechanism with no additional overhead that monitors *confidence* to detect a possible drop in accuracy. We formulate this as a binary classification problem. First, for each baseline we evaluate, we iterate over *only* ID and obtain model predictions while computing the moving average of accuracy of the past m predictions using a sliding-window. The accuracy of the sliding-window is denoted as \mathcal{A}_{SW} and the accuracy distribution thus obtained is denoted as \mathcal{A}_{ID} with mean μ_{ID} and standard-deviation σ_{ID} respectively. In our empirical evaluations, we find that $m = 100$ is reasonable. We then construct corrupted datasets by appending ID with each corrupted-ID version. Next, we iterate over these ID+CID datasets (one for each type of corruption) using a sliding-window while computing the moving average of confidence (\mathcal{C}_{SW}). A *CID-detected* event corresponds to $\mathcal{C}_{SW} < \rho$, and a true positive event occurs when $\mathcal{C}_{SW} < \rho$ and $\mathcal{A}_{SW} \leq \mu_{ID} - 3 \cdot \sigma_{ID}$, where

ρ is a user-defined threshold. We vary ρ from 0 to 1 in steps of 0.1 and repeat the whole experiment. Appendix A.4 provides additional details. Finally, we report the average area under precision-recall curve (AUPRC).

Failure detection Failure detection aims at rejecting both ID misclassifications and OOD instances, as both indicate a potential failure. To this end, we define two binary classification tasks: 1) ID \checkmark | ID \times , which separates correct predictions from incorrect ones within ID/CID samples, and 2) ID \checkmark | OOD, which distinguishes correct predictions from OOD. For the former, we use ID+CID datasets where errors can arise from hard-to-classify ID samples or CID inputs, while for the latter, we use ID and OOD data. In addition, (Xia & Bouganis, 2023) leverages uncertainty estimation for failure detection, optimizing for a user-defined threshold. However, since this threshold may vary *in the field* depending on application requirements, we report the threshold-independent Area Under the Receiver Operating Characteristic curve (AUROC). Finally, (Xia & Bouganis, 2023) shows that Deep Ensembles (DEEP) achieve superior performance on threshold-free metrics, making DEEP the strongest baseline in our comparisons.

Uncertainty quantification For UQ, we report the class-weighted F1 score instead of accuracy because its better for imbalanced datasets, and two proper scoring metrics (Gneiting & Raftery, 2007): Brier score (BS) (Brier, 1950), which measures the accuracy of predicted probabilities and negative log-likelihood (NLL), which measures how close the predictions are to the ground truth (see Appendix C). Also, despite its documented unreliability in measuring uncertainty quality (Nixon et al., 2019), we report expected calibration error (ECE) (Guo et al., 2017) due to its popularity. Appendix C discusses the limitations of ECE.

BASELINES

BASE: Unmodified implementations of models evaluated.

Monte Carlo Dropout (MCD) (Gal & Ghahramani, 2016): computes mean of prediction vectors from K inference passes by activating *dropout*. For fairness, dropout and EE locations are same. We use a dropout rate of 0.1, 0.05, 0.2, 0.2 for MNIST, SpeechCmd, CIFAR10 and Tiny-Imagenet respectively.

Deep Ensembles (DEEP) (Lakshminarayanan et al., 2017): explicit ensemble of models with same base model architecture but with different random weight initializations. The number of ensemble members = K .

Early-exit Ensembles (EE-ensemble) (Qendro et al., 2021): attaches multiple early-exits with additional FC-layers. The size of the FC layers is chosen such that the model delivers the best uncertainty estimation. All early-exit prediction vectors including that of the final exit are averaged to get the final ensemble prediction. We place the

early-exits at the same locations as QUTE. Ablation studies with EE-ensemble are presented in Appendix B.2.

HYDRA (Tran et al., 2020): uses the same architecture as QUTE. Instead of early-exit distillation like QUTE, it employs ensemble distillation, capturing unique predictive behavior of individual ensemble members.

Generalized-ODIN (G-ODIN) (Hsu et al., 2020): a generalized OOD detection method that decomposes the confidence score and introduces a preprocessing layer that perturbs the inputs for better OOD detection. The perturbation magnitude is determined on a small held-out validation set.

6. Results

First, we evaluate our MCU implementations in Section 6.1 on both Small-MCU and Big-MCU to showcase QUTE’s resource efficiency and infeasibility of prior methods on ultra-low-power devices. QUTE achieves an average **27% latency reduction** on Big-MCU compared with the best performing baseline with **58%** and **26% smaller models** compared to EE-ensemble and DEEP respectively. On Small-MCU, high-resource prior methods **cannot even fit** on CIFAR10, further underscoring QUTE’s suitability for constrained environments. Next, we demonstrate QUTE’s superiority over all baselines in detecting accuracy-drop events due to CID in Section 6.2. On failure detection (Section 6.3), QUTE consistently outperforms prior methods in distinguishing between ID \checkmark | ID \times and is competitive on ID \checkmark | OOD with other baselines, even outperforming specialized OOD detectors like G-ODIN on MNIST and SpeechCmd (Section 6.3). Next, we compare QUTE’s calibration in Section 6.4, demonstrating that QUTE predictions are consistently well-calibrated, especially on tiny-sized models with the **59% lower model parameters** on average compared to EE-ensemble. In addition, we conduct an ablation study to assess the effectiveness of EV-assistance method and its effect on model convergence (Section 6.5). Finally, Sections 6.6 and 6.7 demonstrate QUTE’s superiority in uncertainty estimation compared to temperature scaling (TS) (Guo et al., 2017; Rahaman et al., 2021) and single-pass deterministic methods (Charpentier et al., 2020) respectively. Additional ablation studies are presented in Appendix B.

6.1. MCU Fit: QUTE vs Resource-Heavy methods

We evaluate MCU implementations of QUTE on two MCUs of different sizes. Since the primary objective of TinyML systems is to minimize the *energy-per-prediction* (Ghanathe & Wilton, 2023), our target platform is the Small-MCU, as the base network fits comfortably on this smaller device.

We compare QUTE against the two most relevant baselines: EE-ensemble and DEEP. We exclude MCD due to its re-

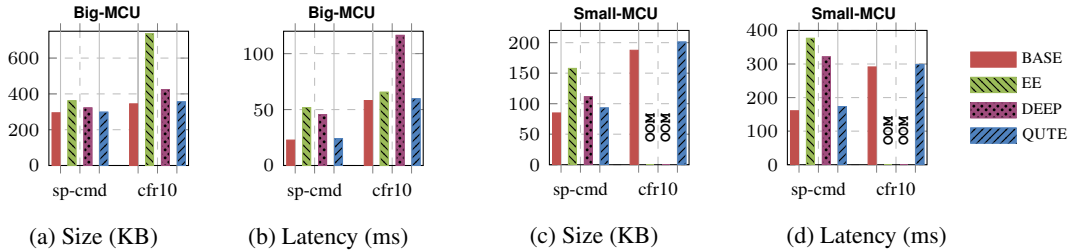


Figure 3: Microcontroller results for SpeechCmd (sp-cmd), CIFAR10 (cfr10) on Big-MCU and Small-MCU (*lower is better*). EE-ensemble and DEEP on CIFAR10 *do not fit* i.e., out-of-memory (OOM)

Model	1stack	2stack	3stack	4stack	QUTE
Fog	30%	24%	18%	11%	25%
Gaussian Noise	7%	-1.2%	12%	13%	20%

Table 2: Percentage reduction in average confidence for Fog (high epistemic) and Gaussian Noise (high aleatoric) corruptions in Resnet{1,2,3,4} and Resnet- QUTE.

liance on a specialized dropout module (Ahmed et al., 2023), which is impractical on MCUs. We also exclude HYDRA because it is often suboptimal, as shown in the following sections. Our evaluation uses CIFAR10 and SpeechCommands, as TinyImageNet/MobileNetV2 exceeds tinyML device limits. On Big-MCU, QUTE achieves **31% and 47% latency reductions** over EE-ensemble and DEEP, respectively, and maintains accuracy parity with both, even with **58% and 26% smaller models**. On Small-MCU, both EE-ensemble and DEEP *do not fit* on the device for CIFAR10; EE-ensemble model size exceeds Small-MCU’s on-device memory by 270KB, whereas DEEP exceeds by 1.9KB. Notably, EE-ensemble has the highest peak RAM usage than DEEP and QUTE due to its need to store large intermediate feature maps during early-exit computation. In contrast, DEEP processes each ensemble member sequentially as there is no parallel compute capability on MCUs. This characteristic allows single-forward-pass methods like QUTE and EE-ensemble to achieve reduced latency, especially as the ensemble size increases. Appendix B.3 delves into more details. These results underscore that the novelty of QUTE lies in its exceptional resource efficiency, making it a practical solution for deploying UQ or model monitoring mechanisms on ultra-low-power MCUs, especially in scenarios where other approaches may be infeasible. This is crucial for the extreme edge because sometimes the alternative is no UQ/monitoring.

6.2. Accuracy-drop detection

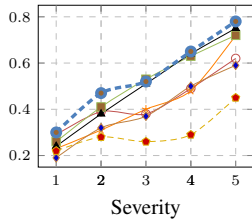
As described in Section 5, we monitor for a drop in confidence to detect accuracy-drop events in a model due to CID. We also explore alternative uncertainty measures like predictive entropy. However, we choose confidence be-

cause it produces similar results without extra processing, a practical advantage for tinyML. Figure 4 reports AUPRC averaged over all corruptions for all datasets. QUTE consistently outperforms all baselines in detecting accuracy-drops caused due to CID across all datasets, which is crucial for maintaining model reliability in deployment. Our empirical evaluations show that early-exits (EE) are more adept at capturing epistemic uncertainty whereas deeper layers are better at capturing aleatoric uncertainty, as illustrated by Table 2. Table 2 reports the percentage reduction in average prediction confidence of models of different sizes on fog (epistemic) and gaussian noise (aleatoric) corruptions compared to the average confidence on CIFAR10-ID. By combining knowledge from EE at the final exits, QUTE exhibits reduced confidence on both types of corruptions, thereby enhancing its accuracy-drop detection capabilities.

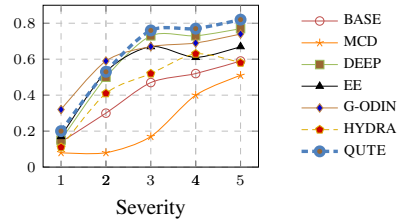
MCD performs poorly in all cases either due to reduced model capacity (due to dropout) or due to extreme calibration. For example, on TinyImagenet-ID, MCD with an accuracy of 0.34 has a median confidence value of 0.345 across all predictions. However, on TinyImagenet-CID, the overparameterized nature of modern neural networks does not allow the median confidence of MCD to drop any further than 0.27 i.e., a drop of 0.075, which is not discernible enough to detect drop in accuracy. This is observed with EE-ensemble too for larger model sizes. In contrast, QUTE’s median confidence on TinyImagenet-ID is 0.69 and the average median confidence on CID is 0.57, which helps create a much clearer distinction between ID and CID. EE-ensemble performs comparably to QUTE for tiny/medium-sized models. However, for large models, the additional learning layers added to balance the learning capacities of all exits cause the early-exits to function like conventional deep networks. Therefore, EE-ensemble becomes overconfident for severe corruptions, which results in a 8.9% drop in performance on TinyImagenet-C. In contrast, the early-view-assistance method allows QUTE to retain early-exit knowledge in its output blocks across different model sizes, enabling it to outperform all baselines on CID detection on all datasets/models. The performance of QUTE and EE-ensemble (especially on tiny-sized models), highlight

AUPRC (\uparrow)	MNIST-C	SpeechCmd-C
- BASE	0.54	0.52
- MCD	0.45	0.56
- DEEP	0.55	0.57
- EE-Ensemble	0.63	0.58
- G-ODIN	0.48	0.37
- HYDRA	0.53	0.51
- QUTE	0.63	0.61

(a) AUPRC MNIST-C & SpeechCmd-C



(b) AUPRC CIFAR10-C



(c) AUPRC TinyImagenet-C

Figure 4: Accuracy-drop detection results. AUPRC is reported for all evaluated baselines (*higher is better*)

AUROC	ID \checkmark ID \times			ID \checkmark OOD		
	Mnist	SpCmd	cfr10	Mnist	SpCmd	cfr10
-BASE	0.75	0.90	0.84	0.07	0.9	0.88
-MCD	0.74	0.89	0.87	0.48	0.89	0.89
-DEEP	0.85	0.91	0.86	0.78	0.91	<u>0.92</u>
-EE-ensemble	0.85	0.90	0.85	0.85	0.90	0.90
-G-ODIN	0.72	0.74	0.83	0.4	0.74	0.95
-HYDRA	0.81	0.90	0.83	0.71	0.9	0.90
-QUTE	0.87	0.91	<u>0.86</u>	<u>0.84</u>	0.91	0.91

Table 3: Failure detection experiments

the utility of early-exit knowledge in detecting corruptions. Similarly, G-ODIN remains overconfident on CID, and is surpassed by most baselines on all corrupted datasets. This suggests that general OOD-detectors are not automatically effective for CID detection, consistent with previous findings that OOD detectors find it hard to detect non-semantic shifts (Hsu et al., 2020). Further, HYDRA is consistently subpar at accuracy-drop detection, often being outperformed even by BASE. This highlights HYDRA’s need for larger classification heads with more parameters to fully integrate ensemble-knowledge, consistent with the original paper’s methodology. Contrarily, QUTE’s early-exit distillation method proves effective in detecting corruptions compared to ensemble distillation mechanisms.

6.3. Failure detection

Table 3 reports the AUROC for failure detection experiments for ID \checkmark | ID \times and ID \checkmark | OOD. As seen, for ID \checkmark | ID \times , QUTE outperforms all baselines on MNIST and SpeechCmd. MCD performs slightly better on CIFAR10, likely due to the higher number of epistemic uncertainty-inducing corruptions present in CIFAR10-C, such as blur and digital distortions. This observation aligns with prior work highlighting MCD’s strength in capturing instance-level epistemic uncertainties (Kendall & Gal, 2017). In contrast, QUTE effectively captures both epistemic and aleatoric uncertainties, resulting in overall superior performance across datasets. For ID \checkmark | OOD, QUTE outperforms all baselines on SpeechCmd, and is a close second on MNIST. Surprisingly, G-ODIN, a specialized OOD detector, is poor on MNIST and SpeechCmd, even outperformed by BASE. Further analysis revealed that G-ODIN is under-

confident on these ID-datasets and over-confident on OOD. As an ablation study, we increased the number of training epochs of G-ODIN by 30 epoch, which resulted a 55% improvement on ID \checkmark | OOD on MNIST. Contrarily, on a larger model, Resnet-8 on CIFAR10, G-ODIN’s performance notably improves on ID \checkmark | OOD, consistent with previous findings. This may suggest the need to rethink OOD detection with extremely tiny models. These results showcase QUTE’s versatility/efficacy in detecting both CID and OOD data.

6.4. Uncertainty quantification

Table 4 reports the calibration metrics for all datasets we evaluate. QUTE *outperforms* all baselines on tiny-sized models i.e., MNIST and SpeechCmd, and is *competitive* for medium and large models i.e., CIFAR10 and TinyImagenet despite using significantly lower resources. All methods perform better than BASE, except MCD on MNIST and SpeechCmd. Further inquiry revealed that applying dropout during inference on the extremely tiny models takes away from its model capacity thereby leading to drop in accuracy and poor calibration. DEEP showcases good calibration on all datasets owing to its larger model capacity. However, both MCD and DEEP are impractical for tinyML because they both require multiple inference passes or specialized hardware. In contrast, QUTE provides comparable uncertainty estimates in a single forward pass with 53% lower model sizes on average compared to DEEP.

QUTE also outperforms EE-ensemble, the most relevant prior work, across all datasets with models that are, on average, 59% smaller, except on larger datasets and models like TinyImagenet on MobileNetV2. Since QUTE’s primary focus is resource efficiency, the limited parameters in its output blocks—restricted to a single depthwise convolutional (CONV) layer—can constrain calibration performance relative to high-resource methods like DEEP and EE-ensemble. To normalize the comparison, we enhanced the QUTE architecture with additional learning layers at each output block, creating QUTE+. For Resnet-8 on CIFAR10, we added two depthwise-separable CONV layers, while for MobileNetV2 on TinyImageNet, we included an additional

Model	F1 (\uparrow)	BS (\downarrow)	NLL (\downarrow)	ECE (\downarrow)
MNIST				
-BASE	0.910 \pm 0.002	0.013 \pm 0.000	0.292 \pm 0.006	0.014 \pm 0.001
-MCD	0.886 \pm 0.004	0.018 \pm 0.000	0.382 \pm 0.004	0.071 \pm 0.006
-DEEP	0.931 \pm 0.005	0.010 \pm 0.000	0.227 \pm 0.002	0.034 \pm 0.004
-EE-ensemble	0.939 \pm 0.002	0.011 \pm 0.000	0.266 \pm 0.005	0.108 \pm 0.002
-HYDRA	0.932 \pm 0.006	0.010 \pm 0.000	0.230 \pm 0.012	0.014\pm0.005
-QUTE	0.941\pm0.004	0.009\pm0.000	0.199\pm0.010	0.026 \pm 0.003
SpeechCmd				
-BASE	0.923 \pm 0.007	0.010 \pm 0.000	0.233 \pm 0.016	0.026 \pm 0.001
-MCD	0.917 \pm 0.006	0.011 \pm 0.000	0.279 \pm 0.013	0.048 \pm 0.002
-DEEP	0.934\pm0.008	0.008 \pm 0.000	0.205 \pm 0.012	0.034 \pm 0.006
-EE-ensemble	0.926 \pm 0.002	0.009 \pm 0.000	0.226 \pm 0.009	0.029 \pm 0.001
-HYDRA	0.932 \pm 0.005	0.008 \pm 0.000	0.203 \pm 0.016	0.018 \pm 0.004
-QUTE	0.933 \pm 0.006	0.008\pm0.000	0.202\pm0.016	0.018\pm0.001
CIFAR10				
-BASE	0.834 \pm 0.005	0.023 \pm 0.000	0.523 \pm 0.016	0.049 \pm 0.003
-MCD	0.867 \pm 0.002	0.019 \pm 0.000	0.396 \pm 0.003	0.017 \pm 0.005
-DEEP	0.877 \pm 0.003	0.017 \pm 0.000	0.365\pm0.015	0.015\pm0.003
-EE-ensemble	0.854 \pm 0.001	0.021 \pm 0.000	0.446 \pm 0.011	0.033 \pm 0.001
-HYDRA	0.818 \pm 0.004	0.026 \pm 0.000	0.632 \pm 0.017	0.069 \pm 0.001
-QUTE	0.858 \pm 0.001	0.020 \pm 0.000	0.428 \pm 0.019	0.025 \pm 0.003
-QUTE +	0.878\pm0.003	0.017\pm0.000	0.369 \pm 0.008	0.026 \pm 0.001
TinyImagenet				
-BASE	0.351 \pm 0.005	0.004 \pm 0.000	5.337 \pm 0.084	0.416 \pm 0.003
-MCD	0.332 \pm 0.004	0.003 \pm 0.000	2.844 \pm 0.028	0.061 \pm 0.005
-DEEP	0.414 \pm 0.006	0.003 \pm 0.000	3.440 \pm 0.049	0.115 \pm 0.003
-EE-ensemble	0.430\pm0.005	0.003 \pm 0.000	2.534\pm0.046	0.032\pm0.006
-HYDRA	0.376 \pm 0.004	0.004 \pm 0.000	3.964 \pm 0.036	0.328 \pm 0.004
-QUTE	0.395 \pm 0.014	0.004 \pm 0.000	3.700 \pm 0.123	0.282 \pm 0.009
-QUTE +	0.381 \pm 0.010	0.003\pm0.000	<u>2.757\pm0.044</u>	0.122 \pm 0.008

Table 4: Calibration Metrics for all baselines on ID data. For F1, higher is better, and for BS and NLL, lower is better. The best results are marked in bold. All results are mean \pm std-dev for 3 independent splits of ID test data.

dense (fully connected) layer to form QUTE +. As seen in Table 4, F1 score of QUTE + exceeds that of DEEP on CIFAR10, however, it slightly degrades on TinyImagenet. Further analysis revealed that since the base network is trained simultaneously with all early-exits (EE) and EV-exits, the weights of the shared base network is negatively impacted since the training routine tries to optimize for all exits. In contrast, EE-ensemble method employs only half as many exits (only EEs) compared to QUTE, and the EE’s regularization effect helps network accuracy (Teerapittayanon et al., 2016). Appendix B.1 studies the effect of ensemble size on QUTE’s calibration performance. Nevertheless, QUTE + significantly improves calibration, the primary focus of this work, outperforming all baselines on BS and achieving NLL on par with the top-performing method. These results indicate that, in an unconstrained setting, QUTE can achieve calibration performance comparable to high-resource methods. Additionally, we assess the effectiveness of QUTE’s early-exit distillation technique by comparing against HYDRA, which utilizes ensemble-distillation. While HYDRA’s lightweight classification heads (similar to QUTE) capture ensemble knowledge in tiny models, they struggle as model size increases, leading to reduced accuracy and poor calibration, even underperforming BASE on CIFAR10. This reveals HYDRA’s reliance on larger classification heads for effective ensemble integration as discussed before. Contrarily, QUTE’s early-exit distillation method consistently proves effective across diverse models and datasets.

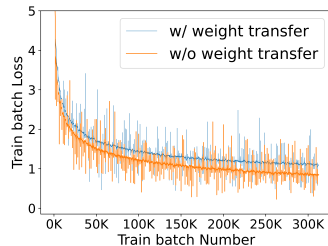


Figure 5: Batch-loss at EV-0 for CIFAR10 on Resnet-8

6.5. Effectiveness of EV-assistance method and its effect on convergence

To investigate the effect of the early-view (EV) assistance method on ensemble quality, we train non-EV versions of QUTE i.e., the same architecture (with early-exits) except without the weight transfer mechanism. Table 5 reports the calibration metrics for both EV and non-EV versions for CIFAR10 and TinyImagenet. As shown, EV-assistance has a massive influence on improving calibration, and improves the network accuracy on TinyImagenet by 8%. Also, we find that the ability of non-EV versions to detect accuracy-drop events reduces by 15% on average compared to EV-versions. These results show that the notably lower NLL obtained with EV-assistance (5% lower on CIFAR10 and 19% lower on TinyImagenet) not only improves uncertainty estimates, but it also proves crucial in both CID and OOD detection. In addition, we investigate the effect of weight transfer on model convergence. Figure 5 plots the loss after each trainbatch at one of the EV-blocks of Resnet-8 on CIFAR10. As seen, model convergence is not affected because the weight transfer mechanism is designed such that the loss at EV-exit is always computed w.r.t. to the copied weights (see Section 4). In essence, EV-assistance uses the EE-weights to extract and reemphasize the EE features at the final exit(s).

6.6. Comparison with Temperature scaling

Temperature scaling (TS) (Guo et al., 2017) is a simple post-training calibration technique that applies a linear scaling factor to logits (pre-softmax) to mitigate model overconfidence. Table 6 and Table 7 present the calibration and accuracy-drop detection results, respectively, for TS versions of BASE (BASE-TS) and QUTE (QUTE-TS). For QUTE-TS, the same architecture as QUTE is used, but without weight transfer. We employed the best-performing *pool-then-calibrate* configuration from (Rahaman et al., 2021), which applies temperature scaling to ensembles to obtain QUTE-TS. As the results show, QUTE consistently outperforms BASE-TS on both calibration and accuracy-drop detection. However, despite improving calibration in some scenarios, QUTE-TS performs worse than QUTE on accuracy-drop detection, particularly with tiny datasets such as MNIST-C and SpeechCmd-C. A deeper analysis revealed

Model	F1 (\uparrow)		BS (\downarrow)		NLL (\downarrow)	
	EV	no-EV	EV	no-EV	EV	no-EV
CIFAR10	0.858 \pm 0.001	0.858 \pm 0.000	0.0205\pm0.000	0.0206 \pm 0.000	0.435\pm0.007	0.459 \pm 0.006
TinyImagenet	0.380\pm0.011	0.35 \pm 0.004	0.0043\pm7.2E-5	0.0046 \pm 5.7E-5	3.813\pm0.105	4.743 \pm 0.111

Table 5: Effectiveness of *early-view* assistance. Results averaged over 3 independent training runs.

Model	F1 (\uparrow)	BS (\downarrow)	NLL (\downarrow)
MNIST			
- BASE-TS	0.937 \pm 0.002	0.009 \pm 0.000	0.203 \pm 0.005
- QUTE-TS	0.942\pm0.001	0.008\pm0.000	0.191\pm0.005
- QUTE	0.941 \pm 0.004	0.009 \pm 0.000	0.199 \pm 0.010
SpeechCmd			
- BASE-TS	0.923 \pm 0.007	0.009 \pm 0.000	0.229 \pm 0.017
- QUTE-TS	0.926 \pm 0.006	0.009 \pm 0.000	0.233 \pm 0.010
- QUTE	0.933\pm0.006	0.008\pm0.000	0.202\pm0.016
CIFAR10			
- BASE-TS	0.834 \pm 0.000	0.023 \pm 0.000	0.493 \pm 0.000
- QUTE-TS	0.852 \pm 0.002	0.021 \pm 0.000	0.444 \pm 0.013
- QUTE	0.858\pm0.001	0.020\pm0.000	0.428\pm0.019

Table 6: Calibration with temperature scaling (TS) on ID

AUPRC (\uparrow)	MNIST-C	SpeechCmd-C	CIFAR10-C (Severity 1 to 5)
- BASE-TS	0.47	0.52	0.30, 0.42, 0.38, 0.50, 0.61
- QUTE-TS	0.52	0.53	0.26, 0.45, 0.51, 0.65, 0.76
- QUTE	0.63	0.61	0.30, 0.47, 0.52, 0.65, 0.78

Table 7: Accuracy-drop detection with temperature-scaling

that QUTE-TS tends to be overconfident for certain corruptions. For instance, under fog corruption on MNIST-C, QUTE-TS exhibits an overconfident behavior with a median confidence of 0.99, while QUTE maintains a much lower confidence of 0.51, leading to superior accuracy-drop detection capability.

6.7. Comparison with single-pass deterministic methods

Postnets (PostN) (Charpentier et al., 2020) is the most relevant single-pass deterministic method compared to QUTE, which uses normalizing flows to predict posterior distribution of any input sample with no additional memory overhead. We evaluate PostN on MNIST and CIFAR10 by substituting the encoder network architectures with the ones used in our evaluations. We use identical hyperparameter settings of (Charpentier et al., 2020) and train for same number of epochs as in our evaluations but with early stopping. Table 8 reports the calibration metrics for PostN and QUTE. Since PostN don’t necessarily focus on accuracy, we found that MNIST-PostN needed a 50% increase in number of epochs to even surpass the accuracy of MNIST-BASE. We report this result. For CIFAR10, we found that early stopping halted the training, and report those results. As seen, QUTE consistently outperforms PostN on all metrics on both datasets, making QUTE a more efficient choice for resource-constrained environments.

7. Conclusion and Discussion

Uncertainty quantification is essential for model monitoring in tinyML, yet many prior works fail to address the

Model	F1 (\uparrow)	BS (\downarrow)	NLL (\downarrow)
MNIST-PostN	0.92	0.012	0.286
MNIST-QUTE	0.94	0.009	0.199
CIFAR-PostN	0.84	0.022	0.462
CIFAR-QUTE	0.858	0.020	0.428

Table 8: Calibration comparison with Postnets on ID

resource constraints inherent in this domain. We propose a novel resource-efficient ensemble architecture for uncertainty estimation that is optimized for tinyML. At the core of our methodology is our finding that model overconfidence decreases with its size. We leverage the better uncertainty estimation quality of early-exits (EE) by injecting EE weights into the multiple lightweight classification heads created at the output. QUTE provides reliable uncertainty estimates in a single forward pass with **59% smaller models** on average and a **31% average reduction in latency** on low-power MCUs. Furthermore, QUTE effectively distills EE knowledge into the final exits, enabling it to capture both aleatoric and epistemic uncertainties more effectively. This capability leads to superior performance in detecting accuracy-drop events (corruptions) across a diverse range of datasets and model complexities. QUTE serves as an excellent and cost-effective *accuracy-monitoring* mechanism for field-deployed models.

Limitations/Discussions The ensemble size (K) is restricted by the depth of the base network, which may not provide uncertainty estimates of sufficient quality for a given safety-critical application. In addition, we find that increasing K beyond a certain point negatively influences the shared base network, leading to degradation in calibration (ablation study in Appendix B.1). Furthermore, while QUTE achieves the best calibration on tiny-sized models, it requires additional learning layers on large models and datasets. Alternatively, methods that improve early-exit calibration (Jazbec et al., 2024) could potentially improve effectiveness of QUTE’s EV-assistance. Interestingly, we observe that excessively increasing EV-assistance, such as by raising the depth multiplier in QUTE’s output blocks, negatively impacts accuracy despite improving calibration. However, its worth noting that the training times with QUTE are higher compared to other methods (except HYDRA) due to weight transfer between train-batches. Finally, despite QUTE’s versatility/effectiveness in OOD detection, specialized OOD detectors should be considered, especially for large models.

References

- Ahmed, S. T., Danouchi, K., Münch, C., Prenat, G., Anghel, L., and Tahoori, M. B. Spindrop: Dropout-based bayesian binary neural networks with spintronic implementation. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 13(1):150–164, 2023. doi: 10.1109/JETCAS.2023.3242146.
- Ahmed, S. T., Hefenbrock, M., and Tahoori, M. B. Tiny deep ensemble: Uncertainty estimation in edge ai accelerators via ensembling normalization layers with shared weights, 2024. URL <https://arxiv.org/abs/2405.05286>.
- Antoran, J., Allingham, J., and Hernández-Lobato, J. M. Depth uncertainty in neural networks. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 10620–10634. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/781877bda0783aac5f1cf765c128b437-Paper.pdf.
- Bai, H., Canal, G., Du, X., Kwon, J., Nowak, R., and Li, Y. Feed two birds with one scone: Exploiting wild data for both out-of-distribution generalization and detection, 2023.
- Baier, L., Schlör, T., Schöffner, J., and Kühl, N. Detecting concept drift with neural network model uncertainty, 2022.
- Banbury, C. R., Reddi, V. J., Torelli, P., Holleman, J., Jeffries, N., Király, C., Montino, P., Kanter, D., Ahmed, S., Pau, D., Thakker, U., Torrini, A., Warden, P., Cordaro, J., Guglielmo, G. D., Duarte, J. M., Gibellini, S., Parekh, V., Tran, H., Tran, N., Niu, W., and Xu, X. Mlperf tiny benchmark. *CoRR*, abs/2106.07597, 2021. URL <https://arxiv.org/abs/2106.07597>.
- Bifet, A. and Gavaldà, R. Learning from time-changing data with adaptive windowing. In *Proceedings of the 2007 SIAM international conference on data mining*, pp. 443–448. SIAM, 2007.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight uncertainty in neural network. In *International conference on machine learning*, pp. 1613–1622. PMLR, 2015.
- Bonato, V. and Bouganis, C.-S. Class-specific early exit design methodology for convolutional neural networks. *Applied Soft Computing*, 107:107316, 2021. ISSN 1568-4946. doi: <https://doi.org/10.1016/j.asoc.2021.107316>. URL <https://www.sciencedirect.com/science/article/pii/S1568494621002398>.
- Brier, G. W. Verification of forecasts expressed in terms of probability. *Monthly weather review*, 78(1):1–3, 1950.
- Charpentier, B., Zügner, D., and Günnemann, S. Posterior network: Uncertainty estimation without ood samples via density-based pseudo-counts. *Advances in neural information processing systems*, 33:1356–1367, 2020.
- Deng, D., Chen, G., Yu, Y., Liu, F., and Heng, P.-A. Uncertainty estimation by fisher information-based evidential deep learning. In *International Conference on Machine Learning*, pp. 7596–7616. PMLR, 2023.
- Dinari, O. and Freifeld, O. Variational- and metric-based deep latent space for out-of-distribution detection. In Cussens, J. and Zhang, K. (eds.), *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence*, volume 180 of *Proceedings of Machine Learning Research*, pp. 569–578. PMLR, 01–05 Aug 2022. URL <https://proceedings.mlr.press/v180/dinari22a.html>.
- Ferianc, M. and Rodrigues, M. Mimmo: Multi-input massive multi-output neural network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pp. 4563–4568, June 2023.
- Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059. PMLR, 2016.
- Gama, J., Medas, P., Castillo, G., and Rodrigues, P. Learning with drift detection. In *Advances in Artificial Intelligence—SBIA 2004: 17th Brazilian Symposium on Artificial Intelligence, Sao Luis, Maranhao, Brazil, September 29–October 1, 2004. Proceedings 17*, pp. 286–295. Springer, 2004.
- Ghanathe, N. P. and Wilton, S. T-recx: Tiny-resource efficient convolutional neural networks with early-exit. In *Proceedings of the 20th ACM International Conference on Computing Frontiers*, pp. 123–133, 2023.
- Gneiting, T. and Raftery, A. E. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. In *International conference on machine learning*, pp. 1321–1330. PMLR, 2017.
- Havasi, M., Jenatton, R., Fort, S., Liu, J. Z., Snoek, J., Lakshminarayanan, B., Dai, A. M., and Tran, D. Training independent subnetworks for robust prediction, 2021.

- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016. doi: 10.1109/CVPR.2016.90.
- Hendrycks, D. and Dietterich, T. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019.
- Hendrycks, D. and Gimpel, K. A baseline for detecting misclassified and out-of-distribution examples in neural networks, 2018.
- Hernández-Lobato, J. M. and Adams, R. P. Probabilistic backpropagation for scalable learning of bayesian neural networks, 2015.
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., and Adam, H. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- Hsu, Y.-C., Shen, Y., Jin, H., and Kira, Z. Generalized odin: Detecting out-of-distribution image without learning from out-of-distribution data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10951–10960, 2020.
- Huang, G., Chen, D., Li, T., Wu, F., van der Maaten, L., and Weinberger, K. Q. Multi-Scale Dense Networks for Resource Efficient Image Classification. *arXiv e-prints*, art. arXiv:1703.09844, March 2017.
- Jaeger, P. F., Lüth, C. T., Klein, L., and Bungert, T. J. A call to reflect on evaluation practices for failure detection in image classification, 2023. URL <https://arxiv.org/abs/2211.15259>.
- Jazbec, M., Allingham, J., Zhang, D., and Nalisnick, E. Towards anytime classification in early-exit architectures by enforcing conditional monotonicity. *Advances in Neural Information Processing Systems*, 36, 2024.
- Jordal, I. Audiomentations, 2024. URL <https://github.com/iver56/audiomentations>. Version 0.32.0.
- Katz-Samuels, J., Nakhleh, J., Nowak, R., and Li, Y. Training ood detectors in their natural habitats, 2022.
- Kaya, Y., Hong, S., and Dumitras, T. Shallow-deep networks: Understanding and mitigating network overthinking. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 3301–3310. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/kaya19a.html>.
- Kendall, A. and Gal, Y. What uncertainties do we need in bayesian deep learning for computer vision? In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/2650d6089a6d640c5e85b2b88265dc2b-Paper.pdf.
- Krizhevsky, A. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, pp. 6405–6416, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- Le, Y. and Yang, X. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Liang, S., Li, Y., and Srikant, R. Enhancing the reliability of out-of-distribution image detection in neural networks, 2020.
- Liu, W., Wang, X., Owens, J. D., and Li, Y. Energy-based out-of-distribution detection, 2021.
- Malinin, A. and Gales, M. Predictive uncertainty estimation via prior networks. *Advances in neural information processing systems*, 31, 2018.
- Malinin, A., Mlodozieniec, B., and Gales, M. Ensemble distribution distillation. *arXiv preprint arXiv:1905.00076*, 2019.
- Meronen, L., Trapp, M., Pilzer, A., Yang, L., and Solin, A. Fixing overconfidence in dynamic neural networks, 2023. URL <https://arxiv.org/abs/2302.06359>.
- Mu, N. and Gilmer, J. Mnist-c: A robustness benchmark for computer vision, 2019.
- Mukhoti, J., Kirsch, A., van Amersfoort, J., Torr, P. H., and Gal, Y. Deep deterministic uncertainty: A new simple baseline. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 24384–24394, 2023.
- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. 2011.

- Nixon, J., Dusenberry, M. W., Zhang, L., Jerfel, G., and Tran, D. Measuring calibration in deep learning. In *CVPR workshops*, volume 2, 2019.
- Ovadia, Y., Fertig, E., Ren, J., Nado, Z., Sculley, D., Nowozin, S., Dillon, J. V., Lakshminarayanan, B., and Snoek, J. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift, 2019.
- Qendro, L., Campbell, A., Lio, P., and Mascolo, C. Early exit ensembles for uncertainty quantification. In *Machine Learning for Health*, pp. 181–195. PMLR, 2021.
- Rahaman, R. et al. Uncertainty quantification and deep ensembles. *Advances in Neural Information Processing Systems*, 34:20063–20075, 2021.
- Sensoy, M., Kaplan, L., and Kandemir, M. Evidential deep learning to quantify classification uncertainty. *Advances in neural information processing systems*, 31, 2018.
- STMicroelectronics. *STM32L432KC Datasheet: Ultra-low-power Arm Cortex-M4 32-bit MCU+FPU, 100DMIPS, up to 256KB Flash, 64KB SRAM, USB FS, analog, audio*, 2018. URL <https://www.st.com/resource/en/datasheet/stm32l432kc.pdf>.
- STMicroelectronics. *STM32F767ZI Datasheet: ARM Cortex-M7 Microcontroller with 512 KB Flash, 216 MHz CPU, ART Accelerator, FPU, and Chrom-ART Accelerator*, 2019. URL <https://www.st.com/resource/en/datasheet/stm32f767zi.pdf>.
- Tang, X., Yang, K., Wang, H., Wu, J., Qin, Y., Yu, W., and Cao, D. Prediction-uncertainty-aware decision-making for autonomous vehicles. *IEEE Transactions on Intelligent Vehicles*, 7(4):849–862, 2022. doi: 10.1109/TIV.2022.3188662.
- Teerapittayanon, S., McDanel, B., and Kung, H. Branchynet: Fast inference via early exiting from deep neural networks. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pp. 2464–2469, 2016. doi: 10.1109/ICPR.2016.7900006.
- Teye, M., Azizpour, H., and Smith, K. Bayesian uncertainty estimation for batch normalized deep networks, 2018.
- Tran, L., Veeling, B. S., Roth, K., Swiatkowski, J., Dillon, J. V., Snoek, J., Mandt, S., Salimans, T., Nowozin, S., and Jenatton, R. Hydra: Preserving ensemble diversity for model distillation. *arXiv preprint arXiv:2001.04694*, 2020.
- Van Amersfoort, J., Smith, L., Teh, Y. W., and Gal, Y. Uncertainty estimation using a single deep deterministic neural network. In *International conference on machine learning*, pp. 9690–9700. PMLR, 2020.
- Vargas, J., Alsweiss, S., Toker, O., Razdan, R., and Santos, J. An overview of autonomous vehicles sensors and their vulnerability to weather conditions. *Sensors*, 21(16):5397, 2021.
- Warden, P. Speech commands: A dataset for limited-vocabulary speech recognition. *CoRR*, abs/1804.03209, 2018. URL <http://arxiv.org/abs/1804.03209>.
- Wenzel, F., Snoek, J., Tran, D., and Jenatton, R. Hyperparameter ensembles for robustness and uncertainty quantification. *Advances in Neural Information Processing Systems*, 33:6514–6527, 2020.
- Xia, G. and Bouganis, C.-S. Augmenting softmax information for selective classification with out-of-distribution data, 2023. URL <https://arxiv.org/abs/2207.07506>.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- Yang, J., Zhou, K., Li, Y., and Liu, Z. Generalized out-of-distribution detection: A survey. *arXiv preprint arXiv:2110.11334*, 2021.
- Yang, J., Wang, P., Zou, D., Zhou, Z., Ding, K., Peng, W., Wang, H., Chen, G., Li, B., Sun, Y., Du, X., Zhou, K., Zhang, W., Hendrycks, D., Li, Y., and Liu, Z. Openood: Benchmarking generalized out-of-distribution detection, 2022.
- Zaidi, S. and Zela, A. Neural ensemble search for performant and calibrated predictions. 2020.
- Zhang, X.-Y., Xie, G.-S., Li, X., Mei, T., and Liu, C.-L. A survey on learning to reject. *Proceedings of the IEEE*, 111(2):185–215, 2023. doi: 10.1109/JPROC.2023.3238024.
- Zhang, Y., Suda, N., Lai, L., and Chandra, V. Hello edge: Keyword spotting on microcontrollers. *CoRR*, abs/1711.07128, 2017. URL <http://arxiv.org/abs/1711.07128>.

A. Appendix

A.1. Training and dataset details

In this section, we describe the details of the models evaluated and specifics of training and the experimental setup.

A.1.1. DATASETS

In our evaluations, we use four in-distribution datasets for training all baselines methods we evaluate: 1) MNIST (LeCun et al., 1998), 2) SpeechCommands (Warden, 2018), 3) CIFAR10 (Krizhevsky, 2009) and 4) TinyImagenet (Le & Yang, 2015).

MNIST MNIST is a dataset of handwritten digits containing 60,000 grayscale images of size 28×28 and a test set of 10,000 images. MNIST contains 10 classes

SpeechCommands Speech-Cmd is a collection of short audio clips, each spanning 1 second. The dataset consists of utterances for 35 words and is commonly used for benchmarking keyword spotting systems. We train our model (DSCNN) to recognize ten words out of 35: *Down, Go, Left, No, Off, On, Right, Stop, Up, Yes*. Thus, the number of classes is 10. The audio files in WAV format are preprocessed to compute Mel-frequency cepstral coefficients (mel-spectograms). The mel-spectograms are of size 49×10 with a single channel. The DSCNN model is trained for 10 epochs with a batch size of 100.

CIFAR10 CIFAR10 dataset consists of 60,000 32×32 rgb images out of which 10,000 images are in the test set. It contains 10 classes and thus 6000 images per class.

TinyImagenet TinyImagenet is a smaller version of the Imagenet dataset containing 200 classes instead of 1000 classes of the original Imagenet. Each class in TinyImagenet has 500 images in the train set and the validation set contains 50 images per class. The size of the images are resized and fixed at $64 \times 64 \times 3$.

For image classification tasks, we use the following corrupted versions of ID: 1) MNIST-C (Mu & Gilmer, 2019), 2) CIFAR10-C and 3) TinyImagenet-C (Hendrycks & Dietterich, 2019). All corruptions are drawn from 4 major sources: noise, blur, weather and digital. The *digital* and *blur* corruptions induce high epistemic uncertainty, whereas, *noise* corruptions induce higher aleatoric uncertainty. In contrast, *weather* corruptions tend to encompass both epistemic and aleatoric components. Most corruptions are systematic transformations of input images, which predominantly introduce epistemic uncertainty. Pure noise corruptions are the primary source of aleatoric uncertainty; however, many corruptions exhibit both components, often leaning toward epistemic.

MNIST-C The MNIST-C dataset contains 15 corrupted versions of MNIST - *shot_noise, impulse_noise, glass_blur, fog, spatter, dotted_line, zigzag, canny_edges, motion_blur, shear, scale, rotate, brightness, translate, stripe, identity*. All the corruptions are of a fixed severity level.

CIFAR10-C CIFAR10-C includes 19 different types of corruptions with 5 severity levels which gives us $19 \times 5 = 95$ corrupted versions of CIFAR-10 ID. We use all 95 corrupted versions in our experiments. The list of corruptions are: *gaussian_noise, brightness, contrast, defocus_blur, elastic, fog, frost, frosted_glass_blur, gaussian_blur, impulse_noise, jpeg_compression, motion_blur, pixelate, saturate, shot_noise, snow, spatter, speckle_noise, zoom_blur*.

TinyImagenet-C TinyImagenet-C includes 15 different types of corruptions with 5 severity levels which gives us $15 \times 5 = 75$ corrupted versions of ID. We use all 75 corrupted versions in our experiments. The list of corruptions are: *gaussian_noise, brightness, contrast, defocus_blur, elastic_transform, fog, frost, glass_blur, impulse_noise, jpeg_compression, motion_blur, pixelate, shot_noise, snow, zoom_blur*.

SpeechCmd-C For the audio classification task of Keyword spotting on SpeechCmd, we use corrupt the audio using the audiomentations library (Jordal, 2024) with the following 11 corruptions: *gaussian noise, air absorption, band pass filter, band stop filter, high pass filter, high shelf filter, low pass filter, low shelf filter, peaking filter, tanh distortion, time mask, time stretch*.

A.1.2. TRAINING DETAILS

We evaluate four models in our experiments: 1) 4-layer CNN on MNIST, 2) 4-layer Depthwise-separable CNN (DSCNN) on SpeechCmd, 3) Resnet-8 with 3 residual stacks from the MLPerf Tiny benchmark suite (Banbury et al., 2021) on CIFAR10 and 4) MobilenetV2 (Howard et al., 2017) on TinyImagenet.

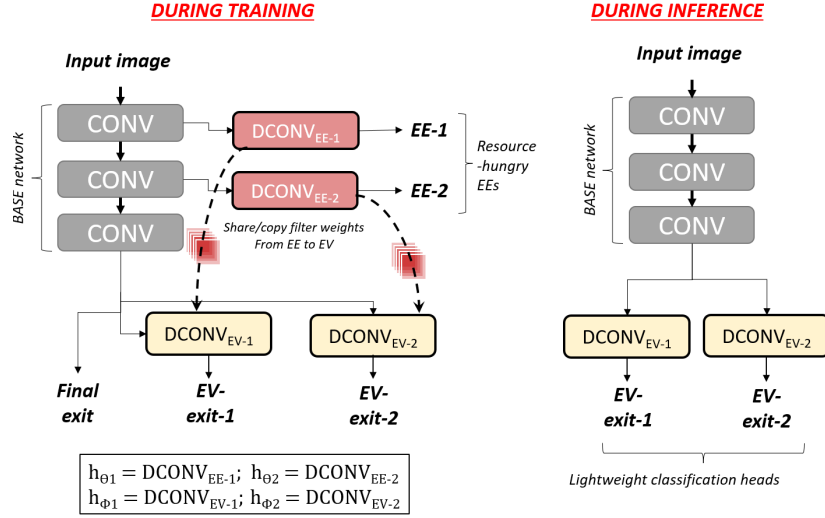


Figure 6: Depiction of QUTE during training and inference. The early-exits (EE) are only used during training for weight-transfer. During inference, the EEs are removed, and only EV-exits are retained. We do not include other layers such as dense and pointwise convolution in the figure for brevity.

The 4-layer CNN is trained for 20 epochs with a batch size of 256, the DSCNN model is trained for 10 epochs with a batch size of 100, the Resnet-8 model is trained for 200 epochs with batch size of 32 and the MobilenetV2 model is trained for 200 epochs with a batch size of 128. All models are trained with Adam optimizer with momentum of 0.9 and an initial learning rate of 0.001, except DSCNN on SpeechCmd which uses an initial learning rate of 0.0005. The learning rate is decayed by a factor of 0.99 every epoch for all image classification tasks. We follow a step function for SpeechCmd that reduces learning rate by half every 2 epochs. For models with QUTE architecture, we achieve the weight transfer from early-exits to all $\{f_{out}(\cdot)\}_{k=1}^K$ using a callback that sets the weights of each f_{out_k} with weights copied from the corresponding assisting early-exit at the beginning of each training batch.

A.2. Weighting the loss at EV-exits

As described in Section 4, we weight the loss at EV-exits with a higher factor w_{EV_k} such that $w_{EV_k} = w_{EV_{k-1}} + \delta$. This is done to further promote diversity across the EV-exits and to vary the influence of each EV-exit on the shared network’s weights. Utilizing distinct weighting factors across EV-exits implies that each EV-exit contributes to the final loss calculation with varying degrees of influence. Consequently, this diversity extends to the weight updates of the shared base network, thereby injecting diversity into the classification heads. We empirically set δ to be 0.5. To determine w_{EV_0} , we repeat training with w_{EV_0} set to $\{2, 3, 4, 5\}$. We find that for $w_{EV_0} > 3$, the NLL starts dropping steadily because the higher loss at EV-exits overshadow the losses of the early-exits. This diminishes the influence of early-exits at EV-exits, which is undesirable. Hence, we set $w_{EV_0} = 3$ in all our experiments.

A.3. Corrupted-in-distribution datasets

We construct the corrupted-in-distribution (CID) datasets for our evaluations as follows. For MNIST-C with a fixed severity level, we append each corrupted version of MNIST to the MNIST-ID dataset to create 15 ID+CID datasets. For CIFAR10-C and TinyImagenet-C, which have 19 and 15 different corruptions respectively with 5 severity levels each, we construct the corrupted datasets as follows. For each type of corruption, we randomly select p samples from each of the 5 severity levels. Next, we concatenate all these samples to create a new corrupted dataset of size $5 \times p$. p is selected such that $5 \times p = \text{size of ID dataset}$. This process is repeated for all corruptions. The datasets thus obtained contain samples from all severity levels. For example, for CIFAR10-C which consists of 19 corruptions, this process yields 19 ID+CID datasets.

A.4. Accuracy-drop detection experiments

For experiments to detect accuracy drop/CID detection described in Section 6.2, we append the ID dataset with each corrupted dataset obtained from the methodology described above. For example, for experiments with CIFAR10-C, we obtain 19 ID+CID datasets that contains both ID samples and corrupted samples (from all severity levels). Next, as described in Section 5, for each prior work we evaluate, we first iterate over *only* ID and obtain predictions from the model while computing the moving average of accuracy of the past m predictions using a sliding-window. In this way, we obtain the accuracy distribution of the sliding-window \mathcal{A}_{SW} on ID, and then compute its mean μ_{ID} and standard-deviation σ_{ID} . Next, we iterate over all ID+CID datasets while computing the moving average of confidence of the past m predictions using a sliding-window (\mathcal{C}_{SW}). We record all instances of \mathcal{C}_{SW} dropping below a certain threshold ρ . These events are denoted as CID-predicted events. At the same time, we also compute \mathcal{A}_{SW} for evaluating how many CID-predicted events are actually a CID-detected event (true positive). Finally, the description of true positive (TP), false positives (FP), false negatives (FN) and true negatives (TN) are as follows.

- True positive: $\mathcal{C}_{SW} < \rho$ and $\mathcal{A}_{SW} \leq \mu_{ID} - 3 \cdot \sigma_{ID}$
- False positive: $\mathcal{C}_{SW} < \rho$ and $\mathcal{A}_{SW} > \mu_{ID} - 3 \cdot \sigma_{ID}$
- True negative: $\mathcal{C}_{SW} > \rho$ and $\mathcal{A}_{SW} > \mu_{ID} - 3 \cdot \sigma_{ID}$
- False negative: $\mathcal{C}_{SW} > \rho$ and $\mathcal{A}_{SW} \leq \mu_{ID} - 3 \cdot \sigma_{ID}$

In this way, we collect the TP, FP, TN, FN from each ID+CID datasets and compute the average precision and recall for a given threshold ρ . Finally, we report the AUPRC. We choose AUPRC because the precision-recall curve is resistant to imbalance in datasets.

B. Ablation Studies

B.1. Uncertainty quantification vs Ensemble size

In this section, we conduct an ablation study to investigate the effect of number of early-exits on uncertainty estimation quality. The ensemble size $|\mathcal{K}|$ is an hyperparameter that depends on the computation/resource budget, and it is bounded above by the depth of the base network. We vary the ensemble size and investigate its effect on calibration quality in MobilenetV2. Figure 7 shows the effect on accuracy and NLL in MobilenetV2 for $|\mathcal{K}|$ ranging from 2-10. The red line




Figure 7: Effect of Ensemble size on Accuracy and NLL in MobilenetV2 on TinyImagenet. Red dotted horizontal line indicates accuracy and NLL of BASE in respective plots

shows the accuracy and NLL of base network in respective plots. As shown, the accuracy does not always improve with ensemble size. This demonstrates that there is a limit to the improvement in accuracy due to early-exits. On the other hand, NLL steadily improves as ensemble size increases with the least NLL obtained for $|\mathcal{K}| = 8$. For $|\mathcal{K}| = 10$, the NLL rises slightly. Our investigations revealed that since the base network, early-exits and the early-view (EV) exits are part of the same network architecture, and are trained simultaneously, the disruption in base network’s weights due to the losses of early-exits and EV-exits reaches a tipping-point after a certain ensemble size thereby, causing a drop off in performance.

B.2. Ablations with EE-ensemble

B.2.1. EXCLUDING FINAL EXIT FROM FINAL PREDICTION

As described in Section 5, we exclude the original output block of the base network when computing the final prediction vector for QUTE because we find it that is overconfident and harms calibration quality. However, we include the original output block’s prediction vector along with that of all early-exits for computing the final prediction vector for EE-ensemble, consistent with the original paper Qendro et al. (2021). We conduct an ablation study on CIFAR10 using Resnet-8, where like QUTE, we exclude the original output block’s prediction vector when computing the final prediction vector. Table 9 reports the calibration metrics on CIFAR10-ID for two configurations: 1) including original output block in computing final prediction vector and 2) excluding original output block in computing final prediction vector.

Resnet-8	Including final-exit			Excluding final-exit		
	F1 (↑)	BS (↓)	NLL (↓)	F1 (↑)	BS (↓)	NLL (↓)
CIFAR10						
- EE-ensemble	0.854	0.021	0.446	0.818	0.0026	0.561

Table 9: Calibration Metrics for Resnet-8 on CIFAR10-ID computed including original output block and excluding original output block from computation of final prediction vector.

As seen, the removal of original output block negatively impacts EE-ensemble leading to poor accuracy and calibration. Therefore, we include the original output block in computing the final prediction vector in our main results for EE-ensemble.

B.2.2. USING QUTE’S OUTPUT BLOCK ARCHITECTURE FOR EARLY-EXITS IN EE-ENSEMBLE

Qendro et al. (2021) adds additional fully-connected layers at the early-exit for EE-ensemble to match the learning capacities of all exits. However, this leads to a large overhead (Table 4). In this ablation study, we show that the additional layers at the early-exit are necessary for EE-ensemble to achieve good accuracy and calibration. We replace the resource-hungry early-exits of EE-ensemble with the architecture of QUTE’s output blocks.

Resnet-8	In-distribution			Corrupted-in-distribution		
	F1 (↑)	BS (↓)	NLL (↓)	F1 (↑)	BS (↓)	NLL (↓)
CIFAR10						
- EE-ensemble w/ additional layers	0.85	0.021	0.446	0.64	0.0049	1.256
- EE-ensemble w/o additional layers	0.78	0.031	0.661	0.57	0.0058	1.459

Table 10: Calibration Metrics for Resnet-8 on CIFAR10-ID computed for EE-ensemble with additional FC layers and without additional FC layers

Table 10 reports the calibration metrics for two configurations: 1) EE-ensemble with additional FC layers at the early-exit and 2) EE-ensemble with QUTE’s lightweight output block architecture at the early-exits. As seen, the configuration with extra learning layers clearly outperforms the one without on both ID and CID in terms of accuracy and calibration. The early-exits placed at different depths work with much less information compared to the final exit. As a result, the accuracy and calibration invariably degrade if the learning capacities of all early-exits do not match that of the final exit.

B.3. MCU results

Model	SpeechCmd			CIFAR10		
	Accuracy	Size (KB)	Latency (ms)	Accuracy	Size (KB)	Latency (ms)
- BASE	0.92	294.5	22.7	0.83	344.4	58.04
- DEEP	0.93	321.6	45.4	0.88	423.4	116.2
- EE-ensemble	0.92	361.5	51.7	0.85	735.5	65.4
- QUTE	0.93	297.8	24.07	0.85	355.3	59.6

Table 11: Microcontroller results on Big-MCU (STM32F767ZI)

Tables 11 and 12 reports the on-device results for Big-MCU and Small-MCU respectively. We report the accuracy, code-size and the latency of each method. On Big-MCU, all methods fit on the device. The low memory-footprint of QUTE allows it

QUTE: Quantifying Uncertainty in TinyML models with Early-exit assisted ensemble

Model	SpeechCmd			CIFAR10			MNIST		
	Accuracy	Size (KB)	Latency (ms)	Accuracy	Size (KB)	Latency (ms)	Accuracy	Size (KB)	Latency (ms)
- BASE	0.92	84.8	160.7	0.83	187.5	291.2	0.906	100.3	5.01
- DEEP	0.93	111.3	321.2		DNF / OOM		0.93	109.3	9.3
- EE-ensemble	0.92	157.7	376.3		DNF / OOM		0.93	113.5	23.7
- QUTE	0.93	93	173.1	0.85	201	298.05	0.92	108.2	5.8

Table 12: Microcontroller results on Small-MCU (STM32L432KC). DNF / OOM indicates *did-not-fit/out-of-memory*

reduce the code-size consistently across datasets, leading to lower latency and, consequently, reduced energy per prediction. In contrast, EE-ensemble has higher memory requirements due to the expensive dense layers at its early exits. Furthermore, with no parallel execution capabilities on an MCU, the latency of DEEP with two ensemble members takes roughly twice as long to execute compared to BASE. This issue could worsen significantly with a larger ensemble size.

We observe a similar trend on the Small-MCU, with QUTE providing the best uncertainty quantification (UQ) per unit of resource used. However, both EE-ensemble and DEEP exceed the available 256KB of ROM on Small-MCU, making them unsuitable for ultra-low-power devices. In such scenarios, QUTE emerges as the only practical solution for effective UQ. Furthermore, we observe an interesting trend on the MNIST dataset. When using a 4-convolution layer model (4-layer CONV), the code sizes of all methods are quite similar, even exceeding that of the SpeechCmd BASE model, which employs a 4-layer depthwise-separable convolutional model (4-layer DSCNN) with 20,000 more parameters. A deeper analysis of the MAP files, which detail memory allocation, reveals that MNIST uses standard convolutions instead of separable convolutions. Therefore, the memory required for code is significantly higher for the 4-layer CONV model compared to the 4-layer DSCNN. In fact, the memory required for operators and code in the 4-layer CONV model is four times greater than the memory needed to store its weights.

These results highlight the efficiency of QUTE. While other methods may struggle with memory limitations, QUTE manages to deliver effective UQ without compromising on resource utilization. Its design allows for optimal performance even in constrained environments, making it a compelling choice for applications where both accuracy and efficiency are of paramount importance.

C. Evaluation metrics

We evaluate uncertainty using the following metrics.

Brier Score (lower is better): It is a proper scoring rule that measures the accuracy of predictive probabilities. Incorrect predictions with high predictive confidence are penalized less by BS. Thus, BS is less sensitive to corruptions and incorrect predictions. Therefore, NLL is a better measure of uncertainty to compare with other methods.

$$BS = \frac{1}{N} \sum_{\substack{n=1 \\ l \in \{1,2,..L\}}}^N (p_{\Theta}(y = l|x) - \mathbb{1}(y = l))^2 \tag{4}$$

where, $\mathbb{1}$ is an identity function and L is the number of classes.

Negative log-likelihood (lower is better): It is a proper scoring rule that measures how probable it is that the predictions obtained are from the in-distribution set. It depends on both the uncertainty (predictive confidence) and the accuracy of the predictions.

$$NLL = - \sum_{l \in \{1,2,..L\}} \mathbb{1}(y = l) \cdot \log p_{\Theta}(y = l|x) + (1 - \mathbb{1}(y = l)) \cdot \log(1 - p(y = l|x)) \tag{5}$$

Expected Calibration Error (ECE) (lower is better): ECE measures the absolute difference between predicted confidence and actual accuracy across different confidence intervals. The confidence is in the range $[0,1]$. ECE divides the confidence range into M intervals/bins of size $\frac{1}{M}$ and computes the bin accuracies and confidence of each bin before averaging them to provide the final ECE score.

$$ECE = \sum_{m=1}^M \frac{|B_m|}{N} \cdot |acc(B_m) - conf(B_m)| \tag{6}$$

where, N is total number of predictions, B_m is the m^{th} bin spanning the interval $(\frac{m-1}{M}, \frac{m}{M})$. A low ECE score is desirable, and indicates less disparity between confidence and accuracy across all intervals. However, the ECE score has several limitations and susceptibilities.

Limitations of ECE

1) ECE divides predicted probabilities into discrete bins. However, there might be a severe imbalance of samples across the bins because neural networks tend to always predict with high confidence (Nixon et al., 2019). This causes only a few bins (concentrated towards the high confidence region) to contribute the most to ECE. In addition, ECE is also sensitive to the bin boundaries i.e., the number of bins.

2) ECE does not differentiate between specific types of miscalibration such as overconfidence/underconfidence, often providing a simplistic and (overly) optimistic view of the calibration quality. Therefore, ECE is not a proper scoring metric. A model might have a lower ECE without genuinely having good calibration.

For these reasons, we do not emphasize ECE in our main results as they are not indicative of the model’s true calibration quality. We focus more on proper scoring metrics like NLL and BS in our main results.