

The Benefit of the Doubt: Uncertainty Aware Sensing for Edge Computing Platforms

Lorena Qendro
University of Cambridge

Jagmohan Chauhan
University of Southampton
University of Cambridge

Alberto Gil C. P. Ramos
Samsung AI Center Cambridge

Cecilia Mascolo
University of Cambridge

ABSTRACT

Neural networks (NNs) have drastically improved the performance of mobile and embedded applications but lack measures of “reliability” estimation that would enable reasoning over their predictions. Despite the vital importance, especially in areas of human well-being and health, state-of-the-art uncertainty estimation techniques are computationally expensive when applied to resource-constrained devices. We propose an efficient framework for predictive uncertainty estimation in NNs deployed on edge computing platforms with no need for fine-tuning or re-training strategies. To meet the energy and latency requirements of these systems the framework is built from the ground up to provide predictive uncertainty based only on one forward pass and a negligible amount of additional matrix multiplications. Our aim is to enable already trained deep learning models to generate uncertainty estimates on resource-limited devices at inference time focusing on classification tasks. This framework is founded on theoretical developments casting dropout training as approximate inference in Bayesian NNs. Our novel layerwise distribution approximation to the convolution layer cascades through the network, providing uncertainty estimates in one single run which ensures minimal overhead, especially compared with uncertainty techniques that require multiple forwards passes and an equal linear rise in energy and latency requirements making them unsuitable in practice. We demonstrate that it yields better performance and flexibility over previous work based on multilayer perceptrons to obtain uncertainty estimates. Our evaluation with mobile applications datasets on Nvidia Jetson TX2 and Nano shows that our approach not only obtains robust and accurate uncertainty estimations but also outperforms state-of-the-art methods in terms of systems performance, reducing energy consumption (up to 28-folds), keeping the memory overhead at a minimum while still improving accuracy (up to 16%).

CCS CONCEPTS

• **Human-centered computing** → **Ubiquitous and mobile computing**; • **Computing methodologies** → **Neural Networks**; **Supervised learning by classification**..

KEYWORDS

Edge Platforms, Sensing, Bayesian Deep Learning, Uncertainty

ACM Reference Format:

Lorena Qendro, Jagmohan Chauhan, Alberto Gil C. P. Ramos, and Cecilia Mascolo. 2021. The Benefit of the Doubt: Uncertainty Aware Sensing for Edge Computing Platforms. In *SEC '21: ACM/IEEE Symposium on Edge Computing, December 14-17, 2021, San Jose, CA*. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

Deep learning models are becoming the de facto standard in mobile and embedded application including activity and context recognition [18, 29, 52], and health and well-being monitoring [7, 16, 37, 46, 47, 50]. However, deep learning models are being also scrutinized due to their lack of interpretability. Understanding predictive uncertainty is important for all types of machine learning tasks but is key in situations when these are relied upon by the medical profession: examples of embedded systems in these domains are soaring, e.g for monitoring gait in Parkinson’s disease patients [27], detection of cardiac arrest via audio from smart devices [11] or revealing sleep apnea [36] through sensing and audio applications.

Predictive uncertainty gives the crucial indication of whether the system is guessing at random or if it is confident about the output inferred. This, as a consequence, has implications on the decision made based on these predictions, perhaps deciding when a human in the loop is necessary to complement the system or balancing the computation in a distributed multi-tier (cloud) architecture in case of uncertainty. High uncertainty predictions might need further computation so the sample might be sent for processing to the cloud or on bigger capacity devices.

Probabilistic approaches exist to provide frameworks for modeling uncertainty towards capturing the erroneous overconfident decisions. However, enabling such approaches on deep neural networks bring significant challenges on embedded devices. Firstly, the most popular approach to provide uncertainty estimates, namely Bayesian Neural Networks (BNNs) require heavy computation. Although, recent efforts have been devoted to making them more efficient, their improvements are still not a good fit for mobile or embedded devices since they are based either on *sampling* [19–21] or model *ensembles* [32]. While sampling demands running a

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SEC '21, December 14-17, 2021, San Jose, CA

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00

<https://doi.org/10.1145/1122445.1122456>

single stochastic neural network multiple times, ensemble methods require training and running multiple neural networks which linearly increases latency if ran in sequence or memory if ran in parallel. Indeed, these solutions are resource agnostic and would incur unfeasible increases in power consumption, latency and memory requirements on many mobile devices with limited resources. Secondly, there is very limited work on alternatives to BNNs providing predictive uncertainty in embedded systems [56, 57] applicable only to multi layer perceptrons (MLPs). However, as is well known in the aforementioned applications areas, using convolution neural networks (CNNs) instead leads to more accurate predictions than MLPs [27, 55]. Indeed, the majority of modern embedded deep learning models do not rely solely on MLPs but are often a combination of different neural layers, CNNs and MLPs [27, 39, 55]. As a consequence, these approaches although suitable for embedded devices are not relevant for the types of deep learning models that are actually being deployed in practice. Moreover, they focus mainly on regression tasks, leaving a considerable amount of questions on how they can be used in classification contexts.

In light of the highlighted challenges, we propose a framework that addresses these limitations by enabling *predictive uncertainty* estimations for mobile and embedded applications and evaluating its efficiency on resource constrained devices. Overall, we make the following contributions in this paper:

- We introduce an efficient framework that directly enables already trained deep learning models to generate uncertainty estimates with no need for re-training or fine-tuning. Its core is based on theoretical developments casting dropout training as approximate inference in Bayesian Convolutional Neural Networks [19]; we consider models that have been already trained with dropout as a regularization method. This assumption is easily satisfiable, since most of the modern deep learning networks use dropout during training [27, 55]. To achieve our goal in providing the uncertainty estimates, we propose an efficient layerwise distribution approximation, which transforms the single deterministic convolutional layer into a stochastic convolutional layer. Unlike previous methods that generate the prediction distribution via multiple runs [19, 21, 32], our layerwise distribution is propagated through the network in a cascaded manner massively reducing the computational complexity by allowing the model to produce uncertainty estimations in one single run. This approach makes it possible therefore to enable predictive uncertainty on a much wider range of small devices where running uncertainty aware deep learning models would be impossible with traditional techniques.
- Our approach focuses on classification tasks which makes obtaining uncertainty estimates challenging. Unlike regression, in a classification scenario, we cannot interpret the output distribution as the model prediction output. To solve this problem, we introduce an efficient way to marginalize over the final distribution to capture the predictive uncertainty and present the class accuracy. Moreover, our approach is able to offer the desired flexibility by enabling predictive uncertainty into CNNs which have better predictive power than MLPs. Combining CNNs with layerwise distribution approximations become a powerful tool to estimate uncertainty




Fig. 1: Density histogram of confidence measures using vanilla DNN. In (a), true (correct) and false (incorrect) predictions; in (b), random data.

while offering higher accuracy compared to the existing works which utilize MLP based models [56, 57].

- We evaluate our framework on the Nvidia Jetson TX2 and Nano embedded platforms on human activity recognition (HAR) and audio sensing applications. We compare our approach with the state-of-the-art Monte Carlo dropout [21], a fully connected network based approach [56] as well as deep ensembles technique [32]. For all approaches, we measure the resource consumption (latency and energy) and model performance, such as the accuracy, and the quality of uncertainty estimations. Our approach can reduce inference and energy consumption by 8-fold to 28-fold, while obtaining robust and accurate uncertainty estimation. We also significantly improve the accuracy of the deep learning models, compared to previous work based on fully connected layered MLP models [57] by a margin of at least 6% to 16% while being more cost-effective computationally. We make sure not too heavily contribute to the memory footprint by adding only a negligible runtime memory overhead (max 5%) compared to the vanilla deep learning model and improving (by 30%) on the MLP baseline. We show that our technique can smoothly run also on CPU only, allowing devices without GPU to still have fast and robust uncertainty aware predictions.

2 MOTIVATION

Limited previous work in mobile and embedded systems [56, 57] has empirically studied ways to provide uncertainty estimations in deep learning models. These techniques mostly focus on regression, leaving the classification scenario relatively unexplored. Classification tasks make the highest percentage of mobile sensing applications [16, 27, 34, 39, 55] but providing uncertainty estimations in the context of these types of data and resource constrained devices is still an open research area. Deterministic DNNs are trained to obtain maximum likelihood estimates and therefore do not consider uncertainty around the model parameters that leads to predictive uncertainty. They provide overconfident decisions as the softmax probability only captures the relative probability that the input is from a particular class compared to the other classes but not the overall model confidence.

To visualize the aforementioned issue, we analyze a human activity recognition scenario. This experiment is performed on the Heterogeneous human activity recognition dataset (HHAR) [1] which consists of readings from two motions sensors (accelerometer and gyroscope) on nine users performing six activities (biking, sitting, standing, walking, climb stairs-up, and stairs-down). We choose a 5 layer network (4 convolutional layers and 1 fully connected) and evaluated the behavior of softmax on the test set as well as on random input data. In Figure 1a, we can see the confidence measures for true (correct) and false (incorrect) predictions. A distribution skewed towards the right (near 1.0 on the x-axis) shows that the model has higher confidence in predictions than the distributions skewed towards left. As we can notice, this model gives high probability for both correct and incorrect predictions. Instead of the current outcome, we would like to see low confidence for the false predictions.

An even more concerning observation can be noticed in 1b where we plot the confidence given by the same model on completely random data which shows that vanilla DNNs are overconfident even when presented with random data the model has not seen during training. This limitation of conventional deep learning approaches motivates our work. We aim to provide accurate predictions while understanding if the model is guessing at random or it is certain about the prediction.

In addition, we want the deep learning models to run on resource-limited devices, therefore, to be latency and energy savvy. Our aim is to overcome the computation overhead of sampling-based Bayesian techniques and other non-Bayesian approaches like deep ensembles. Running a single stochastic NN multiple times for each prediction, or needing to retrain or fine-tune existing model(s), is not feasible for many edge platforms. To close this gap, we build a new framework that can enable uncertainty estimates for currently deployed models under the constraints that it must require only one forward pass for each prediction, no retraining or fine-tuning, and incur only a residual increase in latency and memory. What makes this possible is an approximation to the internal statistics of neural networks, that allows an approximate propagation of the signal and confidence through the network layers. Its core is a layerwise distribution approximation which allows to create a stochastic convolution layer that enables uncertainty estimates to be approximately propagated from inputs to outputs in a single run, which incurs only negligible increased linear algebra cost.

3 RELATED WORK

The benefits that may result from providing uncertainty estimates for predictive models have long been recognized. Indeed, dating back to 1992, seminal work in [38] shows several benefits of stochastic approaches to neural network learning, such as naturally accounting for the model flexibility, aiding in comparison between different models, accurate calibration of predictive uncertainty, and robustness to overfitting to name a few. Naturally, the vast majority of this rich literature aims at small scale problems and precludes deployability considerations. More recently, there has been a significant scale up of the problems such techniques can be applied, but the subject matter of designing algorithms for uncertainty quantification on edge devices is in its infancy. In the following, we attempt

to provide a brief overview of recent developments, and how this work extends those to a wider range of devices.

3.1 Uncertainty Estimations

Modern research in Bayesian Neural Networks (BNN) relies on variational inference to approximate the intractable posterior of Bayesian inference [6, 9, 25, 26, 53]. Pioneering work from Gal et al. [19–21] introduced Monte Carlo dropout (MCDrop) interpreting dropout used to train deterministic neural networks [28, 48] to approximately correspond to variational inference. MCDrop collects the results of stochastic forward passes through a model with dropout enabled at inference time and estimates the predictive uncertainty via sampling-based techniques. Deep Ensembles [32], instead, are a non-Bayesian way to get uncertainty estimations by training and running multiple DNNs.

All the work discussed so far focuses on investigating more accurate uncertainty estimates and does not consider the system implications of mobile and embedded computing; as a consequence, the proposed methods often i) require training new models from scratch or retraining/fine-tuning existing models with a development cost that might inhibit their use and/or ii) are computationally prohibitive i.e., require a linear increase in latency or memory due to multiple forward passes through a single model or one forward pass through several models. We on the other hand take a different approach and focus primarily on providing a simple and effective solution that enriches existing deep learning models with predictive uncertainty estimations in a manner that does not require retraining/fine-tuning and that ensures a latency, memory and energy consumption in the same ballpark as the original model.

3.2 Mobile Applications, Resource Constraints and Uncertainty

Numerous works have investigated the use of deep neural networks for human activity recognition (HAR) [27, 30, 41, 44, 45] and audio sensing [8, 12, 22, 23, 34]. These applications need intelligence at the edge, and therefore, deal with constrained resources. Recently, traditional DNNs are modified to fit in memory, increase execution speed, and decrease energy demand to inference models on the edge [10, 33, 54]. However, there is limited previous work that aims to enable uncertainty estimations on these models and platforms [56, 57]. ApDeepSense [57] proposes a seminal approach on approximating the output distribution using standard dropout [48], which aims to reduce computation time and energy consumption. However, this work only applies to fully-connected NNs, leaving the challenges for more complex models, like Convolutional Neural Networks (CNNs), still to be addressed. Modern architectures are very rarely solely MLP based [27, 55]. This suggests that extending CNNs to obtain uncertainty estimations makes it is possible to not only obtain higher accuracy but also more robust models in their predictive abilities. Moreover, being mainly engineered for regression, it does not use the distribution variance for the uncertainty in the classification task. Using the variance is fundamental to properly capture the uncertainty, without it the variability in the prediction is not taken into consideration. Classification scenarios make the majority of mobile sensing applications, therefore, we intentionally focus on these tasks in our work and we use the

variance by adding sampling to the logit distribution to simulate the sampling in MCDrop but at much lower computational cost.

Also, unlike previous works providing uncertainty estimation only in the last layer or after several runs, our approach provides flexibility through its layerwise approximation which captures the uncertain unobserved quantities at the layer level. This is useful in scenarios where models cannot execute all layers, possibly due to energy and latency constraints [35, 40] or would like to exit early in a densely layered model to save energy, while providing robust uncertainty estimates.

4 UNCERTAINTY ESTIMATION FRAMEWORK

We present a framework that enables pre-trained deterministic deep learning models to generate uncertainty estimates on resource-limited devices i) without increasing their latency, memory and energy profiles except for a minimal amount due to additional linear algebra and ii) without any retraining/fine-tuning which could prevent deployability on the same devices. This is in stark contrast to existing techniques for providing uncertainty estimates which have a linear (often unaffordable) increase in operational costs due essentially to their reliance on multiple forward passes, and which may require training new models.

4.1 Approach at a glance

Our framework is grounded on the fact that when using a neural network trained with dropout, the network does not have a deterministic structure anymore since it is described in part by random variables. Therefore, we enable the convolution operations in the stochastic network generated by dropout to apply to inputs described by probabilistic distributions and generate distributions as outputs (Figure 2). This means that now for every layer the network is described by inputs (and consequently outputs) which are Gaussian distributions with a mean and a variance. The layerwise distribution is going to capture the predictive uncertainty since we are not considering point estimates anymore.

As described in the visual overview in Figure 2, these distributions are then propagated through the whole network carrying and transmitting the information on the prediction and the uncertainty of these predictions. Finally, we integrate over the culminating Gaussian distributions to ultimately sample point values (unaries) and pass them through the softmax function. At this point, we only sample from the final distribution, which is a tiny fraction of the networks compute (see Section 5.6), and therefore, it does not significantly increase the model's inference time. At the end of this process, the classification model produces the class prediction and its predictive uncertainty based on the samples. This operation is extremely fast as we only run the model once, passing inputs to the models to get the output distribution.

While our approach uses the theoretical foundations of MC-Drop [19] casting dropout training as an implicit Bayesian approximation, it radically distinguishes from it as in our technique the distributions are embedded in the network and do not rely on multiple runs to produce them. In addition, we offer a novel mode

to adapt the distribution output to predict the outcome in classification tasks while providing the desired predictive uncertainty estimations.

4.2 Background

4.2.1 Uncertainty definition. Uncertainty is manifested in various forms in mobile applications on the edge. It can be at the physical layer due to node mobility, network topology, routing and resource availability [15]. These factors together with the sensor measurements, calibration and software heterogeneity add variability to the system feeding uncertainty to it. When using deep learning the uncertainty induced by the model architecture and parameters is an additional factor that jeopardizes the trust in the prediction. Observed data can be consistent with many models, and therefore which model is appropriate, given the data, is uncertain. Similarly, predictions about future data and the future consequences of actions are uncertain. Probability theory provides a framework for modelling uncertainty [24].

In our work, we define *predictive uncertainty* as the probability of occurrence of the target variable conditioned on all the available information. We treat the deep learning model predictions as random variables [43]. This means that we will have the predictions as probability distributions instead of point estimates. These probability distributions are used to represent all the uncertain unobserved quantities in a model (including structural, parametric and noise-related). In Bayesian modelling, there are two main types of uncertainty that could be modelled [17]. *Aleatoric* uncertainty captures the noise in the observations which might be sensor or motion noise. This noise gets translated into the uncertainty that cannot be reduced even if more data is collected. *Epistemic* uncertainty (also called *model uncertainty*), instead, represents the uncertainty in the model parameters and captures the ignorance about which model generated the collected data. *Predictive uncertainty* incorporates both *aleatoric* and *epistemic* uncertainty. In this work, we are modelling both uncertainties in one single framework.

Predictive uncertainty provides a better interpretation of the inference because it can indicate if the deep learning model is certain about the prediction or is just guessing at random. If we consider interpretation as an understanding of all the elements that contributed to the prediction, uncertainty estimation is the first step to detecting the anomaly and acts as a trigger to perform further investigations.

4.2.2 Dropout training (standard dropout). [48] proposes dropout as a regularization method to prevent over-fitting. The idea is to drop units from layers to avoid feature co-adaptation. For a *fully connected (FC) neural network* the linear operation can be:

$$\begin{aligned} \mathbf{y}^{(l)} &= \mathbf{x}^{(l)} \mathbf{W}^{(l)} + \mathbf{b}^{(l)} \\ \mathbf{x}^{(l+1)} &= f^{(l)}(\mathbf{y}^{(l)}) \end{aligned} \quad (1)$$

where for each layer l , $\mathbf{x}^{(l)}$ and $\mathbf{y}^{(l)}$ are the input and output of that layer, and $f^{(l)}(\cdot)$ is a nonlinear activation function. $\mathbf{W}^{(l)}$ is the weight matrix of l with dimensions $K^{(l)} \times K^{(l-1)}$ and $\mathbf{b}^{(l)}$ is the bias vector of dimensions $K^{(l)}$. Using dropout at the l^{th} layer is mathematically equivalent to setting the rows of the weight

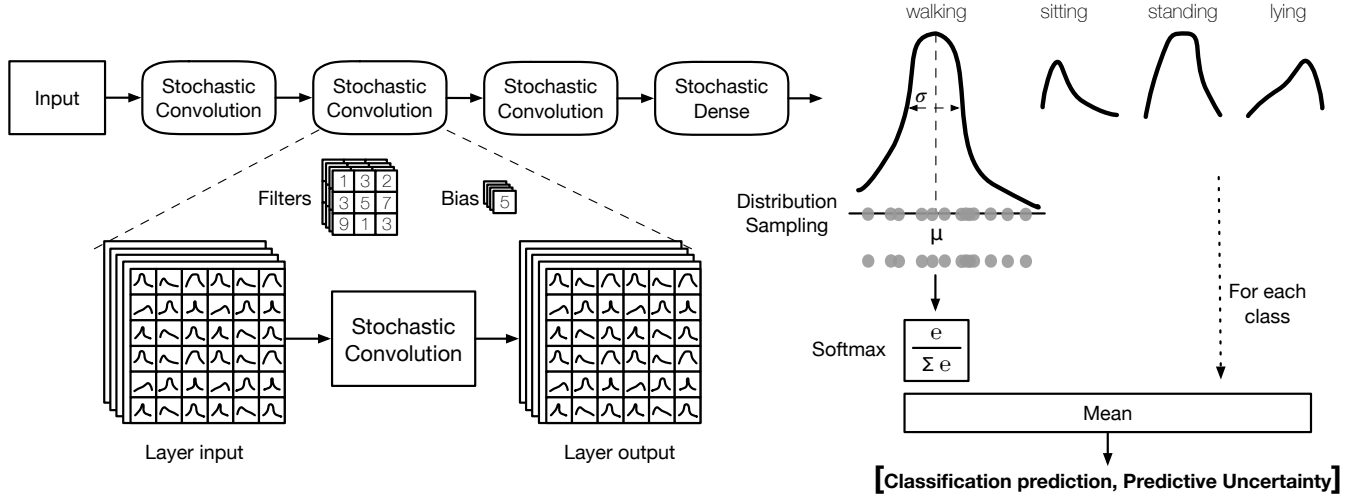


Fig. 2: Overview of the framework: describing the composition of the stochastic convolution, it's propagation throughout the network, and finally the distribution sampling at the softmax layer.

matrix $\mathbf{W}^{(l)}$ for that layer to zero. The FC layer can, therefore, be represented with dropout:

$$\begin{aligned} \mathbf{z}_{[i]}^{(l)} &\sim \text{Bernoulli}(\cdot | \mathbf{p}_{[i]}^{(l)}) \\ \tilde{\mathbf{W}}^{(l)} &= \text{diag}(\mathbf{z}^{(l)}) \mathbf{W}^{(l)} \\ \mathbf{y}^{(l)} &= \mathbf{x}^{(l)} \tilde{\mathbf{W}}^{(l)} + \mathbf{b}^{(l)} \\ \mathbf{x}^{(l+1)} &= f^{(l)}(\mathbf{y}^{(l)}) \end{aligned} \quad (2)$$

Here $\mathbf{z}_{[i]}^{(l)}$ are Bernoulli distributed random variables with some probabilities $\mathbf{p}_{[i]}^{(l)}$. The $\text{diag}(\cdot)$ maps vectors to diagonal matrices.

The described dropout operations convert a deterministic NN with parameters $\mathbf{W}^{(l)}$ into a random Bayesian neural network with random variables $\tilde{\mathbf{W}}^{(l)}$, which equates to a NN with a statistical model without using the Bayesian approach explicitly.

4.2.3 Dropout as Bayesian approximation. [19] proves the equivalence between dropout training in a neural network and approximate inference in a deep Gaussian process (GP). Differently from a non-probabilistic NN, a deep GP is a powerful statistical tool that allows to model distributions over functions. This means that to formulate the neural network layer as a Gaussian process we would define its covariance function

$$\mathbf{K}(\mathbf{x}, \mathbf{x}') = \int p(\mathbf{W}^{(l)}) f^{(l)}(\mathbf{x} \mathbf{W}^{(l)} + \mathbf{b}^{(l)}) f^{(l)}(\mathbf{x}' \mathbf{W}^{(l)} + \mathbf{b}^{(l)}) d\mathbf{W}^{(l)}, \quad (3)$$

with an element-wise linearity $f^{(l)}(\cdot)$ and distribution $p(\mathbf{W}^{(l)})$. If we consider now a deep GP with L layers and covariance function $\mathbf{K}(\mathbf{x}, \mathbf{x}')$, it can be approximated by setting a variational distribution over each element of a spectral decomposition of the covariance function of the GP. This spectral decomposition maps each layer of the deep GP to a layer of hidden units in the NN. For an L -layer neural network, this suggests that we can feed the output of one GP

to the covariance of the next GP matching a deep GP [14]. Hence, the final predictive distribution can be formulated as

$$p(\mathbf{y}|\mathbf{x}, \mathbf{X}, \mathbf{Y}) = \int p(\mathbf{y}|\mathbf{x}, \mathcal{W}) p(\mathcal{W}|\mathbf{X}, \mathbf{Y}) d\mathcal{W}, \quad (4)$$

where $p(\mathbf{y}|\mathbf{x}, \mathcal{W})$ is the whole Bayesian NN posterior with random variables $\mathcal{W} = \{p(\mathbf{W}^{(l)})\}$. To infer the predictive distribution estimation $p(\mathbf{y}|\mathbf{x}, \mathbf{X}, \mathbf{Y})$ we need to calculate $p(\mathbf{y}|\mathbf{x}, \mathcal{W})$ which is intractable. To address this, Gal et al. [19] proposed to use $q(\tilde{\mathbf{W}}^{(l)})$, a distribution over the weight matrices as follows:

$$\begin{aligned} \mathbf{z}_{[i]}^{(l)} &\sim \text{Bernoulli}(\cdot | \mathbf{p}_{[i]}^{(l)}) \\ q(\tilde{\mathbf{W}}^{(l)}) &= \text{diag}(\mathbf{z}^{(l)}) \mathbf{W}^{(l)} \end{aligned} \quad (5)$$

The true posterior distribution is, therefore, approximated by the variational distribution $q(\tilde{\mathbf{W}}^{(l)})$ where $\tilde{\mathbf{W}}^{(l)}$ represents the random variables used in dropout operations as described in (2). Finally, to approximate the predictive distribution $p(\mathbf{y}|\mathbf{x}, \mathbf{X}, \mathbf{Y})$ they perform Monte Carlo (MC) sampling of the random variables \mathcal{W} ,

$$q(\mathbf{y}|\mathbf{x}) = \frac{1}{T} \sum_{t=1}^T q(\mathbf{y}|\mathbf{x}, \mathcal{W}_t), \quad (6)$$

where T is the number of MC samples. This method is called Monte Carlo Dropout (MCDrop) and is equivalent to performing T stochastic passes.

In the same way as using dropout on the FC layer, MCDrop can be applied to the individual convolutions in *convolution neural networks* [19]. The convolution process is an iterative process that takes a sequence of steps in order to compute all elements in the input layer. Similarly to the FC layers, we can sample Bernoulli random variables $\mathbf{z}_{i,j,k}$ and apply them as masks to the weight




Fig. 3: Monte Carlo dropout performed in fully-connected and convolutional neural networks. Keeping dropout during inference creates an implicit ensemble of models. As shown it requires running the same network with different dropout masks in order to provide uncertainty estimations.

matrix $\mathbf{W}_i \cdot \text{diag}([z_i, j, k])$ which is equivalent to setting weights to 0 for different elements of the input.

Figure 3 illustrates how MC dropout is implemented. This technique relies on MC sampling and requires the whole network to run multiple times. Performing the multiple forward passes creates an implicit ensemble of models that differ to one other due to changes in the dropout masks. This implies that different runs will have different nodes and connections and provide the desired stochasticity in the Bayesian deep learning model. The predictive distribution is, therefore, not produced by explicitly placing distribution over the layers but by running multiple stochastic forward passes with dropout activated also during inference. These runs are then averaged and give an indication of the predictive uncertainty provided by the variance in the output, distribution variance in the regression context and softmax variance in the classification. Although MCDrop is a step forward towards optimizing and accelerating Bayesian deep learning techniques, it is not enough for running on edge platforms.

In this paper, we propose a novel solution, based on the theoretical foundations described in (3), (4), and (5), with the benefit of providing predictive uncertainties with solely one forward pass and test it on mobile applications running on embedded platforms.

4.3 Our Approach to Efficient Uncertainty Estimation

Our approach replaces the slow and computationally intense sampling process with a resource-friendly method by enabling convolution neural networks to output predictive uncertainty alongside the prediction.

We firstly propose a layerwise distribution approximation which allows to incorporate the distribution at the layer level and propagate it throughout the network. We choose this distribution to be a multivariate Gaussian distribution based on the equivalence of GPs

and neural networks trained with dropout. In Section 4.2.3, (3), (4), and (5) provide the theoretical proof that a deep neural network trained with dropout could be represented by a deep Gaussian process. In (3) we can see how the GP would be represented at the layer level and (4) formulates it for the whole deep learning network by feeding the output of one Gaussian process to the covariance of the next GP, and therefore matching a deep GP. Accordingly, the internal layer of the deep neural network can be represented by the internal representations of deep Gaussian processes. A GP is a stochastic process such that every finite collection of those random variables has a multivariate Gaussian distribution.

Therefore, we initially considered our layerwise distribution approximation to be represented by multivariate normal distributions. We found, however, that this approximation was not enough for avoiding multiple forward passes, and therefore explored an additional approximation of Gaussian distributions with diagonal covariance matrices. Indeed, by modeling inputs and outputs of each layer as approximately following Normal distributions with diagonal matrices, it is possible to compute the outputs mean and standard deviation in closed-form from those of the inputs as well as the layer operations for deterministic models trained with dropout, without requiring any retraining.

Our approach is based on an approximation to the internal statistics of neural networks, that permits an approximate propagation of the signal and confidence through the network layers, in a manner that is applicable to convolution and dense layers.

To start, we enable the basic operations in convolution neural networks to output the expected value and a probability distribution of the output random variable. Computing the exact output distribution is intractable; consequently, we approximate it with the multivariate Gaussian distribution. This approximation is based on minimizing the Kullback-Leibler (KL) divergence between the exact and approximate distributions.

$$\begin{aligned}
\min_q \text{KL}(p(x) \| q(x)) &= \min_q \int p(x) \log \left(\frac{p(x)}{q(x)} \right) dx \\
&= \min_{\mu, \sigma^2} - \int p(x) \log \mathcal{N}(x | \mu, \sigma^2) dx \quad (7) \\
&= \min_{\mu, \sigma^2} \frac{\log(\sigma^2)}{2} + \frac{\int x p(x) (x - \mu)^2 dx}{2\sigma^2}
\end{aligned}$$

Where $p(x)$ is the exact output distribution and $q(s) \sim \mathcal{N}(x | \mu, \sigma^2)$ is the approximate output distribution. To obtain the optimal approximate output distribution we take the derivative over the μ and σ^2 , therefore, the approximation can be represented as matching the mean and variance between the two distributions.

$$\begin{aligned}
\mu &= \int p(x) x dx \\
\sigma^2 &= \int p(x) (x - \mu)^2 dx \quad (8)
\end{aligned}$$

We reformulate the convolution operation with dropout as follows. The input to the layer is represented by $\mathbf{x} \in \mathbb{R}^{(H, W, C)}$ with height H , width W and C channels. Let $\mathbf{w} \in \mathbb{R}^{(h, w, c, f)}$ be the weight matrix with height h , width w , c channels and f filters, and $\mathbf{b} \in \mathbb{R}^{(f)}$ the bias vector. Consequently, the output will be represented as $\mathbf{y}_{[:, :, k]} := \mathbf{x} * \mathbf{w}_{[:, :, :, k]} + \mathbf{b}[k]$ for $\mathbf{y} \in \mathbb{R}^{(H-h+1, W-w+1, f)}$ and for $k = 1, \dots, f$ where \odot represents the element-wise operation and $*$ the convolution operation. Formally,

$$\mathbf{y}_{[i, j, k]} := \text{vec}(\mathbf{x}_{[i+[0:h], j+[0:w], :]} \odot \mathbf{z}_{[i+[0:h], j+[0:w], :]})^\top \text{vec}(\mathbf{w}_{[:, :, :, k]}) + \mathbf{b}[k]$$

where

$$\begin{aligned}
\mathbf{z}_{[i+[0:h], j+[0:w], l]} &\sim \text{Bernoulli}(\cdot | \mathbf{p}_{[i+[0:h], j+[0:w], l]}) \\
\mathbf{x}_{[i+[0:h], j+[0:w], l]} &\sim \mathcal{N}(\cdot | \mu_{[i+[0:h], j+[0:w], :]}, \sigma_{[i+[0:h], j+[0:w], l]}^2) \quad (9)
\end{aligned}$$

According to (7), we need to calculate the mean and the variance of the output distribution $p(\mathbf{y}_{[i, j, k]})$. Given that the Bernoulli variables \mathbf{z} and the input Gaussian variables \mathbf{x} , as shown in (9), are independent random variables, we can have the mean of the output as follows:

$$\begin{aligned}
\mathbb{E}[\mathbf{y}_{[i, j, k]}] &= \mathbb{E} \left[\sum_{u=i}^{i+h-1} \sum_{v=j}^{j+w-1} \sum_{a=1}^c \mathbf{x}_{[u, v, a]} \mathbf{z}_{[u, v, a]} \mathbf{w}_{[u-i, v-j, a]} \right] + \mathbf{b}[k] \\
&= \sum_{u=i}^{i+h-1} \sum_{v=j}^{j+w-1} \sum_{a=1}^c \mathbb{E}[\mathbf{x}_{[u, v, a]}] \mathbb{E}[\mathbf{z}_{[u, v, a]}] \mathbb{E}[\mathbf{w}_{[u-i, v-j, a]}] + \mathbf{b}[k] \\
&= \sum_{u=i}^{i+h-1} \sum_{v=j}^{j+w-1} \sum_{a=1}^c \mu_{[u, v, a]} \mathbf{p}_{[u, v, a]} \mathbf{w}_{[u-i, v-j, a]} + \mathbf{b}[k]
\end{aligned}$$

Since $\mathbf{x}_{[u, v, a]} \mathbf{z}_{[u, v, a]} \mathbf{w}_{[u-i, v-j, a]}$ are independent variables, we can measure the variance as:

$$\begin{aligned}
\text{Var}[\mathbf{y}_{[i, j, k]}] &= \text{Var} \left[\sum_{u=i}^{i+h-1} \sum_{v=j}^{j+w-1} \sum_{a=1}^c \mathbf{x}_{[u, v, a]} \mathbf{z}_{[u, v, a]} \mathbf{w}_{[u-i, v-j, a]} + \mathbf{b}[k] \right] \\
&= \sum_{u=i}^{i+h-1} \sum_{v=j}^{j+w-1} \sum_{a=1}^c \text{Var} \left[\mathbf{x}_{[u, v, a]} \mathbf{z}_{[u, v, a]} \mathbf{w}_{[u-i, v-j, a]} + \mathbf{b}[k] \right] \\
&= \sum_{u=i}^{i+h-1} \sum_{v=j}^{j+w-1} \sum_{a=1}^c \mathbb{E} \left[\left(\mathbf{x}_{[u, v, a]} \mathbf{z}_{[u, v, a]} \mathbf{w}_{[u-i, v-j, a]} + \mathbf{b}[k] \right)^2 \right] \\
&\quad - \left(\mathbb{E} \left[\mathbf{x}_{[u, v, a]} \mathbf{z}_{[u, v, a]} \mathbf{w}_{[u-i, v-j, a]} + \mathbf{b}[k] \right] \right)^2 \\
&= \sum_{u=i}^{i+h-1} \sum_{v=j}^{j+w-1} \sum_{a=1}^c \left(\left(\mu_{[u, v, a]}^2 + \sigma_{[u, v, a]}^2 \right) \mathbf{p}_{[u, v, a]} \right. \\
&\quad \left. - \mu_{[u, v, a]}^2 \mathbf{p}_{[u, v, a]}^2 \right) \mathbf{w}_{[u-i, v-j, a]}^2
\end{aligned}$$

We can further represent the operations in a simple way and efficiently compute the output distribution, namely

$$\begin{aligned}
\mathbb{E}[\mathbf{y}] &= (\mu \odot \mathbf{p}) * \mathbf{w} + \mathbf{b} \\
\text{Var}[\mathbf{y}] &= \left(\left((\mu \odot \mathbf{p}) + (\sigma \odot \sigma) \right) \odot \mathbf{p} - (\mu \odot \mu) \odot (\mathbf{p} \odot \mathbf{p}) \right) * (\mathbf{w} \odot \mathbf{w}) \quad (10) \\
\mathbf{y}_{[i, j, k]} &\sim \mathcal{N}(\cdot | \mathbb{E}[\mathbf{y}_{[i, j, k]}], \text{Var}[\mathbf{y}_{[i, j, k]}])
\end{aligned}$$

We have provided a mathematically grounded proof on how to calculate the mean and the variance of the output at each convolution layer on networks that have been trained with dropout. This means that we can have CNNs taking probabilistic distributions as inputs and generate distributions as outputs. Hence, avoiding the need for the computationally costly sampling to from these distributions.

Given our approximation of modeling the inputs and the outputs as Gaussian distributions with diagonal covariance matrix, what is required now is to compute the mean and variance of the activation function that follows the linear mapping, which can then be plugged into the Gaussian model. Towards this end, the mean and variance

$$\begin{aligned}
\mu_{\mathbf{y}} &= \mathbb{E}[\mathbf{y}] \\
\sigma_{\mathbf{y}}^2 &= \mathbb{E}[\mathbf{y} - \mu_{\mathbf{y}}]^2 \quad (11)
\end{aligned}$$

can be represented as a sum of expectations of the output activations with respect to the Gaussian input distribution, over the compact intervals where the activation function is linear. This computation can be done in closed-form via $\text{erf}(\cdot)$ for any piece-wise linear activation as demonstrated in [57], and in particular for the ReLU activation used in our work.

Secondly, we propose an efficient way to exploit the output distribution of our stochastic neural network to provide a classification prediction and the predictive uncertainty measure. To this aim, we marginalise these Gaussian distributions in the logit space. We sample unaries (single elements) from the output distribution and then pass the point values from this distribution to the softmax function.

$$\begin{aligned}
\mathbf{y} &\sim \mathcal{N}(\cdot | \mathbb{E}[\mathbf{y}], \text{Var}[\mathbf{y}]) \\
p(y=c | \mathbf{x}, \mathbf{X}, \mathbf{Y}) &\approx \frac{1}{S} \sum_{s=1}^S \text{Softmax}(\mathbf{y}_s) \quad (12) \\
H(\mathbf{y} | \mathbf{x}, D) &= - \sum_{i=1}^S p_{c\mu} * \log p_{c\mu},
\end{aligned}$$

where \mathbf{y} is the output distribution of the model. Therefore, the prediction can be considered as the mean of the categorical distribution obtained from sampling S single values from the Gaussian output distribution and then squashed with the softmax function to obtain the probability vector and the predictive entropy. The sampling operation from the output distribution is extremely fast as we only run the model once, passing inputs to the models to get the output logits. At this point, we only sample from the logits, which is a tiny fraction of the networks compute, and as we can see in Section 5.6, it does not significantly increase the model's inference time.

In conclusion, we add a layerwise approximation to the convolutional layers which is propagated throughout the network to produce a probability distribution in output. With the approximation presented in (10) and the output distribution sampling in

Table 1: Statistical information of the publicly available datasets used for evaluation. *Heterogeneous* human activity recognition (HHAR) [49], *Opportunity* [13] and *Speech Commands* [4] dataset.

| Dataset | Training size | Testing size | Output | One-hot Encoding |
|-----------------|---------------|--------------|---|------------------|
| HHAR | 28,314 | 1,686 | Sit, Stand, Walk, Bike, StairUp, StairDown | ✓ |
| Opportunity | 26,908 | 6,287 | Stand, Walk, Sit and Lie | ✓ |
| Speech Commands | 1,668 | 562 | Yes, No, Up, Down, Left, Right, On, Off, Stop, Go | ✓ |

(16) we can now enable classification models to output predictive uncertainties alongside the class inference in one single run.

5 EVALUATION

To evaluate the performance of our approach, we build a five-layer deep neural network composed of four 2D convolutional layers with a ReLU activation function and one fully-connected output layer. The choice of the architecture was made to make a fair comparison with the other baselines which mainly relied on five-layer deep networks [32, 56, 57].

During training, the model is optimized by ADAM [31] with a learning rate of $1e-4$. We add dropout of 0.5 (default) at each internal layer to stabilize training, avoid overfitting, and fulfill our requirement of having a model trained with dropout regularization. We use cross entropy as the loss function and a batch size of 64. For all the datasets, we use 5% of the training set for validation and hyper-parameter tuning. We employ the described architecture for all the datasets.

During inference, we enable our layerwise approximation to all layers and propagate it all the way to the output layer (Figure 2). This architecture allows to have the output of the model represented as mean and variance of the output distribution. However, since we are dealing with classification tasks, we sample unaries from the output distribution and pass them through the softmax function, as explained in Section 4.3. We have two inputs to the first layer of the NN: the sample we need to do inference on and the standard deviation of the data calculated on the training set. We use this as a prior to feed on the stochastic network.

5.1 Datasets

Heterogeneous human activity recognition dataset (HHAR) [49] contains readings from two motion sensors (accelerometer and gyroscope). The data is gathered from nine users performing six activities (biking, sitting, standing, walking, climb stairs-up, and climb stairs-down) with six types of mobile devices. Conforming to the description on [57], we segment raw measurements into five seconds samples and take Fourier transformation on these samples as the input data. Each sample is further divided into time intervals of length 0.25s.

Opportunity dataset [13] consists of data from multiple accelerometers and gyroscopes placed on participants' body at different locations such as arms, back, and feet. We used three devices on a hip, a left lower arm, and a right shoe by default and target to detect the mode of locomotion: stand, walk, sit and lie. In total, for all users and all recordings, the dataset consists of 3,653 modes of locomotion instances of variable duration (between 0.2 and 280 seconds). Following the preprocessing proposed by [27], we use run 2 from

user 1, runs 4 and 5 from user 2 and 3 in our test set. The remaining data is used for training. For frame-by-frame analysis, we created sliding windows of duration 1 second with 50% overlap.

For audio sensing, we use the *Speech Commands* dataset [4] and the suggested preprocessing by [39]. We train our network with the portion of the dataset that consists of 2,250 one-second long speech files belonging to 10 keyword classes (yes, no, up, down, left, right, on, off, stop, and go). In this task, the goal is to identify the presence of a certain keyword class in a given speech segment. This 10-class dataset was then randomly split into training (75%) and test (25%) class-balanced subsets to make sure we get the same amount of data for each class. The input to the model is a two-dimensional tensor extracted from the keyword recording, consisting of time frames and 24 MFCC features.

All three of our datasets were collected in a controlled environment, therefore, we augmented real-life noise to the datasets in a principled manner to include real-world variability. Data augmentation can encode prior knowledge on the data, result in more robust models, and provide more resources to the deep learning platform. For both HAR datasets, we used the data augmentation techniques proposed in [51] which consist of seven variations for IMU data on wearables: rotation, permutation, time-warping, scaling, magnitude-warping, jittering, and cropping. For example, permutation consists of randomly perturbing the temporal location of events within the same time window. Rotation, instead, consists of enriching the data to match different sensor placements like an upside-down position. These alterations allow us to consider real-life noise which cannot be witnessed in the data collected in a lab-controlled setting, e.g., rotating of a smartwatch. We randomly chose a period between 20 seconds and two minutes and applied a randomly selected augmentation method; the intervals between noise periods were randomly selected between two and five minutes. For audio, we sampled examples of ambient noise from a publicly available environment sound classification dataset [42] and added them to the audio dataset. We assume that only one single type of noise is present at a given time, and each noise lasts between 15 and 60 seconds.

5.2 Baseline algorithms

We tested our proposed method on the three datasets described above and compared the outcome with other four state-of-the-art approaches: conventional DNN, Monte Carlo Dropout (MC-Drop) [20], ApDeepSense [57] and Deep Ensembles [32].

BackboneDNN is a conventional deep neural network. In order to show the benefits of a stochastic approach, we need to compare it to the traditional deep learning network. This network is used as the non-Bayesian baseline. As mentioned before, our technique

Table 2: Accuracy, F1 Score, and Negative Log-Likelihood (NLL) for HHAR, Opportunity (Opp.), and Speech Command (SC) dataset. MCDrop- k represents MCDrop with k (3, 5, 10, and 30) runs (forward passes), instead DeepEnsemble- k represents an ensemble of k (3, 5, and 10) individual NNs. NLL measures the correspondence between the ground truth values and their predicted distributions. Lower NLL means higher correspondence. Bold entries indicate best performance.

| Model | Ours | BackboneDNN | ApDeepSense | MCDrop-3 | MCDrop-5 | MCDrop-10 | MCDrop-30 | Ensemble-3 | Ensemble-5 | Ensemble-10 |
|-----------------|--------------|-------------|-------------|----------|----------|-----------|-------------|------------|------------|-------------|
| Accuracy (HHAR) | 94.79 | 91.5 | 78.55 | 92.67 | 92 | 91.33 | 92.56 | 92.51 | 92.66 | 92.83 |
| F1 Score (HHAR) | 0.792 | 0.671 | 0.601 | 0.671 | 0.678 | 0.677 | 0.712 | 0.702 | 0.704 | 0.708 |
| NLL (HHAR) | 2.558 | 42.14 | 3.75 | 30.14 | 8.01 | 5.58 | 4.13 | 23.54 | 7.18 | 4.36 |
| Accuracy (Opp.) | 86.80 | 83.85 | 80.03 | 83.24 | 82.81 | 83.85 | 84.63 | 82.35 | 82.99 | 84.06 |
| F1 Score (Opp.) | 0.808 | 0.732 | 0.701 | 0.741 | 0.742 | 0.744 | 0.765 | 0.743 | 0.746 | 0.755 |
| NLL (Opp.) | 2.39 | 12.35 | 4.00 | 12.31 | 9.38 | 8.38 | 2.28 | 9.21 | 7.98 | 3.60 |
| Accuracy (SC) | 82.12 | 78.45 | 71.15 | 78.48 | 80.02 | 80.65 | 81.01 | 79.33 | 80.45 | 81.01 |
| F1 Score (SC) | 0.733 | 0.622 | 0.589 | 0.624 | 0.628 | 0.633 | 0.706 | 0.596 | 0.599 | 0.675 |
| NLL (SC) | 1.01 | 2.11 | 1.97 | 1.44 | 1.20 | 1.13 | 1.08 | 1.44 | 1.18 | 1.08 |

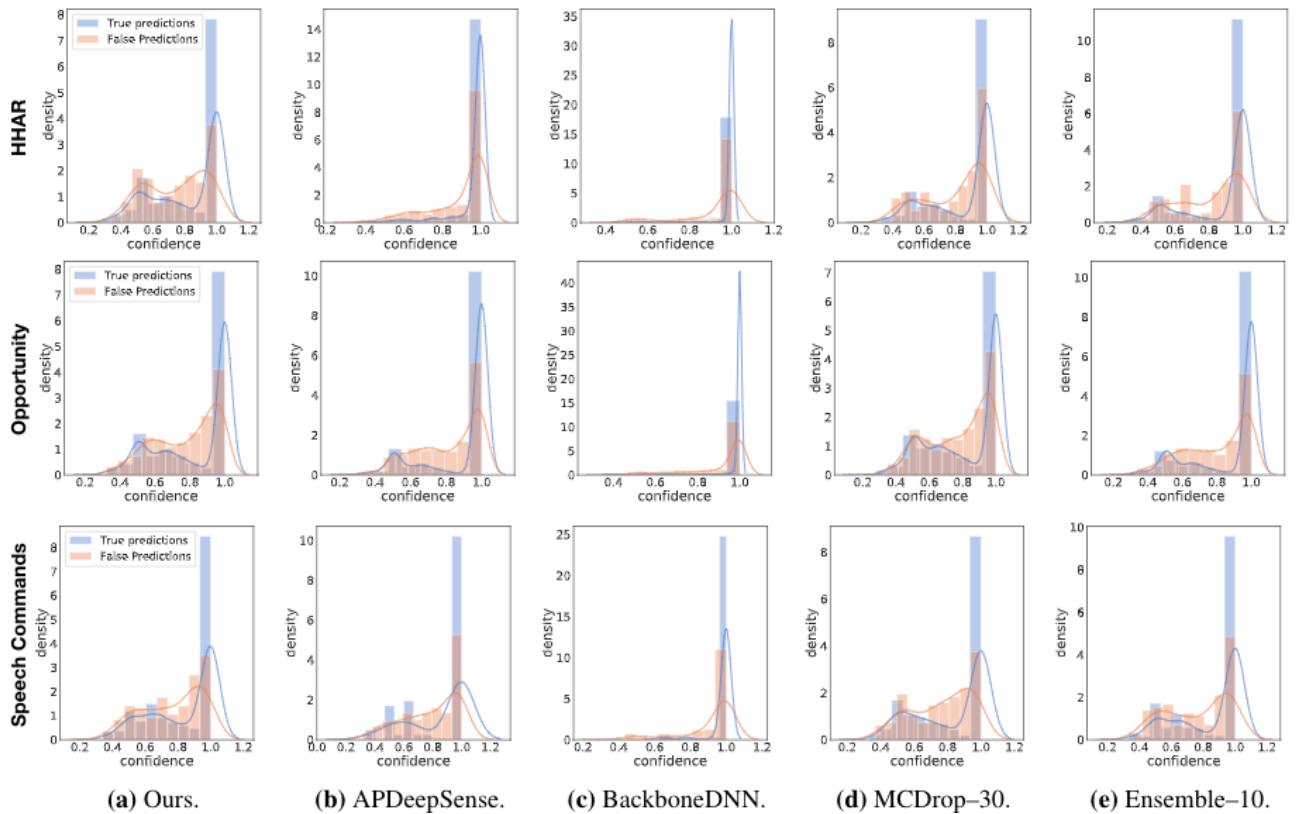


Fig. 4: Density histogram of confidence measures for true (correct) and false (incorrect) predictions. A distribution skewed towards the right (near 1.0 on the x-axis) indicates the model has higher confidence in predictions than the distribution skewed towards left. [The density histogram is a histogram with area normalized to one. Plots are overlaid with kernel density curves for better readability.]

and MCDrop rely on an already trained network, therefore, this network is the one we consider as the already trained network we refer to.

Monte Carlo Dropout (MCDrop) is based on Monte Carlo sampling and, runs the neural network multiple times to generate the uncertainty estimation. Hence, we use MCDrop- k to represent MCDrop with k (3, 5, 10, and 30) runs (forward passes). This approach, like ours, assumes that the model has already been trained

with dropout. For this baseline, we keep dropout activated during inference too.

ApDeepSense is an algorithm that enables fully-connected NNs to provide uncertainty estimations during inference. This technique too uses dropout to perform the basic operations in the FC layers. Compared to our method, ApDeepSense works only with MLPs and considers mainly regression tasks where the uncertainty is represented by the variance of the distribution. Therefore, it does not

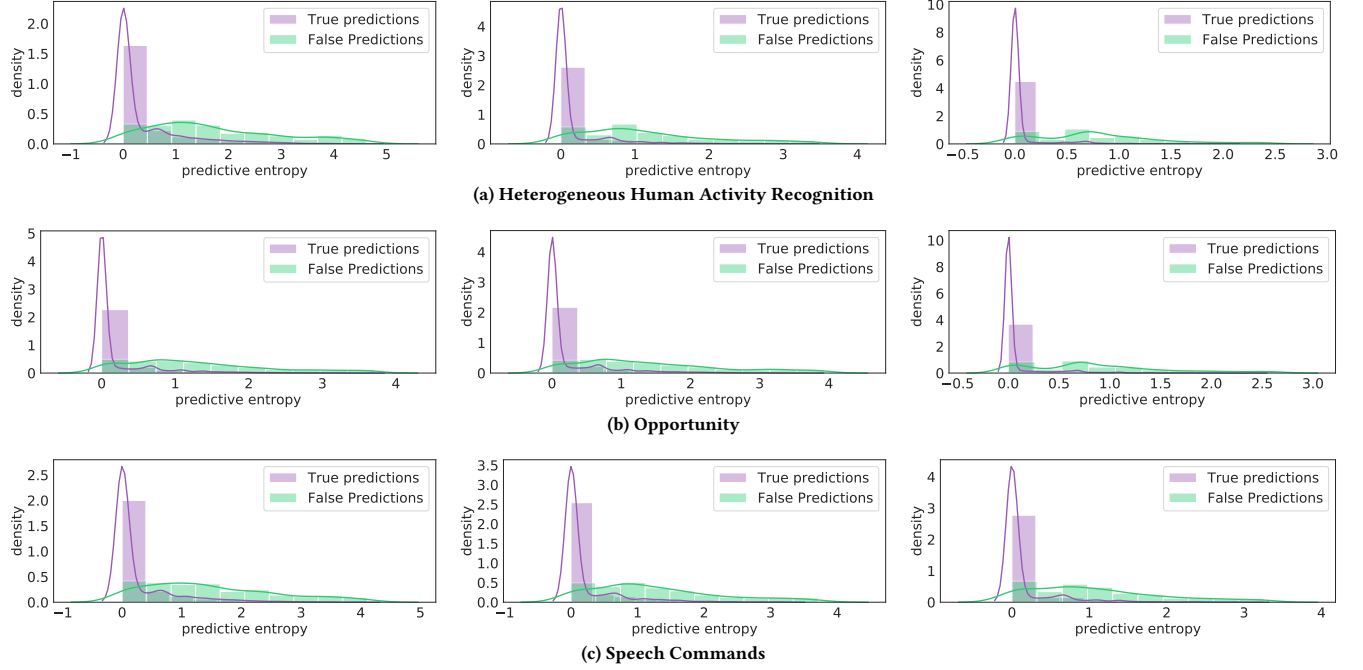


Fig. 5: Density histogram of predictive entropy for true (correct) and false (incorrect) predictions. Predictive entropy $H(y|x, D)$ which captures the average amount of information contained in the predictive distribution. To reliably capture the predictive uncertainty, we aim for a predictive entropy which is low for true and high for false inferences. [The density histogram is a histogram with area normalized to one. Plots are overlaid with kernel density curves for better readability.]

translate very well to classification tasks. For this baseline, as suggested in the original proposal [57], we use a 5-layer neural network composed of fully connected layers with 512 hidden dimensions and ReLU activation function.

Deep Ensembles rely on providing uncertainty estimations by training and running an ensemble of models (multiple networks). Although this baseline requires retraining, we include it to illustrate the upper bound of the uncertainty estimation quality that can be accomplished with retraining. Ensembles are created by training the models with random initialization [32]. To achieve this, we use the Backbone architecture (not its trained model) with random initialization for each model. We use DeepEnsemble- k to represent an ensemble of k (3, 5, and 10) individual NNs.

5.3 Quantitative Evaluations

In this section, we present the results in terms of accuracy, F1 score and negative log-likelihood (NLL). The prediction accuracy expresses the correlation between the prediction of the deep neural network and the actual value, instead the F1 score is the weighted average of precision and recall. NLL measures the correspondence between the ground truth values and their predicted distributions. Lower NLL means higher correspondence.

$$NLL(\mu, \sigma^2; x_1, \dots, x_n) = \frac{n}{2} \ln(2\pi) + \frac{n}{2} \ln(\sigma^2) + \frac{1}{2\sigma^2} \sum_{i=1}^n (x_i - \mu)^2 \quad (13)$$

In addition to the aforementioned metrics, we consider the confidence and predictive entropy as measure of uncertainty. The confidence metric gives a better understanding on how the model behaves during inference. Ideally we want to achieve high confidence for correct predictions and low confidence for incorrect ones. In classification tasks, it is considered as the confidence given from the softmax. In conventional DNN and APDeepsense, it is measured as the result based on only one softmax operation. Instead in the other baselines, including ours, it is the mean of categorical predictive distribution. To evaluate the predictive uncertainty we measure the predictive entropy $H(y|x, D)$ which captures the average amount of information contained in the predictive distribution.

$$H(y|x, D) = -\sum_{i=0}^{K-1} p_{i\mu} \log p_{i\mu}, \quad (14)$$

where $p_{i\mu}$ is the predictive mean probability of the i^{th} class from T Monte Carlo samples in the case of MCDrop, the K model predictions in the case of Deep Ensembles, and finally, from the j unaries sampled from the output distribution in our approach. To reliably capture the predictive uncertainty, we aim for a predictive entropy which is low for true and high for false inferences.

5.4 Embedded Edge Systems Setup

To evaluate our framework's performance, we run the inference on two edge platforms, NVIDIA Jetson TX2 [2] and Nano [3], and measure the aforementioned metrics while computing the latency and energy consumption per testing sample. The latency indicates

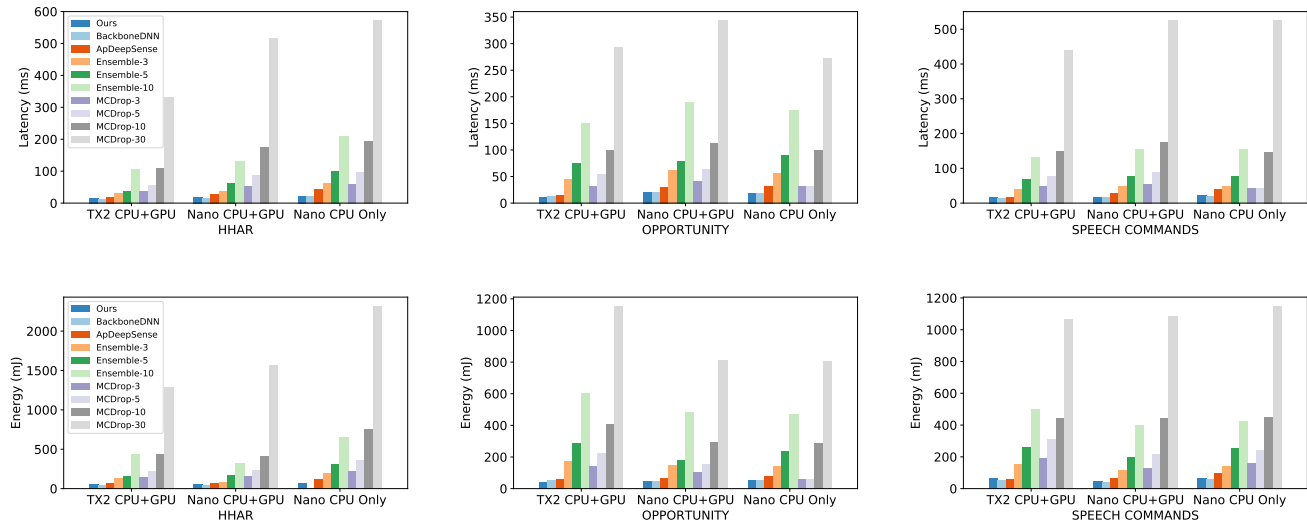


Fig. 6: Inference time and Energy consumption on two edge platforms with different CPU and GPU capacity. The latency indicates the average time it takes for the deep learning model to make a prediction on the provided sample. The energy consumed is expressed as $power \times time$.

the average time it takes for the model to make a prediction on the provided sample. The energy consumed is expressed as $power \times time$. The TX2 is an embedded system-on-module and it is representative of today’s embedded platforms with capabilities (256-core Pascal GPU, 2-core Denver + 4-core ARM A57 CPU, 8GB RAM, input ~19V) similar to high-end smartphones such as Samsung Galaxy 20 and OnePlus 7 Pro (Octa-core CPUs, Adreno 640/650 GPU, and 12GB RAM). The Nano, instead, has lower capabilities (128-core Maxwell GPU, 4-core ARM A57 CPU, 4GB RAM, input ~5V) and highlights the efficiency and feasibility of our approach (see Section 5.6) to run on more constraint embedded platforms. Both platforms have CPU and GPU support for deep learning libraries (we use tensorflow 1.15.2). The platforms used are equipped with power monitoring chips (TI INA226 [5]) which measure voltage supply and current draw and are accessible by software. We compute power and energy from their logs. We perform 10 runs (with the whole test set) and average the results. For an evaluation catering for more resource-limited devices, we evaluate our framework on the Jetson Nano with CPU only enabled (with all the datasets). This validates the applicability to these kind of limited platforms.

5.5 Model Estimation Performance

Table 2 shows the accuracy, F1 scores and negative log-likelihood obtained for the three datasets. Negative log-likelihood (NLL) penalizes both over and under-confident predictions. Overconfident predictions can be infinitely penalized, which discourages probabilities equal to zero or one. The minimum value is zero which can be achieved if the prediction is correct with 100% confidence. Otherwise, the worse the prediction, the larger the loss. Since NLL reflects both the accuracy and the quality of predictive uncertainty, our approach is the best performing across the three baselines: we can provide accurate predictions with high-quality uncertainty

estimates. We achieve higher accuracy compared to the other methods, especially ApDeepSense, because we employ CNNs instead of MLPs. MCDrop does achieve accuracy similar to our method, however, it takes 30 runs to reach a NLL close to the one of our method, which makes it very power hungry and not suitable for resource-constrained devices as highlighted in Section 5.6. The similar argument holds for Ensemble-10. An ensemble not only requires running multiple models in order to get the predictive uncertainty, but it also demands for keeping them in memory. In case of many embedded devices this would involve some kind of scheduling for memory allocation to perform them all. In order to apply this technique, it is required to train multiple models, therefore, it cannot be performed on already trained networks.

To have a closer look at what happens when we compare the best performing approaches at a more fine-grained level, we present the results on confidence and predictive entropy in the following density histograms. We notice how the trust concern in the overconfident predictions of the conventional DNN BackboneDNN is valid for all three datasets (see Figure 4). Even if the model has high accuracy, it loses its credibility when the confidence is so high for false predictions. We see a peak near higher confidence values for true predictions. However, most importantly, we want lower confidence values for false predictions. Our approach conforms to this as the best performing approach especially in the HHAR dataset. ApDeepSense performs worse, as expected, given its limitations in relying only on MLPs and moreover it is not enabled to work well for classification.

In Figure 5, the density histograms illustrate the predictive entropy as uncertainty estimate. Predictive entropy embodies the average amount of information in the predictive distribution. Given a data sample, the predictive entropy attains its maximum value when all classes are predicted to have equal uniform probability (therefore, the model is highly uncertain), and its minimum value

of zero when one class has probability 1 and all others probability 0 (i.e. the prediction is certain). We compare our approach versus MCDrop-30 and Ensemble-10, the best performing so far and best SOTA baselines. Although the three techniques can provide uncertainty, ours performs better than Ensemble-10 on all datasets and better than MCDrop-30 on the HHAR. Our results are similar to MCDrop-30 on the other two datasets, however, we obtain the illustrated uncertainty estimates in one single forward pass with great computational advantage and only a slight computation overhead compared to vanilla DNNs (see Section 5.6). Our approach performs better than the others in this respect since it assigns lower confidence to wrong predictions. Where our technique performs similarly to MCDrop or Ensemble, we still benefit from a significant computational gain.

The ideal confidence distribution would have higher density towards 1.0 for true predictions and towards 0.0 for false predictions. This might be impossible in practice because e.g. a correct prediction can be influenced by noise in input, resulting in low confidence. In our speech command task, although we include different background noises in training, there can still be other types of noise not seen in the training data which will cause uncertainty. The uncertainty in this prediction would trigger a reaction from the application which might involve recording again the command and, ultimately, collect these noisy data to include in the next training cycle in the cloud or smart device for personalization.

In the human activity recognition task, we see two of the prediction classes (stairs down and stairs up) which are occasionally confused by the model reflected in either in a true prediction with high uncertainty or wrong prediction with low uncertainty. While this is an undesired outcome, as the model should be well calibrated to handle these predictions, having the measure of uncertainty gives an indication that for these classes the models should be engineered to extract more meaningful features. This glimpse of interpretability allows to make decisions on the model architecture and how to process it in a federated system e.g., i.e. the user is given momentarily the parent class prediction (stairs) and then it is refined to a specific stairs up or down after further computation and higher confidence.

While how the uncertainty is interpreted and the reaction to this trigger is application-specific, we provide an efficient framework that generalizes to different application allowing to dive deeper in the deep learn model reasoning improving the overall system accuracy.

5.6 Latency and Energy Consumption

Figure 6 shows the latency (inference time) and energy consumption for all datasets. The experiments are conducted on an Nvidia Jetson TX2 and Nano. For fairness, we measure only the time needed to pass a sample and not consider the time needed to upload the model. In most of the considered baselines including ours, the model is uploaded just once and kept in memory. However, for ensembles, this might be different depending on the capacity and scenario, therefore we decide not to add that time and computation to the results. As we can see, our approach adds only a slight (max 20%) overhead over the conventional BackboneDNN, while being able to provide uncertainty estimates. Our latency is around 9-19ms

per inference depending on the dataset and the edge platform. The latency of MCDrop is significantly worse and the time to perform inference increases with the number of runs (forward passes) being 20x times in the best scenario compared to our method and increasingly more in other cases (up to 28-fold). Similar trends can be observed for Ensembles (2x - 8x times) as they require running multiple neural networks.

The energy consumption measurements show similar patterns. Ours is also at least (20%) faster compared to one of the most recent approaches (ApDeepSense). In general, our method always requires less energy than all the other approaches and adds only a negligible or a tiny overhead (depending on the dataset) on the traditional DNN approach which does not provide uncertainty estimates. Noticeably, our approach performs well on Nano, highlighting the fact that the applications can harness the utility of reliable predictions on many modern mobile and embedded devices especially if latency could be slightly sacrificed (which is often the case for critical applications). Additionally, our CPU only results on Nano demonstrates that our framework can run efficiently on resource-constrained devices that do not have a GPU.

As mentioned before, we want to make sure that our models can have a small footprint on these devices. For both embedded platforms, our investigations show that we add only a negligible runtime memory overhead (max 5%) compared to the vanilla deep learning model, while improving on the MLP baseline by 30%. Deep ensembles start heavily using the memory swap on the Nano when passing 5 ensembles, therefore, there is more need for memory sharing mechanisms. MCDrop does not contribute extra to the memory but, of course, this technique relies on a lengthy computation time as seen in all the results in this section.

6 DISCUSSION AND CONCLUSIONS

We have introduced a framework able to provide uncertainty estimates through layerwise distribution approximations using only a single forward pass. It provides a simple and system efficient solution that empowers convolutional neural networks to generate uncertainty estimations during inference for classification tasks. We evaluated our approach on multiple mobile systems datasets and showed that it significantly outperforms state-of-the-art baselines in terms of computation, latency and energy expenditure while providing reliable uncertainty estimates and accurate predictions.

There are many interesting avenues we want to pursue further in the future. Our approach could be extended to recurrent neural networks. However, these alterations require additional effort in providing the right mathematical foundations and test its feasibility on real-life datasets.

A key advantage of our framework is the fact that we model each layer to output predictive distributions. This can be a very flexible instrument for many applications, e.g., early prediction models, a class of conditional computation models that exit once a criterion (e.g., sufficient accuracy and low uncertainty) is satisfied at early layers. Such models can be very useful in intermittent learning systems [35, 40] which are powered by harvested energy.

Our approach, based on Gaussian approximations to the internal statistics of the network, is a feasible solution to providing uncertainty estimates on edge devices. These platforms cannot afford an

increase in latency, memory or energy due to additional forward passes but that would benefit from uncertainty quantification. It is also a powerful solution considering that the approach does not require re-training or fine-tuning, however, it comes with some limitations too. Our approximation could be improved to consider the fact that even if the inputs were perfectly Gaussian, the outputs of non-linear activations naturally yield skewed distributions with values possibly in a limited subset of the domain, which are not perfectly Gaussian distributed. Future work, therefore, could consider other approximations that minimize this discrepancy while simultaneously yielding uncertainty estimates for existing neural networks without increased operational costs, as is the case with our approach.

The probabilistic baselines (MCDrop and Ensembles), due to their nature, can significantly improve their performance by continuously increasing the number of samples or models to apply but with a huge computational overhead. Instead, our approach considers hardware limitations by design by giving the uncertainty in one single forward pass, therefore its performance improvements are limited to these constraints.

Our approach includes the limitations of using Multivariate Gaussian distribution which are common to Bayesian approximations using variational inference methods in general. For example the problem of under-estimation of the variance of the posterior distribution or the fact that the limited capacity of the posterior approximation can also result in biases in model parameters (and this is the case e.g., in time-series models).

To conclude, uncertainty estimations bring the much required element of interpretability and reasoning on the predictions made by neural network models. Such estimates are vital in the area of mobile and embedded systems as these systems deal with different kind of uncertainties. We have offered an avenue to provide them cheaply on these platforms while maintaining the needed level of performance.

7 ACKNOWLEDGMENTS

This work is supported by Nokia Bell Labs through their donation for the Centre of Mobile, Wearable Systems and Augmented Intelligence and by ERC Project 833296 (EAR).

REFERENCES

- [1] [n.d.]. HHAR, Heterogeneous human activity recognition. <https://archive.ics.uci.edu/ml/datasets/Heterogeneity+Activity+Recognition>.
- [2] [n.d.]. jetson. <https://developer.nvidia.com/embedded/jetson-tx2>.
- [3] [n.d.]. nano. <https://developer.nvidia.com/embedded/jetson-nano-developer-kit>.
- [4] [n.d.]. speech-commands. <https://ai.googleblog.com/2017/08/launching-speech-commands-dataset.html>.
- [5] [n.d.]. Texas Instruments INA226. <https://www.ti.com/product/INA226>.
- [6] Sungjin Ahn, Anoop Korattikara, and Max Welling. 2012. Bayesian posterior sampling via stochastic gradient Fisher scoring. *arXiv preprint arXiv:1206.6380* (2012).
- [7] Brandon Ballinger, Johnson Hsieh, Avesh Singh, Nimit Sohoni, Jack Wang, Geoffrey H Tison, Gregory M Marcus, Jose M Sanchez, Carol Maguire, Jeffrey E Olgin, et al. 2018. DeepHeart: semi-supervised sequence learning for cardiovascular risk prediction. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [8] Sourav Bhattacharya and Nicholas D Lane. 2016. Sparsification and separation of deep learning layers for constrained resource inference on wearables. In *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM*. ACM, 176–189.
- [9] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. 2015. Weight uncertainty in neural networks. *arXiv preprint arXiv:1505.05424* (2015).
- [10] Han Cai, Chuang Gan, and Song Han. 2019. Once for all: Train one network and specialize it for efficient deployment. *arXiv preprint arXiv:1908.09791* (2019).
- [11] Justin Chan, Thomas Rea, Shyamnath Gollakota, and Jacob E Sunshine. 2019. Contactless Cardiac Arrest Detection Using Smart Devices. *arXiv preprint arXiv:1902.00062* (2019).
- [12] Jagmohan Chauhan, Jathushan Rajasegaran, Suranga Seneviratne, Archan Misra, Aruna Seneviratne, and Youngki Lee. 2018. Performance Characterization of Deep Learning Models for Breathing-based Authentication on Resource-Constrained Devices. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 4 (2018), 1–24.
- [13] Ricardo Chavarriaga, Hesam Sagha, Alberto Calatroni, Sundara Tejaswi Digmarti, Gerhard Tröster, José del R Millán, and Daniel Roggen. 2013. The opportunity challenge: A benchmark database for on-body sensor-based activity recognition. *Pattern Recognition Letters* 34, 15 (2013), 2033–2042.
- [14] Andreas Damianou and Neil Lawrence. 2013. Deep gaussian processes. In *Artificial Intelligence and Statistics*. 207–215.
- [15] Sajal K Das and Christopher Rose. 2005. Coping with uncertainty in mobile wireless networks. In *Emerging Location Aware Broadband Wireless Ad Hoc Networks*. Springer, 189–204.
- [16] Toon De Pessemier and Luc Martens. 2018. Heart rate monitoring, activity recognition, and recommendation for e-coaching. *Multimedia Tools and Applications* 77, 18 (2018), 23317–23334.
- [17] Armen Der Kiureghian and Ove Ditlevsen. 2009. Aleatory or epistemic? Does it matter? *Structural safety* 31, 2 (2009), 105–112.
- [18] Davide Figo, Pedro C Diniz, Diogo R Ferreira, and João M Cardoso. 2010. Preprocessing techniques for context recognition from accelerometer data. *Personal and Ubiquitous Computing* 14, 7 (2010), 645–662.
- [19] Yarin Gal and Zoubin Ghahramani. 2015. Bayesian convolutional neural networks with Bernoulli approximate variational inference. *arXiv preprint arXiv:1506.02158* (2015).
- [20] Yarin Gal and Zoubin Ghahramani. 2016. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*. 1050–1059.
- [21] Yarin Gal and Zoubin Ghahramani. 2016. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in neural information processing systems*. 1019–1027.
- [22] Petko Georgiev, Sourav Bhattacharya, Nicholas D Lane, and Cecilia Mascolo. 2017. Low-resource multi-task audio sensing for mobile and embedded devices via shared deep neural network representations. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 1–19.
- [23] Petko Georgiev, Nicholas D Lane, Cecilia Mascolo, and David Chu. 2017. Accelerating mobile audio sensing algorithms through on-chip gpu offloading. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*. 306–318.
- [24] Zoubin Ghahramani. 2015. Probabilistic machine learning and artificial intelligence. *Nature* 521, 7553 (2015), 452.
- [25] Wenbo Gong, Yingzhen Li, and José Miguel Hernández-Lobato. 2018. Meta-learning for stochastic gradient MCMC. *arXiv preprint arXiv:1806.04522* (2018).
- [26] Alex Graves. 2011. Practical variational inference for neural networks. In *Advances in neural information processing systems*. 2348–2356.
- [27] Nils Y Hammerla, Shane Halloran, and Thomas Plötz. 2016. Deep, convolutional, and recurrent models for human activity recognition using wearables. *arXiv preprint arXiv:1604.08880* (2016).
- [28] Geoffrey E Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R Salakhutdinov. 2012. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580* (2012).
- [29] Shih-Chung Hsu, Cheng-Hung Chuang, Chung-Lin Huang, Ren Teng, and Miao-Jian Lin. 2018. A video-based abnormal human behavior detection for psychiatric patient monitoring. In *2018 International Workshop on Advanced Image Technology (IWAIT)*. IEEE, 1–4.
- [30] Md Abdullah Al Hafiz Khan, Nirmalya Roy, and Archan Misra. 2018. Scaling human activity recognition via deep learning-based domain adaptation. In *2018 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 1–9.
- [31] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [32] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*. 6402–6413.
- [33] Nicholas D Lane, Sourav Bhattacharya, Petko Georgiev, Claudio Forlivesi, Lei Jiao, Lorena Qendro, and Fahim Kawsar. 2016. DeepX: A software accelerator for low-power deep learning inference on mobile devices. In *Proceedings of the 15th International Conference on Information Processing in Sensor Networks*. IEEE Press, 23.
- [34] Nicholas D Lane, Petko Georgiev, and Lorena Qendro. 2015. DeepEar: robust smartphone audio sensing in unconstrained acoustic environments using deep learning. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 283–294.

- [35] Seulki Lee and Shahriar Nirjon. 2019. Neuro. ZERO: a zero-energy neural network accelerator for embedded sensing and inference systems. In *Proceedings of the 17th Conference on Embedded Networked Sensor Systems*. ACM, 138–152.
- [36] Ji Lin, Yongming Rao, Jiwon Lu, and Jie Zhou. 2017. Runtime neural pruning. In *Advances in Neural Information Processing Systems*. 2181–2191.
- [37] Hong Lu, Denise Frauendorfer, Mashfiqui Rabbi, Marianne Schmid Mast, Gokul T Chittaranjan, Andrew T Campbell, Daniel Gatica-Perez, and Tanzeem Choudhury. 2012. Stresssense: Detecting stress in unconstrained acoustic environments using smartphones. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. ACM, 351–360.
- [38] David JC MacKay. 1992. *Bayesian methods for adaptive models*. Ph.D. Dissertation. California Institute of Technology.
- [39] Akhil Mathur, Anton Isopoussu, Fahim Kawsar, Nadia Berthouze, and Nicholas D Lane. 2019. Mic2Mic: using cycle-consistent generative adversarial networks to overcome microphone variability in speech systems. In *Proceedings of the 18th International Conference on Information Processing in Sensor Networks*. ACM, 169–180.
- [40] Alessandro Montanari, Manuja Sharma, Dainius Jenkus, Mohammed Alloula, Lorena Qendro, and Fahim Kawsar. 2020. ePerceptive: energy reactive embedded intelligence for batteryless sensors. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*. 382–394.
- [41] Liangying Peng, Ling Chen, Zhenan Ye, and Yi Zhang. 2018. AROMA: A Deep Multi-Task Learning Based Simple and Complex Human Activity Recognition Method Using Wearable Sensors. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 2 (2018), 74.
- [42] Karol J Piczak. 2015. ESC: Dataset for environmental sound classification. In *Proceedings of the 23rd ACM international conference on Multimedia*. 1015–1018.
- [43] Joaquin Quinero-Candela, Carl Edward Rasmussen, Fabian Sinz, Olivier Bousquet, and Bernhard Schölkopf. 2005. Evaluating predictive uncertainty challenge. In *Machine Learning Challenges Workshop*. Springer, 1–27.
- [44] Valentin Radu, Catherine Tong, Sourav Bhattacharya, Nicholas D Lane, Cecilia Mascolo, Mahesh K Marina, and Fahim Kawsar. 2018. Multimodal deep learning for activity and context recognition. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 4 (2018), 157.
- [45] Daniele Ravi, Charence Wong, Benny Lo, and Guang-Zhong Yang. 2016. Deep learning for human activity recognition: A resource efficient implementation on low-power devices. In *2016 IEEE 13th International Conference on Wearable and Implantable Body Sensor Networks (BSN)*. IEEE, 71–76.
- [46] Charissa Ann Ronao and Sung-Bae Cho. 2016. Human activity recognition with smartphone sensors using deep learning neural networks. *Expert systems with applications* 59 (2016), 235–244.
- [47] Jason Ryder, Brent Longstaff, Sasank Reddy, and Deborah Estrin. 2009. Ambulation: A tool for monitoring mobility patterns over time using mobile phones. In *2009 International Conference on Computational Science and Engineering*, Vol. 4. IEEE, 927–931.
- [48] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research* 15, 1 (2014), 1929–1958.
- [49] Allan Stisen, Henrik Blunck, Sourav Bhattacharya, Thor Siiger Prentow, Mikkel Baun Kjærgaard, Anind Dey, Tobias Sonne, and Mads Møller Jensen. 2015. Smart devices are different: Assessing and mitigating mobile sensing heterogeneities for activity recognition. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*. ACM, 127–140.
- [50] Abdulhamit Subasi, Mariam Radhwan, Rabea Kurdi, and Kholoud Khateeb. 2018. IoT based mobile healthcare system for human activity recognition. In *2018 15th Learning and Technology Conference (L&T)*. IEEE, 29–34.
- [51] Terry Taewoong Um, Franz Michael Josef Pfister, Daniel Pichler, Satoshi Endo, Muriel Lang, Sandra Hirche, Urban Fietzek, and Dana Kulić. 2017. Data augmentation of wearable sensor data for parkinson’s disease monitoring using convolutional neural networks. *arXiv preprint arXiv:1706.00527* (2017).
- [52] Yan Wang, Shuang Cang, and Hongnian Yu. 2018. A data fusion-based hybrid sensory system for older people’s daily activity and daily routine recognition. *IEEE Sensors Journal* 18, 16 (2018), 6874–6888.
- [53] Max Welling and Yee W Teh. 2011. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*. 681–688.
- [54] Tien-Ju Yang, Andrew Howard, Bo Chen, Xiao Zhang, Alec Go, Mark Sandler, Vivienne Sze, and Hartwig Adam. 2018. NetAdapt: Platform-Aware Neural Network Adaptation for Mobile Applications. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- [55] Shuochao Yao, Shaohan Hu, Yiran Zhao, Aston Zhang, and Tarek Abdelzaher. 2017. Deepsense: A unified deep learning framework for time-series mobile sensing data processing. In *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 351–360.
- [56] Shuochao Yao, Yiran Zhao, Huajie Shao, Aston Zhang, Chao Zhang, Shen Li, and Tarek Abdelzaher. 2018. Rdeepsense: Reliable deep mobile computing models with uncertainty estimations. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 4 (2018), 173.
- [57] Shuochao Yao, Yiran Zhao, Huajie Shao, Chao Zhang, Aston Zhang, Dongxin Liu, Shengzhong Liu, Lu Su, and Tarek Abdelzaher. 2018. Apdeepsense: Deep learning uncertainty estimation without the pain for iot applications. In *2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 334–343.