Uncertainty Aware Mobile Sensing

Lorena Qendro
University of Cambridge
lq223@cam.ac.uk

Jagmohan Chauhan
University of Cambridge
jc2161@cam.ac.uk

Cecilia Mascolo
University of Cambridge
cm542@cam.ac.uk

1. INTRODUCTION

Deep neural networks (DNN) are state-of-the-art machine learning models for mobile and ubiquitous sensing applications such as activity and context recognition, monitoring health and well-being, and location prediction. The key to their success relies on their structure of multiple non-linear hidden layers which makes them very expressive models able to learn complex intricate relationships between the sensory inputs and predicted outputs. However, DNNs lack interpretability which does not allow a proper understanding of how the final predictions were deduced.

Many of the aforementioned mobile sensing applications have a significant role in automated decision making required in the healthcare, autonomous vehicles and life or mission-critical systems. A DNN model in such a system, when encountered with examples that are outside its data distribution, can make irrational decisions or suggestions leading to severe consequences. Having an understanding of the model's confidence gives a useful indication on when the system is guessing at random or is confident about the prediction.

Bayesian Neural Network (BNN) builds a bridge between DNNs and probabilistic Bayesian approach by bringing together the representational flexibility of DNNs with principled parameter estimation of probabilistic models. The Bayesian framework offers uncertainty estimations on predictions made by DNNs which helps in capturing the erroneous overconfident decisions in case of out-of-distribution or noisy data.

2. PROBLEM STATEMENT

Despite the advantages of BNNs, their training and inference is computationally very expensive. Hence running them on devices with limited resources such as smartphones and wearables is a big challenge.

Latest studies from the machine learning community offer enlightening directions on getting robust uncertainty estimations for DNNs but they require multiple iterative runs increasing the computation of training a single model multiple times. Gal et al.[1] introduce Monte Carlo dropout (MCDrop) interpreting dropout[2] to approximately correspond to variational inference. MCDrop collects the results of stochastic forward passes through the model and estimates the predictive uncertainty. As a result, this information can be used with existing NN models trained with dropout. Although this technique does not require retraining of the model it still does not manifest a resource-friendly approach since it requires the stochastic NN to run multiple times in order to average the results. Yao et al.[3] propose an approximation of the output distribution at each NN layer which reduces computation time and energy consumption. However, the authors approach only applies to fully connected NNs leaving the challenges for Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) still open. Capturing temporal dynamics and working with sequential data is an important aspect of many mobile sensing applications and therefore tackling these challenges is a crucial part of our work.

3. PRELIMINARY WORK & DISCUSSION

In this work, we plan to devise and evaluate various techniques to achieve uncertainty aware mobile sensing. The aim is to further develop the state-of-the-art approaches to include uncertainty estimations in CNNs and RNNs in a resource-friendly manner which can execute efficiently on mobile devices and compare the new approaches with the existing baselines (proposed by Yao et al.[3]) in terms of accuracy, execution time and energy efficiency. Our approach includes incorporating layer-wise approximations as well as proposing different NN architectures to guarantee a reasonable trade-off between power consumption and performance.

4. REFERENCES

[1] Gal et al. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. *CoRR*, abs/1506.02142, 2015.

[2] Srivastava et al. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.

[3] Yao et al. Apdeepsense: Deep learning uncertainty estimation without the pain for iot applications. In *38th IEEE ICDCS 2018*, pages 334–343.