

Received 10 July 2024, accepted 30 July 2024, date of publication 7 August 2024, date of current version 19 August 2024.

Digital Object Identifier 10.1109/ACCESS.2024.3439850

RESEARCH ARTICLE

Conformal Prediction Based Confidence for Latency Estimation of DNN Accelerators: A Black-Box Approach

MATTHIAS WESS¹, DANIEL SCHNÖLL¹, (Graduate Student Member, IEEE),
DOMINIK DALLINGER¹, (Graduate Student Member, IEEE), MATTHIAS BITTNER¹,
AND AXEL JANTSCH¹, (Fellow, IEEE)

¹Institute of Computer Technology, TU Wien, 1040 Vienna, Austria

Christian Doppler Laboratory for Embedded Machine Learning, Institute of Computer Technology, TU Wien, 1040 Vienna, Austria

Corresponding author: Matthias Wess (matthias.wess@tuwien.ac.at)

This work was supported in part by Austrian Federal Ministry for Digital and Economic Affairs; in part by the National Foundation for Research, Technology and Development; in part by the Christian Doppler Research Association; and in part by Technical University (TU) Wien Bibliothek.

ABSTRACT Today, there exists a large number of different embedded hardware platforms for accelerating the inference of Deep Neural Networks (DNNs). To enable rapid application development, a number of prediction frameworks have been proposed to estimate the DNN inference latency on a wide range of hardware platforms. This work presents a novel smart padding benchmarking method, which allows the profiling of hardware platforms without requiring detailed per-layer reports. To mitigate the measurement inaccuracies inherent in the black-box approach, a confidence framework comprising three metrics has been developed. These metrics not only enhance the interpretation of prediction results but also significantly contribute to the refinement of the estimation framework itself, as they facilitate to improve the coverage of the training dataset for relevant layers and detect weaknesses in the training dataset. Empirical results demonstrate the method's robustness, with average prediction errors minimized to below 10% for smart padding benchmarking-based ANNETTE predictions for the Jetson Xavier, NXP i.MX93, and NXP i.MX8M+.

INDEX TERMS Estimation, latency, confidence, neural network hardware, conformal prediction.

I. INTRODUCTION

The vast design space of optimization, pruning, quantization and mapping DNNs on embedded hardware platforms makes it almost impossible to quickly find the best fitting solution for an application. Neural Architecture Search (NAS) [1], [2] provides a means to achieve a DNN optimized with regards to certain requirements. Specifically in hardware-aware NAS the inference latency is often used as the target constraint and therefore needs to be computed or measured for each selected DNN architecture. To avoid the need to deploy each DNN on the requested platforms, various approaches have been proposed to predict the inference latency. Solutions to this problem range from the use of simple proxy metrics

The associate editor coordinating the review of this manuscript and approving it for publication was Pinjia Zhang¹.

(such as the number of floating point operations) [3] and analytical models [4], [5] to Graph Convolutional Networks (GCNs) [6], [7]. Some solutions focus on specific design spaces to enable hardware-aware NAS and therefore provide limited generalization capabilities. Other methods (e.g. ANNETTE [5], nn-Meter [8]) aim to provide accurate predictions for a wide range of applications and cover the aspects of graph optimizations in a separate step to correctly model all steps in the deployment flow.

However, the vast amount of different hardware platforms available for DNN inference makes the adaption of estimation algorithms for each hardware platform cumbersome. Therefore, Metrics that provide additional information.

Our goal is to address two challenges related to benchmarking and predicting the inference time of neural networks on constrained devices. First, benchmarking hardware platforms

for specific DNNs is challenging due to layer fusion, dependencies on layer sequences, data loading effects, interference of profiling techniques with execution time, and other complications. Additionally, it is important to gain insights not only for entire networks but also at the per-layer level. Currently, achieving this level of detail necessitates the use of per-layer profiling results to accurately model execution time. However, there are situations where implementing per-layer profiling is not feasible or requires additional implementation effort and possibly generates additional profiling overhead. We tackle this challenge by developing an intelligent benchmarking strategy that allows for the generation of per-layer abstraction models without relying on detailed insights.

Second, a latency estimate is only useful to designers if they know to which extent it can be trusted. How can we ensure the comparability of models and how can we trust models trained on limited data points? To address this concern, we propose three novel confidence metrics. These metrics provide quantitative measures of the reliability of our latency prediction models, enabling informed decision-making when selecting hardware platforms and DNN architectures for the application-specific DNN hardware implementation. Additional applications of latency prediction, such as hardware aware DNN compression [9], [10] and DNN offloading and partitioning [11], [12] can also potentially make use those prediction reliability measures.

Specifically, this paper makes the following contributions:

- We propose a method for profiling the latency of DNN inference on hardware with padded models;
- We propose a conformal prediction framework for DNN latency prediction to quantify the confidence of the predicted values.

II. RELATED WORK

A. LATENCY PREDICTION

The goal of latency prediction is to estimate the total execution time of a network composed of a sequence of N layers $L = \{l_1, l_2, \dots, l_N\}$. Each layer in the DNN has specific attributes and parameters that define its configuration, computation needs, and connections to other layers. These connections determine the data flow through the entire network. Current approaches for DNN latency prediction range from simple analytical models based on the roofline model [5] to elaborate Machine Learning (ML) based latency estimators [13]. These ML based prediction algorithms are trained on collected datasets $Z = \{(\vec{x}_1, y_1), (\vec{x}_2, y_2), \dots\}$, where \vec{x}_i are the feature vectors, describing layer i , and y_i are the values to be predicted. In the case of latency estimation, the target values can represent for example time or compute efficiency. As a result, ML based prediction algorithms are not limited to a specific hardware platform. They show good accuracy [14], [15], [16] but are mostly limited to the selected design space and are usually not designed for general network prediction.

Analytical prediction methods such as those presented in [17] and [18] provide high prediction accuracy for the target hardware platforms. However, they require in-depth hardware knowledge and are therefore not suitable when in-depth architecture details cannot be obtained due to confidentiality or when the required effort is excessive.

The latency prediction framework Blackthorn [4] encompasses analytical models constructed based on several measurement points. The focus of Blackthorn is on finding optimal measurement points to reduce the required amount of overall measurements to profile NVIDIA Graphic Processing Units (GPUs).

The framework ANNETTE [5] provides analytical models based on a refinement of the roofline model which, in addition to the compute and memory boundary, also takes into account the underlying compute architecture. In addition, ANNETTE relies on random forest regression models predicting the per-operator compute efficiency and also deploys decision trees to predict operator fusion rules.

Other similar approaches with iterative improvements and slightly different focus with regard to the profiled hardware [19], [20] have been proposed. nn-Meter [8] focuses on the prediction of mobile devices and deploys similar principles as ANNETTE relying on a larger training dataset. MAPLE-X [21] incorporates explicit prior knowledge of hardware devices to improve the prediction accuracy for newly benchmarked devices.

Finally, Graph Neural Networks (GNNs) offer the option to operate directly on the graph structure of the DNN to be predicted. Sectum [22] deploys a GNN to detect memory over-commitment in addition to an ANNETTE-like structure. While frameworks like DNNPerf [13] and GENNAPE [23] focus on the prediction of other DNN performance parameters (such as accuracy, training time, etc.), PerfSAGE [7] and DIPPM [24] rely on GNNs to predict latency, energy, and memory consumption and promise high prediction accuracy for different classes of network architectures. In both cases, the GraphSAGE architecture is deployed in different variants. Lastly, SLAPP [6] applies GNNs at sub-graph level to preserve the advantage of gained insights through per-operator prediction but still relying on a large number of data points.

The black-box approach using smart padding, presented in this work, can be a valuable method for most of the ML based latency estimation frameworks. Even though the technique does not replace detailed per-layer profiles, it enables in-model latency measurement of single layers or blocks of layers while decreasing the required effort for implementing overhead-free per-layer profiling tools.

B. CONFORMAL PREDICTION

The conformal prediction framework, introduced by Vovk, Grammerman and Shafer [25], [26] provides a general method for quantifying the uncertainty of predictions for arbitrary prediction algorithms and provides guarantees on

the prediction error. Traditionally, confidence intervals are estimated using quantile regression [27], [28] or Bayesian methods [29]. In the context of this work, which leverages random forest regression [30], conformal prediction is particularly beneficial for uncertainty quantification, as it not only demonstrates good efficiency [31] but also ensures broad applicability across different machine learning algorithms. Furthermore, conformal prediction offers many additional advantages, such as its straightforward interpretability, model-agnostic nature, and adaptability.

For uncertainty quantification, conformal prediction relies on the computation of a nonconformity score α for each instance in a calibration set distinct from the initial training set. In regression, α is typically computed as the absolute error $\alpha_j = |\hat{y}_j - y_j|$ [31], where \hat{y} is the predicted value. For the prediction of confidence intervals with significance δ , these calculated nonconformity scores are used to formulate the prediction region for each instance j as $\hat{Y}_j^\delta = \hat{y}_j \pm \alpha(\delta)$ [31]. This means the predicted region \hat{Y} will cover the true value of y with probability $p = 1 - \delta$. In the standard case, this results in confidence intervals of uniform width across all input feature vectors \vec{x} .

Thus, to minimize the average interval width, it is possible to implement normalized nonconformity functions [32]. Here, the nonconformity scores are scaled by σ , an estimate of the model's accuracy for the predicted instance. The resulting prediction regions are then computed as $\hat{Y}_j^\delta = \hat{y}_j \pm \alpha(\delta) \cdot \sigma_j$. This quality estimate can be obtained by various methods, such as predicting the errors with additional trained models or using the errors of the k nearest neighbors. Conformal prediction has been successfully applied in various domains, including medical diagnosis [33], face recognition [34], and financial risk prediction [35]. However, to our knowledge, this work presents the first approach to leverage conformal prediction for confidence estimation in latency estimation of DNNs.

III. METHODOLOGY

Currently, for latency estimation of DNN hardware accelerators, we encounter two primary challenges:

1. Across the broad spectrum of available DNN accelerators the availability of knowledge, insight, and profiling tools varies widely. This diversity necessitates tailored benchmarking and modeling approaches for each type.
2. The accuracy of latency prediction models varies widely due to variations in benchmarking methodologies, dataset size (e.g. limited due to the compilation time), and hardware architectures. These issues compromise the reliability of latency estimates and affect the coverage of the DNN design space.

The following sections address the identified challenges in estimating latency for DNN hardware. Section III-A provides an overview of the model generation process, highlighting the additions to the latency estimation framework. To tackle issue (1), Section III-B presents a flexible methodology

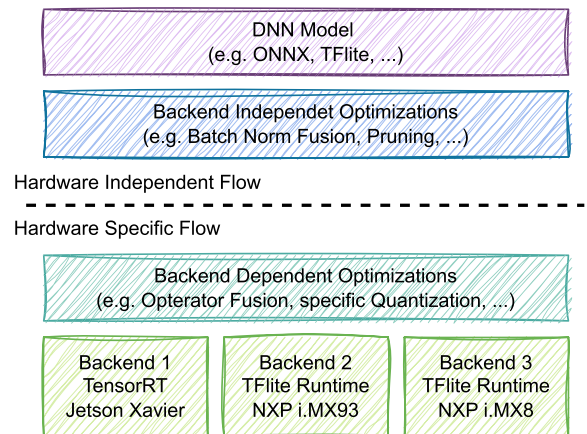


FIGURE 1. Overview of the compilation flow for inference on embedded hardware platforms.

that allows us to profile hardware platforms based on a minimal requirement on the available hardware insights and profiling possibilities. Lastly, to address the diverse latency prediction model quality (2), in Section III-C we propose the application of conformal prediction methods as measures for the confidence of the per-layer and per-network estimation.

A. OVERVIEW

Figure 1 depicts the usual stages to compile a trained Neural Network (NN) for hardware inference:

- The trained **DNN model** is exported from a training framework such as Tensorflow or Pytorch to an intermediate exchange format (e.g. ONNX, TFLite).
- **Backend independent optimizations** are applied to optimize the graph for inference. These can include removing layers from the graph that are only required for training (e.g. Dropout), or fusing layers while still maintaining mathematical equivalency (e.g. Batch Normalization). While most inference frameworks apply this step automatically, it is still recommended to make use of tools such as NVIDIA's ONNX-GraphSurgeon¹ or ONNX-simplifier² in a separate step. Hence, this step can similarly be applied in the latency estimation flow.
- **Backend dependent optimizations** represent the changes applied to the DNN model, that are either required or beneficial with regard to latency and/or efficiency, and which are not executable on all hardware platforms. Since each hardware platform provides a different set of operations and possibly allows for multiple operations in a pipelined manner (*composite layers*) to reduce data transfer these optimizations need to be considered in the estimation framework [5], [8].
- Lastly, the model is compiled and executed on the hardware platform using the hardware-specific inference **backend**. Some compilers provide different

¹<https://docs.nvidia.com/deeplearning/tensorrt/onnx-graphsurgeon>

²<https://github.com/daquexian/onnx-simplifier>

optimization targets (e.g. latency, memory) or optimize the workload for a specific hardware setting. It has to be considered that, with the current methods, each prediction model can only provide the predictions for one specific combination of compiler and hardware settings.

From the inference workflow, there are different levels of insights that can be gathered and used for the latency estimation framework:

- DNN graph before and after backend-dependent optimization
- Per layer latencies
- Overall network latency

Ideally, these insights not only include the hardware mapping of the computational graph but also precise timing for each layer. This would allow for the development of accurate latency estimation models and the identification of further optimizations, such as combining individual layers into *composite layers*. In this context, *Composite layers* refer to the fusion of multiple neural network operations (e.g., Conv2D + ReLU + MaxPool) into a single operation executed as one unit, enhancing processing efficiency and reducing latency.

However, the level of detail available in profiling data can vary significantly across different hardware platforms. Some allow for more detailed analysis than others. Additionally, the generation of per-layer reports can also lead to additional overhead resulting in inaccurate latency measurements. In cases where direct profiling at this level is not feasible, alternative methods, like employing GNNs for overall network latency estimation or block-wise estimation [15], have been explored. However, these approaches do not provide insights at the layer-level and are limited in their coverage of the design space, as they cannot account for all possible blocks and network configurations in the training dataset. The experiments conducted in this study demonstrate that simply benchmarking each layer type through single-layer measurements (profiling NNs consisting of only one layer) does not yield the required level of measurement accuracy. This is due to the overhead associated with data transfer at the start and end of the execution.

While each hardware architecture presents unique complexities, the smart padding method introduced in this work enables benchmarking across a wide array of current hardware platforms. This technique accounts for the data transfer overhead during the measurement process, thus isolating the actual computation time within a DNN with multiple layers. However, it operates under the assumption that the hardware platform performs computations on a per-layer basis. This assumption aligns with the operational characteristics of most modern hardware architectures, where parallel execution of layers typically does not yield substantial performance gains.

As an example, when taking a closer look at the block diagram of an ARM Ethos Neural Processing Unit (NPU) (see Figure 2) and the attached main components of

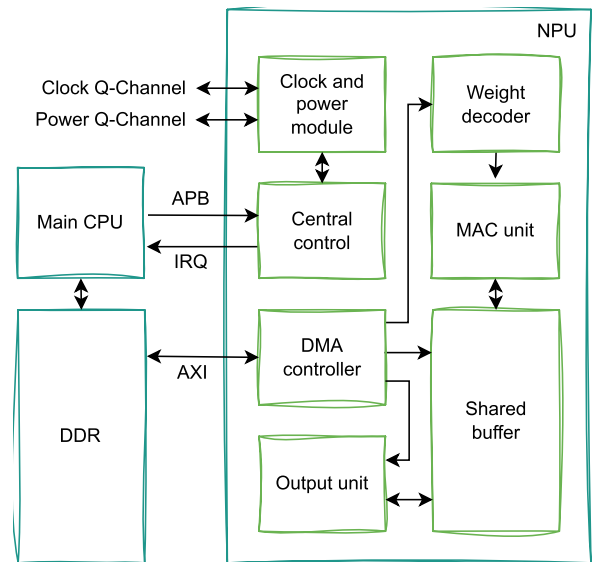


FIGURE 2. Blockdiagram of the ARM Ethos NPU [36] and connections to the main CPU and DDR.

the NXP i.MX93 (Main CPU and DDR Memory), we can gain insight into the underlying cause of this overhead. During the computation of the NN the intermediate feature maps are stored in a shared buffer, which is tightly coupled to the compute units. This setup enables fast data transfer and optimal compute efficiency. However, at the start and end of the NN inference, data must be transferred via the AXI-bus into or from this shared buffer. Additionally, potential data reordering or similar steps can further impede the speed of this process.

Considering the benchmarked hardware as a black-box, without in-depth knowledge of the specific relationships between the amount of transferred or processed data and the resulting latency, the developed methodology therefore needs to be able to account for this overhead. Furthermore, the implemented confidence metrics should reflect the added estimation uncertainty stemming from the black-box approach.

For this work, we build on Accurate Neural Network Execution Time Estimation (ANNETTE) [5], an open-source framework for NN latency estimation on embedded hardware platforms. Figure 3 provides an overview of the modules of ANNETTE and the components that are added for this work.

The ANNETTE workflow comprises two phases: the characterization phase and the estimation phase. Initially, in the characterization phase, **Benchmark Tool** (Fig. 3) executes the benchmarks on the hardware, by autonomously measuring the latencies for a set of parametric dummy network models and stores the results in a data frame. Subsequently, the **Model Generator** utilizes this data to generate prediction models for the assessed layer types and fusion rules. Predominantly, the end user interacts with the **Estimation Tool**, which loads a DNN model description in ONNX format and predicts the latency

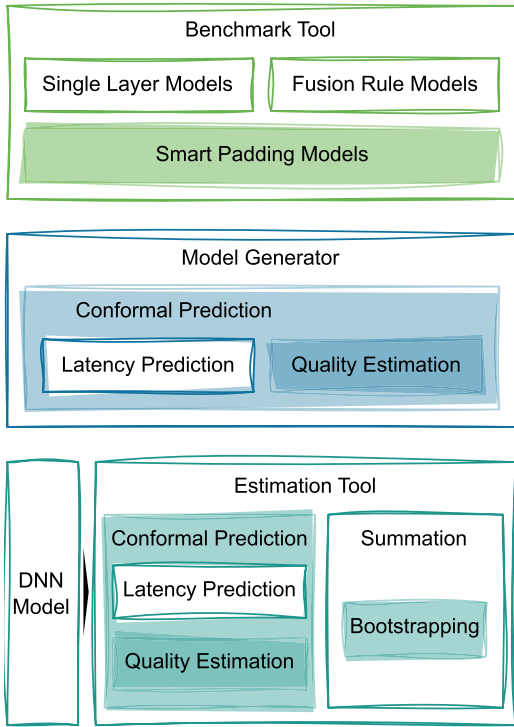


FIGURE 3. Overview of the components in ANNETTE. The color-shaded components are added in this work.

using the previously generated models. This work relies on the random forest-based estimation models of ANNETTE. However, it is possible to apply the same methodology to other latency estimation frameworks, such as nn-Meter or PerfSAGE since they are compatible with the conformal prediction approach [37].

To facilitate the proposed black-box benchmarking approach, the set of benchmarks is expanded with smart padding models (described in Section III-B). To enable uncertainty quantification for latency prediction, we apply conformal prediction methods to the random forest regression models (see Section III-C). This requires modifications to both the **Model Generator** and the **Estimation Tool** to support the conformal prediction framework. Specifically, the **Model Generator** is extended to include support for training the quality estimators and calculating non-conformity scores. Updates to the **Estimation Tool** enable the inference of conformal prediction, including the quality estimators, and the use of bootstrapping to compute the per-network confidence metrics.

B. BLACK-BOX BENCHMARKING

This section describes the techniques used to achieve per-layer prediction models for hardware with limited profiling capabilities. The notation used in this section is outlined in Table 1.

As mentioned in Section III-A, when measuring individual layers, there is additional overhead due to data transfer times,

TABLE 1. Definitions of time-related symbols.

Symbol	Definition
t	Latency
T	Latency Interval
$t_{\text{data_in}} / t_{\text{data_out}}$	Data transfer latency
t_{comp}	Computation latency
LUM	Layer Under Measurement
\hat{t}	Predicted Latency
\hat{T}	Predicted Latency Interval
\mathcal{T}	Measured Latency

complicating the accurate assessment of execution times. The measured execution time \mathcal{T} of a NN on hardware that executes layer by layer is determined by the computation time $t_{\text{comp},i}$ per layer, as well as the additional data transfer times $t_{\text{data_in}}$ and $t_{\text{data_out}}$.

$$\mathcal{T} = t_{\text{data_in}} + t_{\text{data_out}} + \sum_{i=1}^{\text{layers}} t_{\text{comp},i} \quad (1)$$

As a result, when benchmarking single-layer models based on the measured latency of the entire model, the estimator will overestimate the execution time of multi-layer models (see Section IV). Figure 4 illustrates a model with three layers. The effects of pipelining result in the overlaps of the compute and actual data read and data write times ($t'_{\text{data_in}}$ and $t'_{\text{data_out}}$) since the compute unit can start computation without having all the data available. While in most cases $t'_{\text{data_in}}$ and $t'_{\text{data_out}}$ are proportional to the amount of data to be transferred, due to the irregular pipelining effects, estimating $t_{\text{data_in}}$ and $t_{\text{data_out}}$ is more complicated and requires a different approach.




FIGURE 4. Relationship between measured, compute and data transfer times for a DNN with three layers.

At this point, the main challenge lies in disentangling the data read/write times from the computation latency of the Layer Under Measurement (LUM) t_{LUM} . At first glance, this task may seem straightforward; however, the intricate relationship between layer configuration and the dimensions of the resulting input and output feature maps requires a smart approach. To address this issue, we propose a smart padding

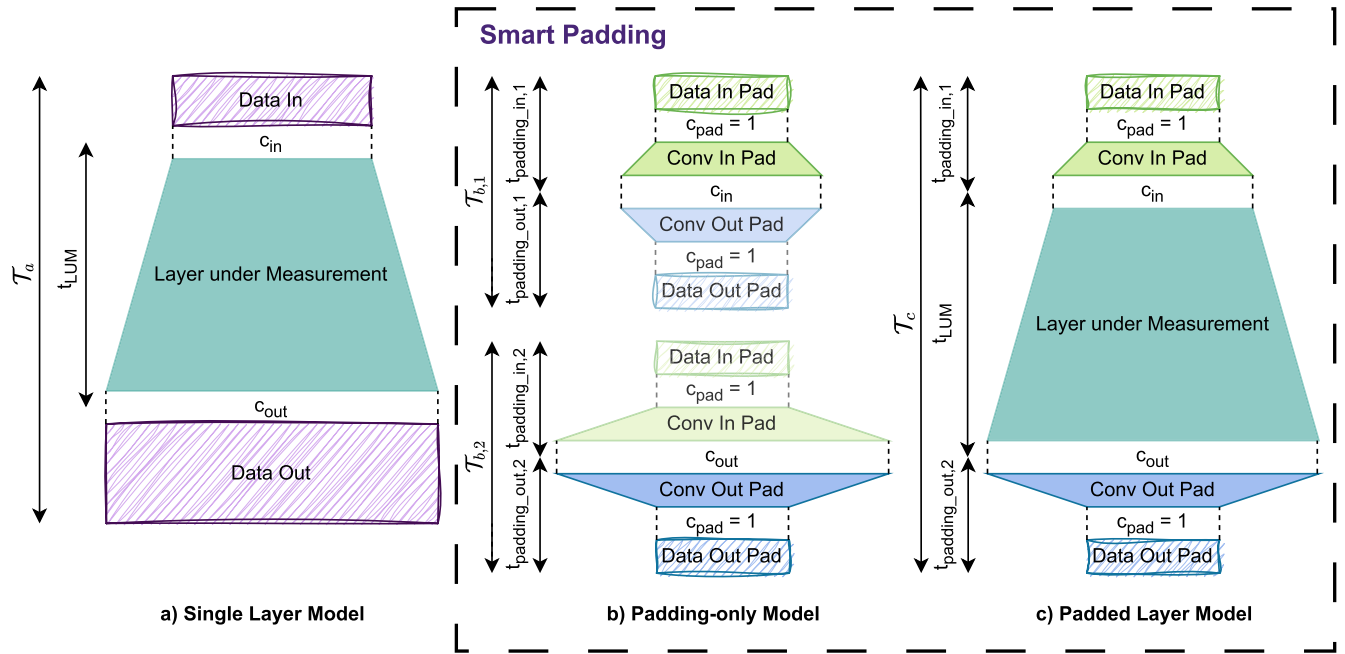


FIGURE 5. The three models used for benchmarking the different platforms. The Single Layer Model (a) is the simplest way but does not provide accurate measurements of t_{LUM} . The black-box benchmarking method makes use of padding-only (b) and padded layer models (c) to solve this problem.

strategy to measure the computation latency of the LUM within a multi-layer model. Figure 5 depicts the models for the smart padding strategy (Figure 5b,c), alongside a simple single-layer model (Figure 5a). Our strategy is based on two key concepts: firstly, to reduce data transfer times (t_{data_in} and t_{data_out}) to a bare minimum, thereby mitigating their impact on the latency measurements. Secondly, we independently measure the execution latency of a padding-only model. This enables the calculation of T_{LUM} by subtracting the times of the padding-only models (Figure 5b) from the padded layer model (Figure 5c).

According to Equation 1, the measured latency T_a for the single layer model includes the data transfer times (t_{data_in} , t_{data_out}) and the actual computation time t_{LUM} . The padded model consists of the LUM padded by an input and output padding 1×1 2D convolution layer with $c_{in}=1$ input channels, for the input padding layer and $c_{out}=1$ output channels for the output padding layer. However, to calculate t_{LUM} accurately, those padding layers also need to be benchmarked separately. To minimize the error when calculating the latency of the LUM, we construct those padding-only models by pairing input and output padding convolution layers with matching dimensions. Therefore, each measured padding-only model consists of two convolution layers with the same number of operations and equal data input and output dimensions. Equations 2 and 3 describe $T_{b,n}$ and T_c .

$$T_{b,n} = t_{padding_in,n} + t_{padding_out,n} \quad (2)$$

$$T_c = t_{padding_in,1} + t_{LUM} + t_{padding_out,2} \quad (3)$$

Padding the LUM with convolutional layers at the input and output offers two major advantages: Firstly, it reduces

Algorithm 1 Smart Padding for Latency Benchmarking

```

Initialize look-up table for padding-only models
for each required combination of padding-only model do
    Construct the padding-only models as in Fig. 5b
    Measure total latency  $T_{b,1}$ 
    Measure total latency  $T_{b,2}$ 
    Store  $T_{b,1}$  and  $T_{b,2}$  in the look-up table
end for
for each LUM to be measured do
    Construct a padded layer model as in Fig. 5c
    Measure total latency  $T_c$ 
    Load correct  $T_{b,1}$  and  $T_{b,2}$  from look-up table
    Compute  $T_{LUM}$  using the Equation 4
end for
    
```

the amount of input data transfer to a minimum since c_{in} and c_{out} can be set to 1. Consequently, we only need to determine the latency of the padding layers including the data transfer times. Secondly, using padding layers allows us to profile layers with different input and output dimensions (e.g. convolution layers with stride) compared to other solutions such as repeating the same layer multiple times.

The algorithm for computing all T_{LUM} is summarized in Algorithm 1. Firstly, we measure the latency of the padding-only models with configuration sets characterized by the height (h), width (w), and channel dimensions (c_{in} and c_{out}). These measurements allow the creation of an exhaustive look-up table that accounts for any combination of input and output padding dimensions required for the padded layer models. Secondly, the latencies for the padded layer

models are measured. However, by using Equations 2 and 3 it is neither possible to determine t_{LUM} nor the distribution between the input and output padding layers ($t_{padding_in,1}$ and $t_{padding_out,2}$).

For layers where the dimensions of the input padding layer and the output padding layer are not identical, without per-layer profiling, the exact numbers for $t_{padding_in,1}$ and $t_{padding_out,2}$ are not obtainable. However, it is possible to compute the upper and lower bound of the latency interval of the LUM with:

$$\begin{aligned} t_{LUM_upper} &:= \mathcal{T}_c - \min_{n \in \{1,2\}} (\mathcal{T}_{b,n}) \\ t_{LUM_lower} &:= \mathcal{T}_c - \max_{n \in \{1,2\}} (\mathcal{T}_{b,n}) \\ T_{LUM} &:= [t_{LUM_lower}, t_{LUM_upper}] \end{aligned} \quad (4)$$

Compared to alternative methods, smart padding drastically reduces the interval width of T_{LUM} . This is due to the small number of additional operations and minimized data transfer of the padding layers. For example, when measuring the single layer model of a 2D convolution layer with a stride of 1, the time required for data transfer is proportional to $w \cdot h \cdot (c_{in} + c_{out})$. Here, w and h represent the width and height of the layer, while c_{in} and c_{out} refer to the number of input and output channels, respectively. The computation time scales with $w \cdot h \cdot c_{in} \cdot c_{out} \cdot k_h \cdot k_w$, where k_h and k_w are the kernel height and width. In contrast, when considering the padded layer model, the data transfer time remains the same for the input and output padding. However, the computation time for the entire model is now determined by the computation time of the LUM and an additional term that accounts for the computation time padding layers.

$$\begin{aligned} t_{comp_padded} &\propto w \cdot h \cdot c_{in} \cdot c_{out} \cdot k_h \cdot k_w \\ &+ w_{in} \cdot h_{in} \cdot c_{in} + w_{out} \cdot h_{out} \cdot c_{out} \end{aligned} \quad (5)$$

This means that the resulting width of the possible latency interval is the difference between $\mathcal{T}_{b,1}$ and $\mathcal{T}_{b,2}$. As a result, there are three major possible outcomes:

- 1) $c_{in} = c_{out}$, $w_{in} = w_{out}$, $h_{in} = h_{out}$: $t_{LUM_upper} = t_{LUM_lower}$ as a result of $\mathcal{T}_{b,1} = \mathcal{T}_{b,2}$
- 2) $c_{in} \neq c_{out}$, $w_{in} = w_{out}$, $h_{in} = h_{out}$: The error margin is dominated by the difference in computation time of the input and output padding layers
- 3) $c_{in} \neq c_{out}$, $w_{in} \neq w_{out}$, $h_{in} \neq h_{out}$: The error margin is composed of the difference in computation time and data transfer time of the input and output padding layers

As an example, Figures 6 and 7 depict the computed T_{LUM} for the case 2 ($c_{in} \neq c_{out}$, $w_{in} = w_{out}$, $h_{in} = h_{out}$). We note that the computed median and error interval of t_{LUM} are magnitudes smaller than the measured \mathcal{T}_a for the single layer model on the NXP i.MX93 development board (i.MX93) and the NVIDIA Jetson Xavier AGX (Jetson Xavier).

For the final dataset generation, T_{LUM} is computed for each individual padded model measurement alongside the calculated interval. Using this method, we can utilize the smart padding benchmarks for the layer model generation as




FIGURE 6. Computed and measured times for the three models of a 2D convolution layer with $c_{in} = 64$, $c_{out} \in [1, 256]$, $h = 64$, $w = 64$ for the i.MX93.




FIGURE 7. Computed and measured times for the three models of a 2D convolution layer with $c_{in} = 64$, $c_{out} \in [1, 256]$, $h = 64$, $w = 64$ for the Jetson Xavier.

described in [4], [5], and [8]. In general, the decision to use 2D convolution layers (including a non-linearity) as padding layers is motivated by two main factors.

Firstly, unlike when padding with slicing or concatenation operations, it ensures that there is no possibility for the compiler to further simplify the computation graph. Secondly, the same procedure can be applied to 1D and 3D convolutions while still achieving a similar reduction in operations and data transfer. Lastly, based on our understanding, the presented method of smart padding could be applied to other operators that meet those specific requirements.

Table 2 contains a list of all used features for the benchmarked layers. The width and height of the images are limited to 1024 and the kernel sizes (k_h , k_w) are limited to 11. The resulting parameter space is sub-sampled randomly and then balanced across the FLOPs dimension, to ensure a balanced dataset. The resulting number of data points is around 1000 per layer type.

C. LATENCY ESTIMATION WITH CONFIDENCE

The primary target of latency estimation frameworks is to accurately predict the application of optimization strategies layer execution time of DNNs. However, incorporating confidence metrics into latency prediction frameworks

TABLE 2. Layer feature specifications and sample sizes.

Layer	Features
Conv2D + relu	$h, w, c_{in}, c_{out}, k_w, k_h, stride_h, stride_w, FLOPs, params$
DWConv2d + relu	$h, w, c, k_w, k_h, stride_h, stride_w, FLOPs, params$
FC	$c_{in}, c_{out}, FLOPs, params$
maxpool	$h, w, c, k_w, k_h, stride_h, stride_w$
avgpool	$h, w, c, k_w, k_h, stride_h, stride_w$
add + relu	h, w, c_{in}
concat	h, w, c_{in1}, c_{in2}

significantly improves their interpretability and practical usefulness. This enhancement not only provides insights into the precision of the predictions but also informs downstream decision-making processes by flagging areas of low certainty. The implementation of confidence metrics should improve the usability of the predictors for two primary applications:

- **Hardware Platform Selection:** Since the modeling process does not work with the same accuracy for each hardware platform, providing a confidence level helps in selecting the most appropriate hardware platform.
- **Network Architecture Comparison:** When comparing DNNs with different layer types and configurations, it is important to understand which layers are outside the distribution of the training datasets and therefore not correctly predicted by the estimator.

Based on these two major use-cases, the *confidence metrics* should demonstrate several key properties at both the layer and accelerator levels. The confidence metric should:

- 1) Take into account the method of data acquisition, providing insights into the reliability of the data, especially in cases where the black-box measurement method from Section III-B is deployed.
- 2) Enable comparison of confidence in estimation at the levels of per-layer compute efficiency and per-layer latency.
- 3) Assess the coverage of the benchmark dataset and identify configurations of layers that are outside of the benchmarked design space.

To implement such confidence metrics, we rely on the conformal prediction framework which offers various options to generate statistically valid prediction regions for any underlying point predictor [25], [26]. As a result, we implement three confidence metrics that enable the comparison of DNNs prediction results and the underlying prediction models, on layer and network level:

- Confidence Metric Throughput Variance (CM_{TV})
- Confidence Metric Latency Variance (CM_{LV})
- Confidence Metric Outliers (CM_O)

For these confidence metrics, the concepts of quality estimation for conformal prediction are used to estimate different systematic uncertainties in the latency prediction models.

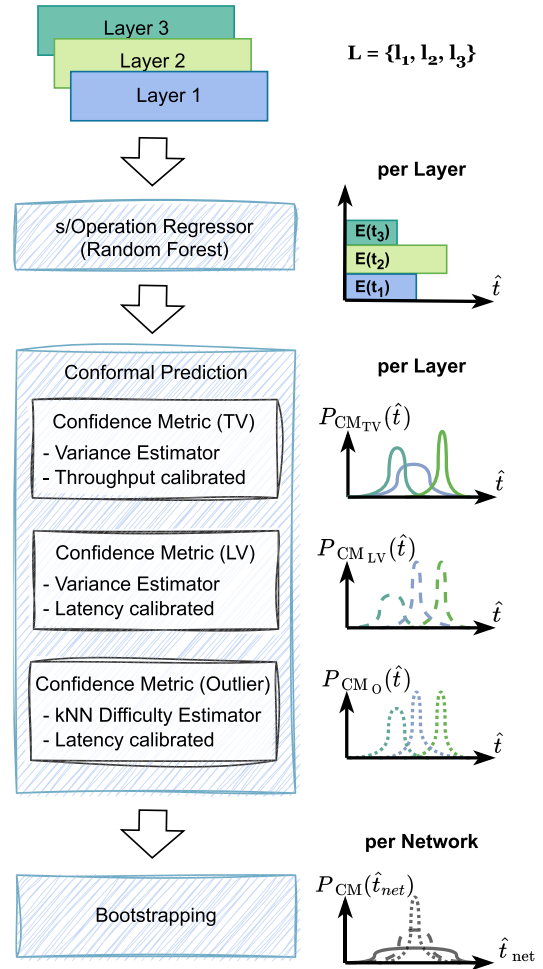


FIGURE 8. Overview of the confidence prediction methodology.

Additional confidence metrics could be easily integrated by following the same principles. In this work, the primary emphasis is on the prediction confidence of layer time predictors. Although the prediction of model optimizations performed by the optimization toolchain is also crucial in accurately estimating total network execution times, it is not the central focus of this study. The motivation behind this decision is that the correct prediction of fusion rules represents a simpler challenge than the per-layer latency prediction, due to the limited amount of possible and useful combinations of layers, in comparison to the myriad configurations of each layer type. Nevertheless, the presented concepts have the potential to be applied to the model optimization predictors in future work. Furthermore, the following methodology requires that all occurring layer types within the investigated networks are benchmarked and modeled with the statistical method of ANNETTE.

1) INFERENCE

Figure 8 depicts an overview of the confidence estimation extension for ANNETTE. The network topology is described by a set of N layers where each layer is described by a feature

vector \vec{x} which is composed of the configuration parameters describing the layer. Furthermore, \vec{x} also includes additional high-level features such as the number of parameters, number of input features, etc.:

$$L = \{l_1, l_2, l_3 \dots l_N\} \text{ where } l_i = \vec{x}_i \quad (6)$$

Based on the measured times for each layer type an individual random forest regressor is trained. For this work, the target value is the time per operation t_{op} instead of the number of operations per second. This adaption is required to avoid potential zero division for broad confidence intervals since the final computation time of a layer t_{comp} is now computed with $t_{comp} = t_{op} \cdot \text{num}_{ops}$. Based on the predictions of the regressors the expected value for the computation time of the entire network is computed as the sum of the predicted computation times of all layers:

$$\hat{t}_{net} = \sum_i^N \text{num}_{ops,i} \cdot \hat{t}_{op,i} \quad (7)$$

Consequently, from a probabilistic perspective if $P_{CM}(\hat{t}_i)$ with $i \in \{1, 2, \dots, N\}$ are the probability distributions for all layers of the network computed by the three different conformal interval predictors, the probability distribution for the total computation time for each interval predictor is computed by the convolution of the probability distributions. With $*$ as the notation of the convolution operator this results in the Equation:

$$P_{CM}(\hat{t}_{net}) = P_{CM}(\hat{t}_1) * P_{CM}(\hat{t}_2) * \dots * P_{CM}(\hat{t}_N) \quad (8)$$

To ensure that $P_{CM}(\hat{t}_{net})$ is computed correctly in all cases, we apply bootstrapping. This helps overcome limitations in the case that only a few data points are used in the calibration step for the uncertainty quantification.

2) TRAINING

For the training of the conformal regressors this work relies on the techniques implemented in CREPES [38] a Python package for generating conformal regressors and predictive systems. To ensure the robustness of our latency prediction models, the predictors are specifically trained on the median of the measured times, focusing on the quantification of the predictor's uncertainty rather than variations in latency for the same network. For each trained regressor CREPES provides a multitude of methods for the generation of confidence intervals. Firstly, to avoid splitting the training data into calibration and proper training dataset, we apply out-of-bag calibration. In contrast to standard non-normalized conformal regressors, which predict constant confidence intervals for all instances, normalized conformal regressors produce instance-specific confidence intervals based on difficulty estimates.

As mentioned in Section II there are several ways to perform the difficulty estimate. For CM_{TV} and CM_{LV} , variance-based difficulty estimation is applied. For CM_O , k-nearest neighbors (k-NN)-based difficulty estimation is

used. Additionally, while the difficulty estimation in CM_{LV} is calibrated based on the absolute prediction error of the layer latency, for CM_{TV} it is calibrated based on the absolute prediction error of the layer efficiency (s/operation). The difficulty estimation for CM_O is solely based on the feature vectors \vec{x} of the calibration data.

The effects of applying the three different normalization methods are depicted in Figure 9, which shows the 95% confidence intervals around the predicted value for the measurements performed on the Jetson Xavier from Section III-B Figure 7.

The confidence interval for CM_{TV} is depicted in Figure 9a. Since the confidence interval estimation is calibrated via the absolute error of the time per operation, the resulting confidence intervals increase with the number of operations. Hence, this CM_{TV} is more useful when comparing the prediction quality of the compute efficiency rather than the overall layer execution time.

To address this limitation, we introduce CM_{LV} . For this measure, the confidence intervals are computed based on the residuals of the computed layer execution time (see Figure 9b). It is worth noting that the confidence intervals for this measure closely align with the error interval extracted earlier, as shown in Figure 7. As a result, this CM_{LV} is most useful for comparing the confidence intervals of the overall layer execution times.

For CM_O (Fig. 9c), it can be observed that the width of the confidence intervals increase towards the boundary values of c_{out} within the example dataset. This is because the distance to the k-Nearest Neighbor data points increases for predictions in those regions. This indicates a sparse local coverage by the benchmark data, which may compromise the prediction accuracy. Thus, CM_O serves as a tool to pinpoint predictions for layers with feature vectors that are not well covered by the training dataset. For feature vectors far beyond the dataset's scope, the resulting confidence intervals might extend to negative values. However, as it is unrealistic for a layer to be computed in negative time, such wide confidence intervals should rather be viewed as indicators of underrepresented areas in the dataset than as precise latency ranges.

IV. RESULTS

For the evaluation of the methodology presented in Section III we conduct a series of experiments. First, we compare the smart padding (see Section III-B) benchmarking method with padded models to simple single-layer benchmarking in terms of overall prediction quality. Secondly, to evaluate the confidence prediction method, we perform a series of experiments to determine if the desired properties listed in Section III-C are met. The experiments include the results for three different hardware platforms: the NVIDIA Jetson Xavier AGX, NXP i.MX 93, and NXP i.MX8M+ development board (i.MX8M+). The Jetson Xavier was operated at maximum power setting with TensorRT as the inference runtime, using integer 8-bit precision and offering 22 TOPs, not

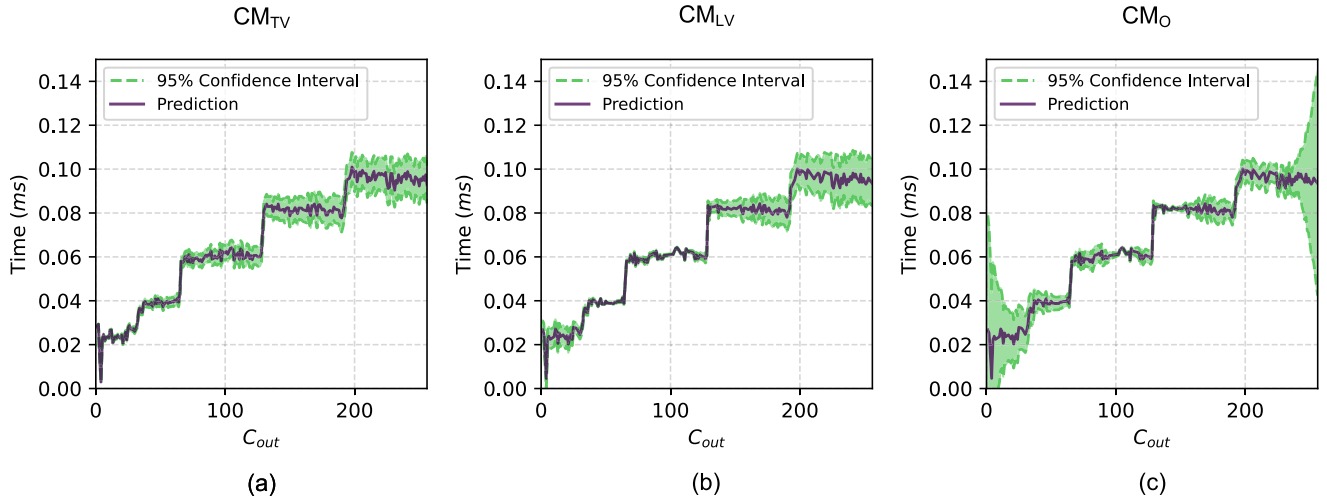


FIGURE 9. Overview of predicted Confidence Intervals for 2D convolution layer with $c_{in} = 64$, $c_{out} \in [1, 256]$, $h = 64$, $w = 64$ on the Jetson Xavier: (a) Confidence s/Operations Normalized, (b) Confidence Normalized with respect to time, (c) Confidence with k-NN difficulty estimation.

TABLE 3. Percentage prediction errors for the ANNETTE models for the i.MX93 in comparison to the Vela compiler estimates.

Network	Measured [ms]	Vela Compiler	Single Layer	Smart Padding
YOLOv5s	103.6	+30.5%	+292.5%	-9.5%
YOLOv5m	213.0	+21.1%	+260.8%	-3.6%
YOLOv5l	383.1	+22.2%	+233.2%	-0.8%
YOLOv8n	67.6	+104.1%	+210.6%	-3.9%
YOLOv8s	138.9	+83.0%	+192.4%	-5.1%
YOLOv8m	294.8	+64.6%	+172.5%	-1.3%
YOLOv8l	486.0	+77.4%	+200.6%	+5.8%
YOLOv8x	763.8	+56.1%	+173.6%	+0.5%
MobilenetV1	4.97	+78.8%	+759.6%	+24.6%
InceptionV4	66.4	+40.5%	+492.4%	-1.5%
Avg. error		57.8%	298.8%	5.7%

TABLE 4. Percentage prediction errors for the ANNETTE models for the Jetson Xavier in comparison to the ANNETTE model based on the per-layer profiling.

Network	Measured [ms]	Per-Layer Profiling	Single Layer	Smart Padding
YOLOv5s	4.6	+0.1%	+55.3%	-2.1%
YOLOv5m	9.4	+0.0%	+65.9%	-4.8%
YOLOv5l	14.2	+12.1%	+79.1%	-7.4%
YOLOv8n	5.5	-21.9%	+17.9%	-18.3%
YOLOv8s	7.29	-2.5%	+53.1%	-3.1%
YOLOv8m	13.5	-3.6%	+64.3%	-6.5%
YOLOv8l	19	+9.9%	+80.5%	+4.2%
YOLOv8x	28.9	-1.6%	+53.9%	-9.7%
MobilenetV1	0.45	-17.5%	+59.3%	-6.1%
InceptionV4	4.82	-14.4%	+116.2%	+6.0%
Avg. abs. error		8.4%	64.6%	6.9%

considering the Deep Learning Accelerator (DLA) cores. The i.MX93 utilized TensorFlow Lite with 8-bit quantization and the TensorFlow Lite inference runtime delegate, providing up to 1 TOPS using the ARM Ethos U65 microNPU. The i.MX8M+ employed the VeriSilicon VIP9000 NPU, delivering up to 2.3 TOPS also using the TensorFlow Lite inference runtime.

A. BLACK-BOX BENCHMARKING

The goal of the following experiments is to compare the quality of the collected smart padding benchmark data with the single-layer benchmark data and assess how well they serve as ground-truth data for prediction models. For the presented results, we generate ANNETTE prediction models using both the single-layer and smart padding methods. These generated prediction models are then compared in terms of total network latency against the measured network latencies. Additionally, we compare the results to the predictions provided by the ARM Vela compiler³ for the Ethos U65 NPU on the i.MX93.

³<https://pypi.org/project/ethos-u-vela>

TABLE 5. Percentage prediction errors for the ANNETTE models for the i.MX8M+.

Network	Measured [ms]	Single Layer	Smart Padding
YOLOv5s	87.3	+96.7%	-6.6%
YOLOv5m	165.3	+88.4%	-5.7%
YOLOv5l	260.1	+84.8%	-3.2%
YOLOv8n	54.4	+56.2%	-2.5%
YOLOv8s	100.9	+43.7%	-2.5%
YOLOv8m	186.4	+37.1%	-6.2%
YOLOv8l	286.1	+36.3%	-4.7%
YOLOv8x	363.0	+36.6%	-7.3%
MobilenetV1	3.69	+305.2%	+2.4%
InceptionV4	63.2	+110.8%	-51.7%
Avg. abs. error		89.6%	9.3%

Table 3 shows the prediction accuracy for a set of state-of-the-art DNNs for the i.MX93. In the case of the i.MX93, the smart padding-based ANNETTE prediction demonstrates superior performance compared to the single-layer ANNETTE prediction and the Vela estimates, achieving higher prediction accuracy across all networks. The average prediction errors for the smart padding-based ANNETTE

prediction, single layer-based ANNETTE prediction, and Vela estimates are 5.7%, 298.8%, and 57.8% respectively. Further in-depth analysis revealed that benchmarking individual layers on the i.MX93 results in additional time overhead due to an extra quantization step. This leads to a more substantial improvement than expected, thanks to the smart padding method. Likewise, for the smart padding-based and single layer-based ANNETTE prediction, the average percentage errors are 6.9% and 64.6% for the Jetson Xavier, and 9.2% and 89.6% for the i.MX8M+. The detailed results for the Jetson Xavier and i.MX8M+ are shown in Tables 4 and 5, respectively.

Notably, only in 3 instances, the smart padding-based ANNETTE estimation errors are larger than 10%. These errors can be explained by the limited dataset used for this work, which does not cover a stride different than 1 and asymmetric convolution kernels. These limitations result in not optimal prediction results for InceptionV4, MobilenetV1, and YOLOv8n but also allow us to evaluate the confidence metrics.

To compare the accuracy of the smart padding strategy with per-layer profiling, we utilized the built-in per-layer profiling from the Nvidia benchmarking script for the Jetson Xavier. The Pearson correlation coefficient between the Nvidia per-layer profiling and the computed mean value of T_{LUM} was found to be 0.975. The ANNETTE predictions based on per-layer profiling data for the Jetson Xavier showed similar accuracy to those from the smart padding-based ANNETTE models. Therefore, we conclude that smart padding benchmarking provides profiling accuracy comparable to per-layer benchmarking, while significantly reducing implementation efforts. The overhead associated with the smart padding benchmarking strategy is limited to measuring 854 padding-only models.

B. CONFIDENCE METRICS

For the evaluation of the *confidence metrics*, we display the results on model, network, and layer levels. Firstly, since CM_{TV} is throughput calibrated, it mostly serves to compare the normalized per-layer confidence interval size for different models. This can, for example, be used to compare the overall confidence of the previously computed models.

1) MODEL-LEVEL COMPARISON

Figure 10 displays the average normalized 90% confidence interval size for the generated models for all tested networks. To evaluate the influence of the smart padding method on the generated latency prediction models, we also generate models based on the mean value without including the previously computed intervals (see Section III-B).

As outlined in Section III-B, the CM_{TV} , which is used for this comparison, is calibrated with regard to the layer throughput. As a result, this metric provides a measure for comparing the confidence for the compute efficiency predictions across the tested hardware platforms and all tested layers. As expected, including the smart padding




FIGURE 10. Average normalized 90% confidence intervals for CM_{TV} on the all tested networks.

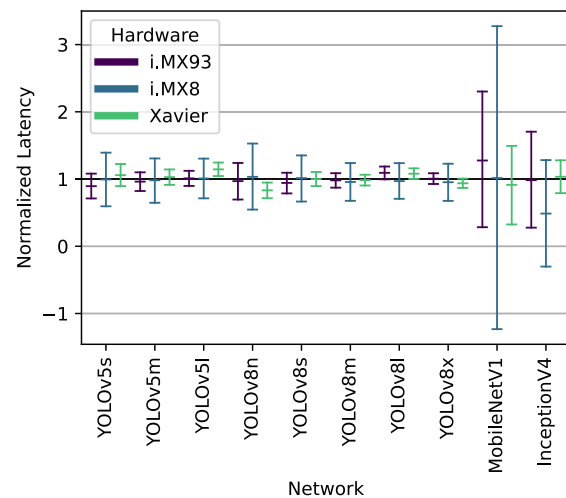


FIGURE 11. CM_{LV} for the tested networks.

intervals in the calibration of the confidence metrics leads to larger confidence intervals. Notably, the increase of the confidence interval widths differs for the different hardware platforms. We conclude that CM_{TV} can be used to determine which hardware platforms would profit the most from implementing per-layer profiling and for which hardware platforms, the smart padding method is sufficient. Furthermore, CM_{TV} can guide engineers in situations where a platform may show consistently high throughput variance across different network layers. By using CM_{TV} , engineers can prioritize hardware that demonstrates lower throughput variance, suggesting more stable performance across diverse workloads.

2) NETWORK-LEVEL COMPARISON

For the network-level comparison, the 90% confidence intervals of CM_{LV} and CM_O are displayed for all networks in Figure 11 and 12 respectively. As mentioned in Section IV-A, these confidence metrics provide a deeper understanding of the predictions performed for each individual network. It is noticeable that the confidence intervals for MobilenetV1 and InceptionV4 are particularly large, which aligns with the occurrence of inaccurate predictions in certain cases. A large confidence interval for CM_{LV} indicates sub-optimal




FIGURE 12. CM_O for the tested networks.

prediction accuracy due to high variances in the dataset within the prediction region. To address this, engineers can use CM_{LV} as an indicator to refine the training datasets by incorporating more diverse network configurations that mirror the operational settings. This can help in reducing the confidence intervals and thus improving the accuracy of latency predictions. Conversely, a large confidence interval for CM_O suggests inadequate coverage of one or more layers in the collected dataset, potentially leading to inaccurate prediction results. In this case, we can go one step further and analyze the prediction results on a layer level.

3) LAYER-LEVEL COMPARISON

CM_{LV} and CM_O provide insights into the root causes of potentially inaccurate prediction on a layer level. Figure 13 displays the confidence interval widths for the latency predictions of YOLOv8n on the Jetson Xavier. The layers with large confidence interval widths almost exclusively have a stride of 2 which is not covered well in the example benchmark dataset. This can be detected by the large confidence interval widths of the CM_O for those specific layers.

On the other hand, the CM_{LV} interval widths for hint layers 1 and 4 at prediction regions with high variance of measured latencies. However, compared with the per-network CM_{LV} and CM_O , we see that the prediction confidence for InceptionV4 for the Jetson Xavier could be improved by extending the dataset so that the required layer configurations are well covered.

For example, in a hardware-aware NAS process, the CM_O can prevent the NAS from settling on a seemingly optimal architecture that performs poorly in untested conditions. Meanwhile, CM_{TV} and CM_{LV} ensure the chosen architecture consistently meets performance expectations across a range of architecture variations, thereby avoiding costly misestimations of network efficiency




FIGURE 13. YOLOv8n per-layer confidence interval widths of CM_{LV} and CM_O for predictions for the Jetson Xavier.

V. CONCLUSION

This study introduces a novel approach for benchmarking DNN accelerators that eliminates the need for per-layer profiling for existing latency estimation frameworks. As a result, the setup for benchmarking new hardware is simplified, and the potential profiling overhead can be eliminated. The experiments underscore the method’s effectiveness across three distinct hardware platforms (Jetson Xavier, i.MX8M+ and i.MX93), improving the latency prediction accuracy by a large margin in comparison to single-layer benchmarking and outperforming the latency prediction of the ARM Vela compiler. Furthermore, this study integrates three confidence metrics to improve the usability and interpretability of latency prediction frameworks.

From the perspective of developers, the introduction of smart padding not only decreases the implementation effort when benchmarking new hardware platforms but also allows benchmarking without profiling overhead. Furthermore, the adoption of our confidence framework has already yielded significant insights into the prediction models for certain hardware platforms. With the guidance of the confidence metrics, we were able to precisely identify and correct inaccuracies in layer-specific predictions.

For end-users, the introduced confidence metrics offer a more informed basis for selecting hardware and network models for DNN deployment.

Future research could further refine the smart padding method by exploring its application across diverse network layers and other domains like time series and 3D modeling. Expanding the use of conformal prediction methods to include GNN based latency prediction methods and developing automated benchmark point selection based on confidence levels are also promising directions.

REFERENCES

- [1] H. Cai, L. Zhu, and S. Han, "ProxylessNAS: Direct neural architecture search on target task and hardware," in *Proc. ICLR*, 2019.
- [2] M. Tan, B. Chen, R. Pang, V. Vasudevan, M. Sandler, A. Howard, and Q. V. Le, "MnasNet: Platform-aware neural architecture search for mobile," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 2815–2823.
- [3] A. Anderson, J. Su, R. Dahyot, and D. Gregg, "Performance-oriented neural architecture search," in *Proc. Int. Conf. High Perform. Comput. Simul. (HPCS)*, Jul. 2019, pp. 177–184.
- [4] M. Lechner and A. Jantsch, "Blackthorn: Latency estimation framework for CNNs on embedded NVIDIA platforms," *IEEE Access*, vol. 9, pp. 110074–110084, 2021.
- [5] M. Wess, M. Ivanov, C. Unger, A. Nookala, A. Wendt, and A. Jantsch, "ANNETTE: Accurate neural network execution time estimation with stacked models," *IEEE Access*, vol. 9, pp. 3545–3556, 2021.
- [6] Z. Wang, P. Yang, L. Hu, B. Zhang, C. Lin, W. Lv, and Q. Wang, "SLAPP: Subgraph-level attention-based performance prediction for deep learning models," *Neural Netw.*, vol. 170, pp. 285–297, Feb. 2024.
- [7] Y. Chai, D. Tripathy, C. Zhou, D. Gope, I. Fedorov, R. Matas, D. Brooks, G.-Y. Wei, and P. Whatmough, "PerfSAGE: Generalized inference performance predictor for arbitrary deep learning models on edge devices," 2023, *arXiv:2301.10999*.
- [8] L. L. Zhang, S. Han, J. Wei, N. Zheng, T. Cao, Y. Yang, and Y. Liu, "Nn-meter: Towards accurate latency prediction of deep-learning model inference on diverse edge devices," in *Proc. 19th Annu. Int. Conf. Mobile Syst., Appl., Services*, Jun. 2021, pp. 81–93.
- [9] M. Sabih, A. Mishra, F. Hannig, and J. Teich, "MOSP: Multi-objective sensitivity pruning of deep neural networks," in *Proc. IEEE 13th Int. Green Sustain. Comput. Conf. (IGSC)*, Oct. 2022, pp. 1–8.
- [10] Z. Chen, C. Liu, W. Yang, K. Li, and K. Li, "LAP: Latency-aware automated pruning with dynamic-based filter selection," *Neural Netw.*, vol. 152, pp. 407–418, Aug. 2022.
- [11] Y. Xu, T. Mohammed, M. Di Francesco, and C. Fischione, "Distributed assignment with load balancing for DNN inference at the edge," *IEEE Internet Things J.*, vol. 10, no. 2, pp. 1053–1065, Jan. 2023.
- [12] Y. Duan and J. Wu, "Optimizing job offloading schedule for collaborative DNN inference," *IEEE Trans. Mobile Comput.*, vol. 23, no. 4, pp. 3436–3451, Apr. 2023.
- [13] Y. Gao, X. Gu, H. Zhang, H. Lin, and M. Yang, "Runtime performance prediction for deep learning models with graph neural network," in *Proc. IEEE/ACM 45th Int. Conf. Softw. Eng., Softw. Eng. Pract. (ICSE-SEIP)*, May 2023, pp. 368–380.
- [14] L. Dudziak, T. Chau, M. S. Abdelfattah, R. Lee, H. Kim, and N. D. Lane, "BRP-NAS: Prediction-based NAS using GCNs," 2020, *arXiv:2007.08668*.
- [15] H. Cai, C. Gan, T. Wang, Z. Zhang, and S. Han, "Once-for-all: Train one network and specialize it for efficient deployment," in *Proc. ICLR*, 2020.
- [16] H. Lee, S. Lee, S. Chong, and S. J. Hwang, "HELP: Hardware-adaptive efficient latency prediction for NAS via meta-learning," 2021, *arXiv:2106.08630*.
- [17] K. Lübeck, A. L. Jung, F. Wedlich, and O. Bringmann, "Work-in-progress: Ultra-fast yet accurate performance prediction for deep neural network accelerators," in *Proc. Int. Conf. Compil., Archit., Synth. Embedded Syst. (CASES)*, Oct. 2022, pp. 27–28.
- [18] L. Mei, H. Liu, T. Wu, H. E. Sumbul, M. Verhelst, and E. Beigne, "A uniform latency model for DNN accelerators with diverse architectures and dataflows," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, Mar. 2022, pp. 220–225.
- [19] J. Li, R. Ma, V. S. Malthody, C. Samplawski, B. Marlin, S. Chen, S. Yao, and T. Abdelzaher, "Towards an accurate latency model for convolutional neural network layers on GPUs," in *Proc. IEEE Mil. Commun. Conf. (MILCOM)*, Nov. 2021, pp. 904–909.
- [20] S. Nair, S. Abbasi, A. Wong, and M. J. Shafiee, "MAPLE-edge: A runtime latency predictor for edge devices," 2022, *arXiv:2204.12950*.
- [21] S. Abbasi, A. Wong, and M. J. Shafiee, "MAPLE-X: Latency prediction with explicit microprocessor prior knowledge," 2022, *arXiv:2205.12660*.
- [22] Y. Li, J. Ma, D. Cao, and H. Mei, "Sectum: Accurate latency prediction for TEE-hosted deep learning inference," in *Proc. IEEE 42nd Int. Conf. Distrib. Comput. Syst. (ICDCS)*, Jul. 2022, pp. 906–916.
- [23] K. G. Mills, F. X. Han, J. Zhang, F. Chudak, A. S. Mamaghani, M. Salameh, W. Lu, S. Jui, and D. Niu, "GENNAPE: Towards generalized neural architecture performance estimators," in *Proc. AAAI*, 1609, pp. 9190–9199.
- [24] K. P. Selvam and M. Brorsson, "DIPPM: A deep learning inference performance predictive model using graph neural networks," in *Proc. Euro-Par. Cham, Switzerland: Springer*, 2023, pp. 3–16.
- [25] A. Gammerman, V. Vovk, and V. Vapnik, "Learning by transduction," in *Proc. UAI*, G. F. Cooper and S. Moral, Eds. San Mateo, CA, USA: Morgan Kaufmann, 1998, pp. 148–155.
- [26] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic Learning in a Random World*, vol. 29. New York, NY, USA: Springer, 2005.
- [27] N. Meinshausen, "Quantile regression forests," *J. Mach. Learn. Res.*, vol. 7, pp. 983–999, Nov. 2006.
- [28] T. Gneiting, "Quantiles as optimal point forecasts," *Int. J. Forecasting*, vol. 27, no. 2, pp. 197–207, Apr. 2011.
- [29] L. Hespanhol, C. S. Vallio, L. M. Costa, and B. T. Saragiotto, "Understanding and interpreting confidence and credible intervals around effect estimates," *Brazilian J. Phys. Therapy*, vol. 23, no. 4, pp. 290–301, Jul. 2019.
- [30] L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001.
- [31] U. Johansson, H. Boström, T. Löfström, and H. Linusson, "Regression conformal prediction with random forests," *Mach. Learn.*, vol. 97, nos. 1–2, pp. 155–176, Oct. 2014.
- [32] H. Boström, H. Linusson, T. Löfström, and U. Johansson, "Accelerating difficulty estimation for conformal regression forests," *Ann. Math. Artif. Intell.*, vol. 81, nos. 1–2, pp. 125–144, Mar. 2017.
- [33] C. Lu, A. Lemay, K. Chang, K. Höbel, and J. Kalpathy-Cramer, "Fair conformal predictors for applications in medical imaging," in *Proc. AAAI*, 2022, pp. 12008–12016.
- [34] C. Eliades and H. Papadopoulos, "Conformal prediction for automatic face recognition," in *Proc. COPA*, in Proceedings of Machine Learning Research, vol. 60, A. Gammerman, V. Vovk, Z. Luo, and H. Papadopoulos, Eds., 2017, pp. 62–81.
- [35] W. Wisniewski, D. Lindsay, and S. Lindsay, "Application of conformal prediction interval estimations to market makers' net positions," in *Proc. COPA*, in Proceedings of Machine Learning Research, vol. 128, A. Gammerman, V. Vovk, Z. Luo, E. N. Smirnov, G. Cherubin, and M. Christini, Eds., 2020, pp. 285–301.
- [36] ARM. *ARM Ethos NPU Technical Reference Manual*. Accessed: Mar. 1, 2024. [Online]. Available: <https://developer.arm.com/documentation/102420/0200/Functional-description/Functional-blocks>
- [37] S. H. Zargarbashi, S. Antonelli, and A. Bojchevski, "Conformal prediction sets for graph neural networks," in *Proc. IJML*, in Proceedings of Machine Learning Research, vol. 202, A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, Eds., 2023, pp. 12292–12318.
- [38] H. Boström, "crepes: A Python package for generating conformal regressors and predictive systems," in *Proc. COPA*, in Proceedings of Machine Learning Research, vol. 179, U. Johansson, H. Boström, K. A. Nguyen, Z. Luo, and L. Carlsson, Eds., 2022, pp. 24–41.



MATTHIAS WESS received the B.Sc. and M.Sc. degrees from the Department of Electrical Engineering, TU Wien, Vienna, Austria, in 2013 and 2017, respectively, where he is currently pursuing the Ph.D. degree with the Institute for Computer Technology. As a member of the Christian Doppler Laboratory for Embedded Machine Learning, his research is primarily focused on the latency estimation of deep neural networks and enhancing the energy efficiency of machine learning algorithms.



networks for embedded inference.

DANIEL SCHNÖLL (Graduate Student Member, IEEE) received the M.Sc. degree in embedded systems from TU Wien, Vienna, Austria, in 2021. He is currently pursuing the Ph.D. degree with the Institute for Computer Technology, Christian Doppler Laboratory for Embedded Machine Learning, TU Wien. He is a part of the Christian Doppler Laboratory for Embedded Machine Learning, TU Wien. His current research interests include TinyML and optimization of deep neural



research interests include energy-efficient machine learning for time-series applications and leveraging artificial intelligence for sustainability.

MATTHIAS BITTNER received the M.Sc. degree in automation and control from TU Wien, Vienna, Austria, in 2021, and the M.Sc. degree in artificial intelligence from Johannes Kepler University, Linz, Austria, in 2024. He is currently pursuing the Ph.D. degree with the Institute for Computer Technology, Christian Doppler Laboratory for Embedded Machine Learning, TU Wien. He is affiliated with the Christian Doppler Laboratory for Embedded Machine Learning, TU Wien. His



DOMINIK DALLINGER (Graduate Student Member, IEEE) received the Bachelor of Science degree in electrical engineering from TU Wien, Vienna, Austria, in 2021, where he is currently pursuing the master's degree in embedded systems, with a broad focus on mechatronics, machine vision, computer systems, and electronics design. He is also engaged with the Christian Doppler Laboratory for Embedded Machine Learning, focusing his research on TinyML.



published five books as an editor and one as an author and over 300 peer-reviewed contributions in journals, books, and conference proceedings. His current research interests include systems on chips, self-aware cyber-physical systems, and embedded machine learning. He has given over 100 invited presentations at conferences, universities, and companies.

AXEL JANTSCH (Fellow, IEEE) received the Dipl.-Ing. and Ph.D. degrees in computer science from TU Wien, Vienna, Austria, in 1987 and 1992, respectively. From 1997 to 2002, he was an Associate Professor with the KTH Royal Institute of Technology, Stockholm. From 2002 to 2014, he was a Full Professor of electronic systems design with KTH. Since 2014, he has been a Professor of systems on chips with the Institute of Computer Technology, TU Wien. He has

...