

# NACHOS: Neural Architecture Search for Hardware Constrained Early Exit Neural Networks

Matteo Gambella, Jary Pomponi, Simone Scardapane, and Manuel Roveri

Dipartimento di Elettronica, Informazione e Bioingegneria,

Politecnico di Milano, Milan, Italy

Sapienza Università di Roma, Rome, Italy

Email: {matteo.gambella,manuel.roveri}@polimi.it

Email: {jary.pomponi,simone.scardapane}@uniroma1.it

**Abstract**—Early Exit Neural Networks (EENNs) endow a standard Deep Neural Network (DNN) with Early Exit Classifiers (EECs), to provide predictions at intermediate points of the processing when enough confidence in classification is achieved. This leads to many benefits in terms of effectiveness and efficiency. Currently, the design of EENNs is carried out manually by experts, a complex and time-consuming task that requires accounting for many aspects, including the correct placement, the thresholding, and the computational overhead of the EECs. For this reason, the research is exploring the use of Neural Architecture Search (NAS) to automatize the design of EENNs. Currently, few comprehensive NAS solutions for EENNs have been proposed in the literature, and a fully automated, joint design strategy taking into consideration both the backbone and the EECs remains an open problem. To this end, this work presents Neural Architecture Search for Hardware Constrained Early Exit Neural Networks (NACHOS), the first NAS framework for the design of optimal EENNs satisfying constraints on the accuracy and the number of Multiply and Accumulate (MAC) operations performed by the EENNs at inference time. In particular, this provides the joint design of backbone and EECs to select a set of admissible (i.e., respecting the constraints) Pareto Optimal Solutions in terms of best trade-off between the accuracy and number of MACs. The results show that the models designed by NACHOS are competitive with the state-of-the-art EENNs. Additionally, this work investigates the effectiveness of two novel regularization terms designed for the optimization of the auxiliary classifiers of the EENN.

**Index Terms**—Neural Architecture Search (NAS), Once-For-All Network (OFA), Constrained optimization, Early Exit Neural Networks, Conditional Computation, Adaptive Neural Networks.

## I. INTRODUCTION

Early Exit Neural Networks (EENNs) are an emerging class of neural networks characterized by the presence of multiple early exits, called Early Exit Classifiers (EECs), placed at different points of the neural network architecture (backbone) to provide intermediate predictions when enough confidence in the classification is achieved [1], [2]. This leads to significant benefits in terms of effectiveness and efficiency, such as the mitigation of overfitting [3], overthinking [4], [5], and vanishing gradient phenomena [6].

The motivating idea underlying EENNs is that the majority of input samples can be classified by resorting to smaller architectures [7], even in very complex datasets such as ImageNet [8]. Indeed, previous works have shown that the

majority of input samples can be classified by the EENNs in their earlier stages, resulting in a significant reduction of the average inference time [1] [9]. This is particularly beneficial for embedded applications implemented by Internet-of-Things (IoT) units, embedded systems, and Edge computing units from the computation and energetic point of view.

For all these reasons, the research interest in designing effective and efficient EENNs is steadily growing with many solutions available in the literature enhancing state-of-the-art neural networks (e.g., Convolutional Neural Networks and Transformers) and addressing different tasks ranging from image classification and semantic segmentation to Natural Language Processing [10]–[14]. Currently, the design of EENNs is carried out manually by experts. Unfortunately, this is a very complex and time-consuming task since many aspects must be accounted for, such as how many EECs to consider, their placement, the setting of per-exit confidence thresholds [15], and the calibration of the EECs [16].

Neural Architecture Search (NAS) is a prominent research area [17]–[19] aiming at automating the design of a neural network by exploring different architectural configurations given a neural network topology, a task to be accomplished, and a dataset used for training and validation so as to provide the optimal architecture. The majority of works in this field focus on the design of deep neural network architectures while taking into account a specific target hardware. This family of NAS solutions is called Hardware-Aware NAS [20]. An emerging direction of these NAS solutions regards the use of constraints in the NAS exploration [21]–[24]. Specifically, the works belonging to this research direction are able to consider both functional constraints (i.e., the type of operations that can be carried out in the neural network) and technological constraints (i.e., constraints on the computational and memory demand of the designed neural network in the exploration of the neural networks). The constraints are usually handled by a genetic algorithm applying a penalization to the fitness of the architectures not satisfying the hardware constraints. Interestingly, there are very few examples of NAS for EENNs in the current literature and these works are not able to provide an effective joint design of the backbone networks and the EECs [23], [25]–[27]. These NAS solutions are commented in Section II.

In this paper, we propose NACHOS, a shorthand for Neural Architecture Search for Hardware Constrained Early Exit Neural Networks, a novel NAS framework able to co-design the network architecture and the EECs and impose constraints on the classification accuracy and the number of Multiply and Accumulate (MAC) operations performed by the EENN at inference time. This solution, which is inspired by what is proposed in [27], automatically designs and places the EECs in the candidate backbone of an EENN network according to a criterion regarding the number of MACs and tunes their thresholds accordingly. Moreover, NACHOS introduces two regularization terms to enforce confidences around the operating point of the EECs to enhance the accuracy and enforce the respect of the constraint during the training of a candidate network. Lastly, a novel figure of merit named  $F_{CM}$  accounts for the number of MACs of an EENN and the related constraint adaptively during the NAS search. Overall, NACHOS shows a significant improvement in terms of the effectiveness and efficiency of the designed EENNs with respect to baseline state-of-the-art approaches. We highlight that NACHOS falls into the family of Hardware-Aware NAS.

To sum up, the novel contributions introduced in this study are the following:

- a novel NAS solution able to co-design the backbone and the EECs in EENNs satisfying user-defined constraints by automatically placing the EECs in the backbone of an EENN accounting for their computational complexity;
- the introduction of a regularization term to ensure that the local certainties of the EECs follow their operating point in order to enhance the accuracy;
- the use of a regularization approach to increase the computational efficiency of the early-exit strategy and account for the constraint;
- a novel figure of merit optimized by the genetic algorithm of the NAS to adaptively handle the constraint.

The rest of this paper is organized as follows. Section II introduces the related works regarding EENNs designed by NAS solutions. Section III provides the background. Section IV introduces the NACHOS framework developed in this study, while in Section V the experimental results aiming at evaluating the effectiveness of NACHOS are shown. Conclusions are finally drawn in Section VI. To facilitate comparisons and reproducibility, the source code of NACHOS is released to the scientific community as a public repository.<sup>1</sup>

## II. RELATED LITERATURE

In the literature, many NAS solutions address explicitly the problem of energy efficiency of deep neural networks by optimizing both the accuracy in a given task and the hardware requirements (e.g., the memory and the computational demand). These works belong to a family of solutions called Hardware-Aware NAS [20]. In this field, a promising research direction regards the design of neural networks in constrained environments. These NAS solutions permit to

impose constraints on the neural network exploration to focus the search in the region of feasible solutions accommodating for the hardware requirements [21]–[24]. Among the Hardware-Aware NAS solutions, some recent works propose NAS solutions for optimizing the design of neural networks for EENNs, which are briefly reviewed here.

EExNAS [25] is a first attempt to optimize an EENN with NAS. It accounts for only one EEC at a fixed position. Differently, ESAI [26] is a NAS for EENNs in a Cloud setting. The EENN is organized into two sub-parts, i.e., a lighter one is processed locally and another heavier one on cloud. The early exit mechanism is used to decide where to send the samples to units (e.g., local or remote server) of a technological ecosystem. Both works operate on medical applications, the first focusing on Myocardial Infarction and Human Activity Recognition, and the second on the classification of dermoscopic images. HADAS [23] presents a framework to design the backbone network, the EECs, and the dynamic voltage and frequency scaling settings, by a two-level optimization with two different genetic algorithms. On the first level, HADAS selects the optimal backbone neural networks by sampling from a Once-For-All (OFA) supernet [28] and taking into consideration the hardware constraints, while on the second level, it enhances the chosen backbones with the EECs by automating the choice of their number and placement, trains the obtained EENNs, and then selects the optimal EENNs. The EENNs are trained by using the Knowledge Distillation (KD) approach. This NAS takes 2-3 GPU days on a cluster of 32 GPUs on CIFAR-100.

A different approach suggests that the OFA can be enhanced and trained with early exits [29]. The work presents OFAv2, a novel OFA extending the OFAMobileNetv3 architecture by adding parallel blocks, additional dense skip connections, and implementing early exits. To support the redesign of the OFA, the training algorithm of the OFA, named Progressive Shrinking, has been expanded by adding two additional steps to handle the parallel blocks and the early exits, respectively. The NAS approach used in the experimental campaign relying on the novel OFA is evaluated on the Tiny Imagenet Dataset, a smaller version of Imagenet, which provides a significant increase in accuracy with respect to the original OFA architecture and maintains the computational efficiency. All in all, OFAV2 performs the search on a previously trained supernet specifically designed with multi-exits.

Currently, there are only few NAS solutions that jointly design the backbones and the EECs while taking into account the computational demand. EDANAS [27] is a framework for the automatic design of both the backbone architecture and the parameters managing the early exit mechanism of the EENNs so as to optimize both the accuracy of the classification tasks and the computational demand. The problem addressed by EDANAS is formulated as a joint optimization problem consisting in selecting the architecture of an EENN aiming at providing a high accuracy on a given task and reducing the MACs associated with the EE by using the genetic algorithm NSGA-II [30]. Similarly, HarvNet [31] addresses

<sup>1</sup><https://github.com/AI-Tech-Research-Lab/CNAS>

the joint design of backbone and EECs while meeting memory and energy constraints in dynamic harvesting environments. It relies on reinforcement learning to adjust the confidence threshold of each exit to maximize accuracy by considering the device’s runtime factors. This NAS takes 47 GPU days on CIFAR-10, since they integrate the runtime aspects of dynamic input mappings and dynamic harvesting environments in the design of NNs respectively.

Lastly, we emphasize that our work shows a hardware-aware NAS for EENNs on computer vision applications working on top of a standard single-exit OFA supernet independent of input mappings and the environment with a search time of 2 GPU days on a single GPU (NVIDIA A40). This overcomes some of the limitations of the previous works such as limited application scenarios, specific design of the supernet, and high computational requirements.

### III. BACKGROUND

#### A. Early Exit Neural Networks

We can see a neural network as a composition of  $b$  sequential blocks:

$$f(x) = B_b \circ B_{b-1} \circ \dots \circ B_1(x) \quad (1)$$

where  $\circ$  denotes function composition. This is a generic definition which is agnostic about the internal topology of a block, and that applies to most neural networks. Moreover, we can apply a classification layer to each block  $i$ -th of the model, which leads to multiple prediction functions:

$$f_i(x) = C_i \circ (B_i \circ B_{i-1} \circ \dots \circ B_1(x)) \quad (2)$$

This formulation leads to an early exit model which is capable of classifying input samples at any stage of the computation. The result is a sequence of predictions for a given input sample, and different exits can mistake different subsets of the dataset. Thus, combining them can improve the overall performance, or reduce the computational time required to calculate the prediction associated with a sample, under a given measure of certainty, which is used to halt the computation in a given branch, if the prediction is good enough.

A common way to train such models is by jointly updating all the exits:

$$L_{\text{joint}}(x, y) = l(f_b(x), y) + \sum_{i=1}^{b-1} \lambda_i \cdot \underbrace{[l(f_i(x), y)]}_{i\text{-th branch loss}}, \quad (3)$$

where  $l(\cdot, \cdot)$  is the loss (e.g. cross-entropy loss), and  $\lambda_i$  the weight associated to the  $i$ -th exit’s loss. Other training approaches are *Layer-wise Training* (LT), aiming at separately optimizing each layer in the NN architecture [32], and *Knowledge Distillation* (KD) where the backbone network is trained separately from the EECs and then the EECs are trained on top of it by relying on KD [33].

After training the model using the aforementioned approach, we can take advantage of the formulation proposed in Eq. 2

to select the earliest feasible exit which is assumed to output the correct prediction. This leads to many benefits such as the reduction of the inference time [1] [9], overfitting [3] and overthinking [1]. To do this we can use different early exit mechanisms, such as entropy, proposed in [1], the largest value of the softmax, or the so-called Score Margin (SM) [15] [34], which is the difference between the largest and the second largest value on the softmax on the prediction. These mechanisms are simple to implement, but they introduce a potential mismatch between training and test behaviours, since the early exit mechanism is not optimized on the training data.

In this paper we follow the alternative formulation proposed in [35], which, using a hierarchical differentiable mechanism, includes the early exit mechanism into the exit itself. To do that, along with the EEC  $C_i$ , we define also a confidence score  $c_i(\cdot)$  which outputs a scalar in  $[0, 1]$ , which follows the same composition rule proposed in Equation (2). So we have that, at a given branch  $i$ , we have:

$$f_i^c(x) = \begin{bmatrix} f_i(x) \\ c_i(x) \end{bmatrix}. \quad (4)$$

The idea is that the confidence score should output a value which measures how much the exit is confident of its prediction. We recursively combine these confidences to produce a cumulative confidence score defined as:

$$c_i^r(x) = c_i(x) \prod_{k=1}^{i-1} (1 - c_k(x)), \quad (5)$$

with  $c_b(\cdot) = 1$  to force the computation to halt when reaching the final exit. To produce the final output, we integrate the just defined confidence scores into the forward step:

$$\bar{f}_k(x) = \sum_{i=1}^k c_i^r(x) f_i(x). \quad (6)$$

The overall model is trained to reduce the following loss

$$L_{\text{acc}}(x, y) = l(f_b(x), y) + l(\bar{f}_{b-1}(x), y) \quad (7)$$

We divide the loss into two parts to avoid the intermediate branches, the second term, overshadowing the final output, the first term, which can be small if the model is highly confident in the early stage of the computation. This training approach allows us to halt the inference process at a given exit  $i$  when the following criterion is met:

$$c_i(x) \prod_{k=1}^{i-1} (1 - c_k(x)) \geq \epsilon, \quad (8)$$

where  $\epsilon \in [0, 1]$  is a confidence threshold. In order to give representative values, the early exits must be regularized. To do that, the following regularization loss, which is added to Equation (7), is employed for each exit  $i < b$ :

$$\begin{aligned} \mathcal{R}_i(x, y) = & -c_i(x) \log(\mathbb{1}(y, \bar{f}_i(x))) \\ & - (1 - c_i(x)) \log(\mathbb{1}(y, \bar{f}_i(x))), \end{aligned} \quad (9)$$

with  $\mathbb{1}(y, \bar{f}_i(x))$  a function that returns 1 if the  $i$ -th exit's prediction is correct with respect to the true label  $y$ , otherwise 0. This function is not differentiable, and to overcome this we calculate these ground truth values before the training step associated to the sample  $x$ , without interfering with the gradient computations. Moreover, a sampling trick is used to ensure that the resulting model is not overconfident, and we remand to the original paper for further details.

### B. Neural Architecture Search

The NAS solutions can be categorized according to three different dimensions [36].

The first dimension is the *Search Space*, which defines the architecture to be represented. The most straightforward *Search Space* is *entire-structured* which is represented layer-wise with one node representing a layer. Motivated by hand-crafted architectures consisting of repeated motifs [37], the *cell-based Search Space* defines each node as a cell (also dubbed block), a group of layers. Other types of search space show cells with different levels (hierarchical) or identity transformations between the layers (morphism-based).

The second dimension is the *Search Strategy*, which establishes how to explore the search space. Common strategies are based on reinforcement learning (RL) [38], [39], gradient optimization [40], and evolutionary algorithms (EAs) [41], [42]. The latter, in particular the genetic methods, represents the most popular choice [43]. The Non-dominated Sorting Genetic Algorithm II (NSGA-II) is widely used in NAS solutions thanks to its ability to efficiently explore the search space [30]. It generates offspring through crossover and mutation operations, and then applies a non-dominated sorting procedure combined with a crowding distance mechanism to select the most optimal individuals for the next generation. This approach ensures that the selected architectures are not only characterized by high fitness but also maintain a diverse set of solutions, helping to avoid premature convergence. Moreover, it is a multi-objective algorithm, i.e., being able to optimize jointly different figures of merit, and it is able to effectively handle constraints [44]. The constraints are handled by redefining the objective function [45], and the simplest scheme works as follows:

$$\phi(x) = f(x) + pG(x) \quad (10)$$

where  $\phi(x)$  is the joint objective,  $f(x)$  is the original objective without constraints,  $p$  is the penalty term and  $G(x)$  is a function deciding whether to apply the penalty by accounting for the constraints. The main idea is that the application of a penalty results in a bad fitness in the genetic algorithm and precisely in a bad ranking in the case of NSGA-II proportional to the constraint violation. The penalty may be static, i.e., a constant value, or dynamic, i.e. adapted during the evolutionary process [45].

The third dimension is the *Performance Estimation Strategy*, which refers to the process of estimating the performance to be optimized. The purpose of this family of methods is to avoid training every single neural network architecture in the search space which would be prohibitive in terms of computational time. Weight sharing is a popular estimation strategy, and in particular, Once-For-All (OFA) [28] represents a state-of-the-art. OFA trains a supernet comprehensive of many network configurations once and evaluates the candidate networks by fine-tuning from the supernet weights during the search of the NAS. The training of the OFA supernet is performed by a technique called *progressive shrinking*, which first trains the largest neural network with maximum depth, width, and kernel size, then progressively fine-tunes the OFA network to support smaller sub-networks that share weights with the larger ones. While OFA provides three variants of CNN-based supernets, inspired by MobileNetV3, ProxylessNAS, and Resnet50 respectively, NASViT proposed a Vision Transformer-based supernet [46]. The choice of the supernet accounts for the different desired accuracy-efficiency trade-off in a target scenario. Another solution is the use of surrogate models such as a Gaussian Process [47] and Radial Basis Function [48]. For instance, MSuNAS [49] proposes a method called *adaptive-switching* which selects the best accuracy predictor among four possible surrogate models according to a correlation metric named Kendall's Tau [50] at each iteration of the NAS process. Moreover, another technique called *weight inheritance* (network morphism) [51] evaluates the performance of a network obtained by an identity transformation (morphism) on a parent model by inheriting the weights of the parent model, i.e., being initialized with the same weights.

## IV. THE PROPOSED NACHOS

This section introduces and details the presented NACHOS. More specifically, Section IV-A provides the problem formulation of NACHOS, Section IV-B provides an overall description of NACHOS, Section IV-C details the automatic placement of the EECs along the backbone, Section IV-D describes briefly the training phase by detailing the introduced regularization terms and the inference phase of the EENNs. Finally, Section IV-E introduces a novel figure of merit handling for the number of MACs of the EENN and the corresponding constraint enforced. To improve the readability, Table I contains all symbols we use to define the problem and our proposal.

### A. Problem formulation

The problem of NACHOS consists in selecting an EENN optimizing the classification accuracy and number of MAC operations performed by the EENNs at inference time while respecting constraints on these two figures of merit. Formally,

NACHOS can be defined as the following constrained optimization problem:

$$\begin{aligned} & \text{minimize } \mathcal{G}(F_A(\tilde{x}), F_M(\tilde{x})) \\ & \text{s. t. } F_A(\tilde{x}) > \bar{F}_A, \\ & \quad F_M(\tilde{x}) < \bar{F}_M, \\ & \quad \tilde{x} \in \Omega_{\tilde{x}} \end{aligned} \quad (11)$$

where  $\mathcal{G}$  is a multi-objective optimization function,  $\tilde{x}$  and  $\Omega_{\tilde{x}}$  represent a candidate neural network architecture and the search space of the neural network exploration, respectively,  $F_A(\tilde{x})$  and  $\bar{F}_A$  are the metric computing the top1 accuracy and the corresponding constraint,  $F_M(\tilde{x})$  and  $\bar{F}_M$  are, respectively, the metric computing the number of MACs of an EENN and the corresponding constraint. The optimization problem described in Equation (11) will be tackled by the proposed NACHOS framework that is introduced in the next Section.

### B. The overall view of NACHOS

This section introduces the proposed NACHOS, which is a NAS solution able to jointly optimize the design of EENN and its EECs.

An overall description of the NACHOS framework is given in Fig. 1. In more detail, NACHOS receives as input a dataset  $\mathcal{DS}$ , comprising a training set (used to train the candidate networks) and a validation set (used to validate the candidate networks), the set of constraints  $\{\bar{F}_A, \bar{F}_M\}$  enforced in the NAS exploration, the set  $\{\mu_1, \mu_2, \mu_3\}$  consisting of the number of epochs for the steps of the training process, the set  $\{\omega_1, \omega_2, \omega_3\}$  consisting of the weights of the training loss function  $L$ , the set  $\Omega_x$  containing the candidate backbones  $x$  of the EENNs sampled by the OFA supernet, and the set  $\Omega_\theta$  containing the hyperparameters  $\theta$  of the EECs.

In more detail, the *OFA Adapter* module receives as inputs  $\Omega_\theta$  and  $\Omega_x$  and enhances each  $x$ , candidate network of the OFA, by placing EECs along the backbone in the positions given by the configuration  $\theta$  sampled from the set  $\Omega_\theta$ . The output is a new set  $\Omega_{\tilde{x}}$  consisting of EENNs represented with  $\tilde{x}$ . Moreover, the *OFA Adapter* automatically designs the architecture of the EECs as introduced in Section IV-C.

The *Trainer* module takes in input the EENNs from the set  $\Omega_{\tilde{x}}$  and trains the EENNs on the  $\mathcal{DS}$  dataset. Each step of the training process is run for  $\mu_i$  epochs with  $i$  being the number of the step. Each component of the loss function is weighted by a weight  $\omega_i$ . After the training, the EENNs are evaluated by inference with a set of thresholds, and the best configuration of the thresholds is chosen according to a criterion accounting for the set of constraints. The output is a set, named the *archive*, consisting of the tuples  $\langle \tilde{x}, F_A(\tilde{x}), F_M(\tilde{x}) \rangle$ , representing the configuration of the EENN and the related figure of merits  $F_A(\tilde{x})$  and  $F_M(\tilde{x})$  on accuracy and computation, respectively. Before starting the search, a representative set of the OFA  $\Omega_{\tilde{x}}$ , is sampled. The number  $N_{start}$  of EENNs sampled is a parameter chosen by the user (usually  $N_{start} = 100$ ). This means that the cardinality of the archive is initially equal to  $N_{start}$ .

In NACHOS, the optimization problem defined in Eq. (11) is reformulated as follows by handling the constraints implicitly:

$$\begin{aligned} & \text{minimize } \mathcal{G}(\mathcal{S}_A(\tilde{x}), \mathcal{S}_{CM}(\tilde{x})) \\ & \text{s. t. } \tilde{x} \in \Omega_{\tilde{x}} \end{aligned} \quad (12)$$

where  $\mathcal{S}_A(\tilde{x})$  is the predicted classification accuracy of  $\tilde{x}$  computed by the accuracy predictor introduced in MSuNAS [49] and  $\mathcal{S}_{CM}(\tilde{x})$  is  $F_{CM}(\tilde{x})$ , introduced in Eq. (19), computed on  $\mathcal{S}_{\mathcal{M}}(\tilde{x})$ , the predicted  $F_M(\tilde{x})$  taken by the surrogate model (i.e., the MACs predictor).

The fundamental part of NACHOS is the *Search Strategy* module. It adopts the genetic algorithm NSGA-II [30] to solve the bi-objective problem introduced in Equation (12), by jointly optimizing the objectives  $\mathcal{S}_A(\tilde{x})$  and  $\mathcal{S}_{CM}(\tilde{x})$  on the dataset  $\mathcal{DS}$ , and by evaluating network architectures candidates from  $\Omega_{\tilde{x}}$ . The search process is iterative: at each iteration, the two surrogates model, computing  $\mathcal{S}_A(\tilde{x})$  and  $\mathcal{S}_{CM}(\tilde{x})$ , respectively, are chosen employing a mechanism called *adaptive-switching*, the same method used by MSuNAS [49], which selects the best surrogate model according to a correlation metric (i.e., Kendall's Tau [50]). The surrogate models are trained by using the *archive* as the dataset. Then, a ranking of the candidate networks, based on  $\mathcal{S}_A(\tilde{x})$  and  $\mathcal{S}_{CM}(\tilde{x})$ , is computed and a new set of candidates is obtained by NSGA-II and forwarded to the *Trainer*. The *Trainer* updates the *archive*, which becomes available for evaluation in the next iteration. The search ends if a number of iterations  $N_{iter}$  is reached or the hypervolume of the Pareto front is not changing for  $N_{tol}$  iterations (convergence).

At the end of the search, NACHOS returns the set of the  $k$  EENN architectures  $X^\circ = \{x_1^\circ, \dots, x_k^\circ\}$  characterized by the best trade-off among the objectives, being  $k$  a value specified by the user. We emphasize that the output of NACHOS is a set of networks enhanced with EECs, while the OFA given in input consists in regular neural networks (i.e., no EECs).

### C. The OFA Adapter module: automatic placement of EECs

This section details the OFA Adapter module whose main goal is the automatic placement of the EECs along the backbone by taking into account the computational cost of the EECs.

We define  $\gamma_i$  as the computational cost of exiting at  $i^{th}$  exit. This cost can be computed as the sum of the computational cost of the backbone up to the  $i^{th}$  exit and the computational cost of the  $i^{th}$  EEC. In our framework, the computational cost  $\gamma_i(\tilde{x})$  is computed as:

$$\gamma_i(\tilde{x}) = F_M(x_i) + F_M(EEC_i) \quad (13)$$

where  $x_i$  is the backbone  $x$  up to the  $i^{th}$  exit and  $EEC_i$  is the  $i^{th}$  EEC.

The basic design of these EECs is a sequence of ReLU activation functions and Convolutional layers followed by a flattening operation and a Fully-Connected layer [35].

The placement of the EECs is given by sampling from  $\Omega_\theta$  a binary encoding configuration  $\theta$  of length  $\bar{B} - 1$  being  $\bar{B}$

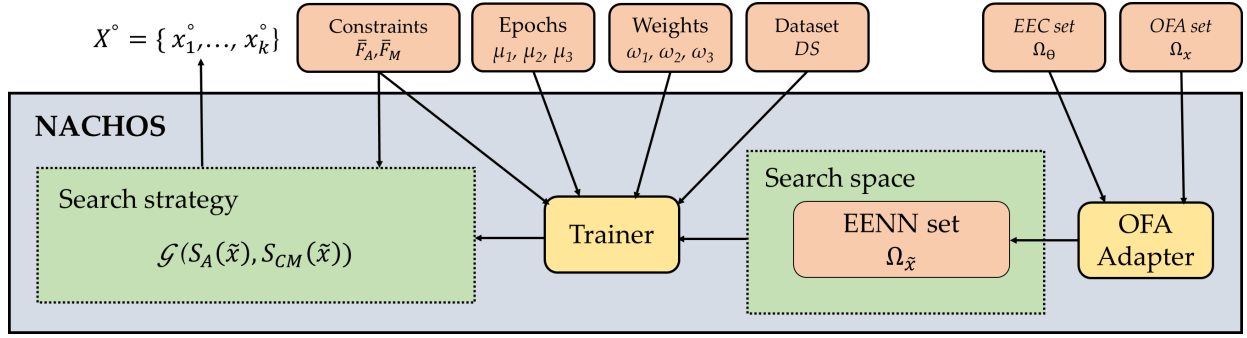


Fig. 1. The proposed NACHOS framework, which is composed of an OFA Adapter, Trainer and a Search Strategy module.

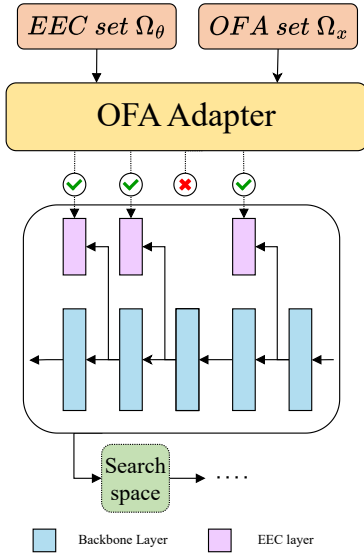


Fig. 2. The process of creating an EENN and how such models are used to populate the searching space. The best model is searched within this space. For further details, please refer to Section IV-C.

a user-defined parameter representing the maximum number of EECs, distributed along the backbone equidistantly and including the final classifier. The uniform placement ensures that the computational complexity is distributed uniformly along the backbone since the layers have similar computational complexity. For instance, with  $\bar{B} = 5$ , a sample  $\theta = [1, 0, 1, 0]$  extracted from  $\Omega_\theta$  means that EECs are placed at  $1/5$  and  $3/5$  of the backbone.

The steps of the placement of the EECs along the backbone are shown in Algorithm 1, while the process is visually depicted in Figure 2. In this Algorithm, the EECs are placed along the backbone  $x$  according to the configuration  $\theta$  (Line 2). The vector of computational costs  $\gamma$  is computed on  $x$  and EECs (Line 3). In the *For* statement each EEC is scanned and the condition  $\gamma_i \leq \gamma_{i+1}$  is checked (Lines 5-9). The last exit is the only one fixed a priori (i.e., basic design by default). If an EEC violates this condition, a *Max Pooling* layer is added through the operator  $\text{AddMaxPooling}(\cdot)$ , where the window of the *Max Pooling* layer is the minimum s.t. this EEC satisfies

TABLE I  
LIST OF THE MAIN SYMBOLS USED IN THE EXPERIMENTS

Symbol	Meaning
$\Omega_x$	Set of OFA backbones.
$\Omega_\theta$	Set of EEC hyperparameters.
$\Omega_{\tilde{x}}$	Set of EENNs built on top of $\Omega_x$ and $\Omega_\theta$ .
$x$	Neural network architecture (backbone) sampled by $\Omega_x$ .
$\theta$	The configuration of hyperparameters sampled by the EEC set.
$\epsilon$	The set of confidence thresholds.
$\tilde{x}$	EENN sampled by $\Omega_{\tilde{x}}$ .
$F_A(\tilde{x})$	The top1 accuracy of $\tilde{x}$ .
$F_M(\tilde{x})$	The number of MACs of $\tilde{x}$ .
$F_{CM}(\tilde{x})$	Novel formulation of $F_M(\tilde{x})$ endowed by handling the constraint $\bar{F}_M$ .
$\bar{F}_A$	The constraint on the figure of merit $F_A(\tilde{x})$ .
$\bar{F}_M$	The constraint on the figure of merit $F_M(\tilde{x})$ .
$S_A(\tilde{x})$	The top1 accuracy of $\tilde{x}$ predicted by the surrogate model.
$S_M(\tilde{x})$	The number of MACs of $\tilde{x}$ predicted by the surrogate model.
$S_{CM}(\tilde{x})$	The value of $F_{CM}(\tilde{x})$ computed on $S_M(\tilde{x})$ .
$U_i(\tilde{x})$	The ratio of samples exiting at the $i^{\text{th}}$ of $\tilde{x}$ .
$\gamma_i(\tilde{x})$	The computational cost of exiting at the $i^{\text{th}}$ of $\tilde{x}$ .
$\mu_i$	The number of epochs of the $i^{\text{th}}$ phase of the training process of an EENN.
$\omega_i$	The weight of the $i^{\text{th}}$ component of the loss training function.

the condition. Finally, the designed EECs are attached to the backbone leading to the creation of an EENN (Line 10) and the EENN is added to  $\Omega_{\tilde{x}}$  (Line 11).

Other works propose an alternative condition to handle the complexity of the EECs [2] [52]. This is applied after training when thresholds and percentage of samples exiting every EEC are known. Differently, our criterion is applied before training and is agnostic of the thresholds and exit ratios.

An OFA Adapter is present also in the EDANAS architecture. However, in the latter work, the structure of the EECs is fixed a priori and the module does not account for their complexity.

#### D. Training and inference of the EENNs in the Trainer module

This section explains how the training of the EENN and the inference are performed in the Trainer module. The training is performed in an end-to-end way optimizing the loss consisting

---

**Algorithm 1:** OFA Adapter update: the update of the  $\Omega_{\tilde{x}}$  performed by the OFA adapter

---

**Input:** backbone  $x$ , configuration  $\theta$ , the set of the EENNs  $\Omega_{\tilde{x}}$

- 1 Design a set of *EECs* based on  $x$  and  $\theta$ ;
- 2 Initialize a variable  $B$  with the number of *EECs*;
- 3 Define the vector  $\gamma$  of computational costs based on  $x$  and *EECs*;
- 4 **for**  $i = B - 1$  to 0 **do**
- 5     **if**  $\gamma_i > \gamma_{i+1}$  **then**
- 6          $EECs_i \leftarrow \text{AddMaxPooling}(EECs_i, \gamma_{i+1})$ ;
- 7     **end**
- 8 **end**
- 9 Attach *EECs* to the backbone  $x$  to create a EENN  $\tilde{x}$ ;
- 10 Add  $\tilde{x}$  to  $\Omega_{\tilde{x}}$ ;
- 11 **return**  $\Omega_{\tilde{x}}$ ;

---

of an accuracy loss and the regularization terms regarding the computational constraint and the operating point of the EECs.

---

**Algorithm 2:** Training of an EENN

---

**Input:** EENN  $\tilde{x}$ , dataset  $DS$ , constraints  $\{\overline{F}_A, \overline{F}_M\}$ , epochs  $\{\mu_1, \mu_2, \mu_3\}$ , weights  $\{\omega_1, \omega_2, \omega_3\}$

- 1 Split the dataset  $DS$  into a training and a validation set;
- 2 Get the backbone  $x$  from  $\tilde{x}$  with the pretrained weights from the OFA supernet;
- 3 Perform fine-tuning on  $x$  for  $\mu_1$  epochs;
- 4 Define the vector of computational costs  $\gamma$  from  $\tilde{x}$ ;
- 5 Train  $\tilde{x}$  applying a regularization  $L_{cost}$  based on  $\gamma$  and the confidences  $c_i$  for  $\mu_2$  epochs;
- 6 **for**  $\mu_3$  epochs **do**
- 7     Extract a set from the training set (Support Set);
- 8     Compute the average accuracies per class on the Support Set;
- 9     For each batch, applies the regularization  $L_{cost}$  and the regularization  $L_{peak}$  based on the average accuracies and the confidences  $c_i$  computed on the batch samples;
- 10 **end**
- 11 Set  $\tilde{x}$  with the weights for the best validation loss;
- 12 **return**  $\tilde{x}$ ;

---

The training performed by the Trainer module is shown in Algorithm 2. Firstly, the Trainer extracts the training and validation set from the dataset  $DS$  (Line 1). Then, it performs fine-tuning on the backbone of the EENN for a number of epochs  $\mu_1$  starting from the pre-trained weights of the OFA supernet (Lines 2-3). Once trained the backbone, the Trainer trains the EENN with the soft exit mechanism presented in Section III-A with two regularization terms,  $L_{cost}$  and  $L_{peak}$ . In particular, we can distinguish between a phase where only a regularization  $L_{cost}$  on computational costs is applied for a

$\mu_2$  number of epochs (Lines 4-5) and a second phase where an additional regularization  $L_{peak}$  for enhancing the accuracy is applied lasting  $\mu_3$  number of epochs. This is done to trigger the latter regularization only after the EECs have been trained for a sufficient number of epochs. (Lines 6-10). We then select the weights corresponding to the best loss computed on the validation set (Line 11).

The overall train loss function is the following:

$$L = \omega_1 * L_{acc} + \omega_2 * L_{cost} + \omega_3 * L_{peak} \quad (14)$$

where:

- $L_{acc}$  is the accuracy term shown in Equation (7)
- $L_{cost}$  is the regularization term to enforce the constraint on the complexity shown in Equation (17)
- $L_{peak}$  is the regularization term introduced to ensure that the confidences of the EECs follow their operating point shown in Equation (18)

The weights  $\omega_1$ ,  $\omega_2$  and  $\omega_3$ , given by the user (see *Weights* in Fig. 1), are fixed in our setting but they could be learned during training to provide a simple way of forcing a desired trade-off between accuracy and energy-efficiency. We emphasize that the number of epochs of each step of the training is provided by the user as shown in Fig. 1.  $\omega_3$  is set to zero to avoid applying the  $L_{peak}$  during the step following the fine-tuning. In the rest of the Section, we explain the aforementioned regularization terms.

First, we examine the regularization to enforce the constraint on the complexity term. We aim at training more the fraction of the EENN not violating the constraint. As introduced in EDANAS [27], the metric  $F_M(\tilde{x})$  can be computed as the sum of  $\gamma_i$  weighted by the utilization of the corresponding *EEC* <sub>$i$</sub>  and can be computed as follows:

$$F_M(\tilde{x}) = \sum_{i=1}^{B+1} \gamma_i(\tilde{x}) U_i(\tilde{x}) \quad (15)$$

where  $U_i$  accounts for the percentage of samples classified by the  $i$ -th classifier. For simplicity, we use the operator  $F_M(\cdot)$  to compute the MACs of both EENNs and neural networks without EECs, where  $U_{B+1}(x) = 1$  and  $B = 0$ .

Following the considerations of a previous work [35], we can define  $\tilde{\gamma}_i(\tilde{x})$  as the recursive definition of the previously defined computational cost  $\gamma_i(\tilde{x})$  in an analogous way to the soft exit mechanism. To take into account the energy efficiency, since  $\tilde{\gamma}_i(\tilde{x})$  is differentiable, we can include it as a regularization term, training our network on a weighted combination of the costs:

$$\tilde{\gamma}_i(\tilde{x}) = c_i \gamma_i(\tilde{x}) + (1 - c_i) \tilde{\gamma}_{i+1}(\tilde{x}) \quad (16)$$

where  $B$  is the number of EECs and  $\tilde{\gamma}_{B+1}(\tilde{x})$  is the computational cost of exiting at the final exit of the EENN  $\tilde{x}$ . For simplicity, we refer to  $c_i$  as the confidence output at the  $i^{th}$  EEC for a generic sample.

To compute the regularization, we adopt  $\tilde{\gamma}_{B+1}(\tilde{x})$  instead of  $F_M(\tilde{x})$  since the ratio of exiting samples from a branch

can be computed only at inference time. In the loss  $L_{cost}$ , the local confidences  $c_i$  act as a proxy of these ratios.

The loss function  $L_{cost}$  is implemented as follows:

$$L_{cost} = \frac{\max(0, \tilde{\gamma}_{B+1}(\tilde{x}) - \bar{F}_M)}{(F_M(x) - \bar{F}_M)} \quad (17)$$

where  $F_M(x)$  is the number of MACs computed by the backbone  $x$  (i.e., no EECs) to process the input and  $\bar{F}_M$  is the constraint enforced on  $F_M(\tilde{x})$ .

In the right term of Equation (17), the numerator relies on a max operator to apply the penalization in case of a violation of the constraint and the denominator computes the maximum possible violation since  $F_M(x)$  is equal to  $F_M(\tilde{x})$  in the worst case, i.e., when all samples exit from the last EEC. In this perspective,  $L_{cost}$  can be seen as the ratio of the constraint violation.

Second, we explain the regularization introduced to force the confidences to follow the operating point of the EENN. We aim at improving the accuracy of the EENN by tuning the confidences while considering the performance on a representative set. First, part of the original training set, containing an equal number of samples for each class, is extracted. We call this extracted set as *Support Set* and the number of classes as  $C$ . The *Support Set* is used to compute a matrix  $C \times B$  named *Support Matrix*. The *Support Matrix* contains the average accuracies per class computed for each epoch on top of the support set. For each class, the average accuracy is the ratio of the corrected predicted samples in the set of samples of a given class in the support set. Once defined the *Support Matrix*, we use a regression loss as a regularization term to force the local certainties (calculated on the batch samples) to follow the global accuracies (calculated on the *Support Set*). Basically, we compute the Mean Squared Error (MSE) on two vectors of  $B$  values computed per each sample as follows:

$$L_{peak} = \text{MSE}(c_{local}, SM_j) \quad (18)$$

where the first vector,  $c_{local}$  is the vector of local certainties and the second vector,  $SM_j$  is the  $j$ -th row of the *Support Matrix* corresponding to the label of the class of that sample. We initially train the EENN without the  $L_{peak}$  term to train all the EECs and then apply this loss after  $\mu_1 + \mu_2$  epochs when the EECs are working properly.

We now detail how the inference of the EENNs and the automatic setting of the thresholds are performed in NACHOS.

The inference is performed as in the scheme shown in Equation (8). If the confidence threshold in the prediction is met, a prediction is computed and the sample is not forward propagated. The confidence measure is computed as the Cumulative in the prediction [35].

The thresholds are automatically set according to the constraints regarding the accuracy and the number of MACs. This is done by at first selecting the configuration of thresholds leading to the best accuracy in inference through a grid search. Then, if possible, starting from the second last exit the thresholds are reduced up to when the number of MACs is less than the constraint value while respecting the constraint

regarding the accuracy. Note that the lower the thresholds, the higher the number of samples exiting from that EEC, hence reducing the overall number of MACs.

### E. MACs' formulation in the Search Strategy module

In this section, we describe the novel figure of merit named  $F_{CM}(\tilde{x})$  that endows  $F_M(\tilde{x})$  by handling the  $\bar{F}_M$  constraint. This objective is optimized by the Search Strategy module. We define  $\phi$  as the ratio of the admissible solutions in the population during the evolutionary process in the Search Strategy. The  $F_{CM}(\tilde{x})$  objective can be formulated as follows:

$$F_{CM}(\tilde{x}) = \phi F_M(\tilde{x}) + (1 - \phi) \max(0, F_M(\tilde{x}) - \bar{F}_M) \quad (19)$$

The max operator eventually applies a penalty as in Equation (17).  $\phi$  can be seen as an adaptive penalty term changing during the execution of the genetic algorithm. Typically  $\phi$  is  $\cong 0$  at the early stages of the NAS when there are a lot of unfeasible solutions and is  $\cong 1$  at the final stages when there are more feasible solutions. The rationale is that  $F_{CM}(\tilde{x})$  aims at applying a stronger penalization to rapidly eliminate unfeasible solutions at the early stages and focus on optimizing the  $F_M(\tilde{x})$  at the final stages.

## V. EXPERIMENTAL RESULTS

This section shows the results of the experimental campaign aiming at assessing the effectiveness of NACHOS. In more detail, the purpose of this section is to point out that NACHOS can design effective network architectures (i.e., providing high accuracy for the considered task), competitive with respect to state-of-the-art EENNs that have been manually designed by experts, while satisfying a constraint on the number of Multiply and Accumulate (MAC) operations performed by the EENNs at inference time.

### A. Dataset

The datasets used in the experimental comparisons are CIFAR-10 [53], CINIC-10 [54], SVHN [55], and Imagenette [56], a subset of 10 easily classified classes from Imagenet [8], whose number of classes is 1000 and image resolution is 224. The first three datasets consist of  $32 \times 32$  color images while the resolution of Imagenette is  $160 \times 160$ , and represent 10 classes of objects. Regarding CIFAR-10, the training set comprises 60000 images, while the testing set comprises 10000 images. Regarding CINIC-10, both the training set and the test set comprise 90000 images. Regarding SVHN, the training set comprises 73257 images, while the testing set comprises 26032 images. Regarding Imagenette, the training set comprises 9469 images while the testing set consists of 3925 images.

### B. Experimental setup

The setup of the experiments is the following:

- the 1st objective is  $F_A(\tilde{x})$  and the 2nd objective is  $F_M(\tilde{x})$ ;
- the training epochs are set as  $\{\mu_1 = 10, \mu_2 = 5, \mu_3 = 5\}$ ;

TABLE II  
RESULTS FOR THE EXPERIMENTAL CAMPAIGN OF NACHOS ON ALL THE DATASETS. COLUMN LABELS EXPLAINED IN TABLE I.

Dataset	Model	NAS	B	$F_M(\bar{x})$ (M)	$F_M(x)$ (M)	$F_A(\bar{x})$ (%)	$F_A(x)$ (%)	$\epsilon$	Confidence
SVHN	NACHOS	yes	$3.00 \pm 0.71$	<b><math>1.46 \pm 0.03</math></b>	$4.46 \pm 0.18$	$79.96 \pm 0.84$	$88.62 \pm 0.23$	$[0.26 \pm 0.05, 0.13 \pm 0.10, 0.10, -]$	Cumulative
	EDANAS	yes	$4.60 \pm 0.55$	<u><math>1.47 \pm 0.02</math></u>	$5.95 \pm 0.05$	$77.98 \pm 0.24$	$86.08 \pm 0.19$	$[0.20 \pm 0.00, 0.14 \pm 0.05, 0.10 \pm 0.00, 0.10 \pm 0.00]$	Score margin
	EEAlexnet	no	5.00	263.51	464.48	<b>89.00</b>	93.10	$[0.70, 0.60, 0.60, 0.60]$	Cumulative
	EEResnet20	no	10.00	70.20	315.26	<u>88.53</u>	93.28	$[0.70, 0.60, 0.60, 0.60, 0.60, 0.60, 0.60, 0.60, 0.60]$	Cumulative
CIFAR-10	NACHOS	yes	$2.80 \pm 0.84$	<b><math>2.44 \pm 0.02</math></b>	$5.78 \pm 0.19$	$72.65 \pm 0.49$	$71.68 \pm 0.31$	$[0.20 \pm 0.00, 0.24 \pm 0.05, 0.10 \pm 0.00, -]$	Cumulative
	EDANAS	yes	$4.60 \pm 0.55$	<u><math>2.47 \pm 0.01</math></u>	$6.00 \pm 0.04$	$67.78 \pm 0.12$	$67.48 \pm 0.23$	$[0.20 \pm 0.00, 0.18 \pm 0.04, 0.12 \pm 0.04, 0.10 \pm 0.00]$	Score margin
	EENCN	no	2.00	7.48	25.14	<b>76.50</b>	77.20	$[0.10]$	Max softmax
	EEAlexnet	no	5.00	233.46	464.48	<u>73.03</u>	83.85	$[0.70, 0.60, 0.60, 0.60]$	Cumulative
CINIC-10	NACHOS	yes	$2.80 \pm 0.84$	<b><math>2.42 \pm 0.01</math></b>	$5.87 \pm 0.10$	$60.85 \pm 0.30$	$59.65 \pm 0.11$	$[0.20 \pm 0.04, 0.22 \pm 0.04, 0.10 \pm 0.00, -]$	Cumulative
	EDANAS	yes	$4.20 \pm 0.45$	<u><math>2.45 \pm 0.03</math></u>	$6.03 \pm 0.02$	$60.65 \pm 0.16$	$59.67 \pm 0.17$	$[0.20 \pm 0.00, 0.18 \pm 0.04, 0.10 \pm 0.00, 0.10 \pm 0.00]$	Score margin
	EEAlexnet	no	5.00	300.20	464.48	<b>64.35</b>	73.30	$[0.80, 0.80, 0.80, 0.80]$	Cumulative
	EEResnet20	no	10.00	142.57	315.26	<u>61.69</u>	69.01	$[0.70, 0.60, 0.60, 0.60, 0.60, 0.60, 0.60, 0.60, 0.60]$	Cumulative
Imagenette	NACHOS	yes	$3.00 \pm 0.71$	<u><math>92.34 \pm 0.58</math></u>	$104.09 \pm 1.09$	<u><math>94.89 \pm 0.15</math></u>	$95.53 \pm 0.08$	$[-, 0.63 \pm 0.06, 0.68 \pm 0.04, 0.65 \pm 0.07]$	Cumulative
	EDANAS	yes	$4.40 \pm 0.55$	<b><math>81.92 \pm 1.11</math></b>	$231.34 \pm 4.49$	<u><math>90.76 \pm 0.54</math></u>	$96.53 \pm 0.38$	$[0.10 \pm 0.00, 0.20 \pm 0.00, 0.14 \pm 0.05, 0.15 \pm 0.07]$	Score margin
	EEEfficientNet-B0	no	4.00	182.77	201.13	<b>95.14</b>	97.25	$[0.50, 0.50, 0.50]$	Cumulative

TABLE III  
COMPUTATIONAL COSTS ( $\gamma$ ) AND UTILIZATION (U) FOR THE EXPERIMENTAL CAMPAIGN OF NACHOS ON THE FOUR DATASETS.

Dataset	Model	Computational Costs $\gamma$ (M)	Utilization U (%)
SVHN	NACHOS	$[1.12 \pm 0.08, 1.85 \pm 0.11, 2.5 \pm 0.00, -, 4.46 \pm 0.18]$	$[58.78 \pm 21.89, 39.91 \pm 24.51, 0.00 \pm 0.00, -, 1.31 \pm 2.92]$
	EDANAS	$[1.13 \pm 0.12, 1.51 \pm 0.06, 2.07 \pm 0.04, 4.01 \pm 0.03, 5.95 \pm 0.05]$	$[24.97 \pm 16.69, 82.94 \pm 20.04, 1.28 \pm 2.64, 0.24 \pm 0.43, 0.56 \pm 0.55]$
	EEAlexnet	$[87.27, 201.59, 380.82, 440.54, 464.48]$	$[23.00, 34.00, 27.00, 10.00, 6.00]$
	EEResnet20	$[19.57, 24.73, 29.89, 35.06, 43.99, 64.63, 85.28, 158.97, 239.48, 315.26]$	$[0.00, 0.00, 0.00, 34.00, 13.00, 34.00, 11.00, 2.00, 2.00, 4.00]$
CIFAR-10	NACHOS	$[1.60 \pm 0.0, 2.18 \pm 0.31, 3.40 \pm 0.44, -, 5.78 \pm 0.19]$	$[41.31 \pm 30.37, 74.31 \pm 30.32, 0.00 \pm 0.00, -, 0.90 \pm 1.52]$
	EDANAS	$[1.11 \pm 0.02, 1.49 \pm 0.04, 2.03 \pm 0.02, 4.03 \pm 0.03, 6.00 \pm 0.04]$	$[0.87 \pm 1.17, 24.58 \pm 7.56, 51.06 \pm 12.65, 18.78 \pm 10.17, 5.08 \pm 4.65]$
	EEAlexnet	$[87.27, 201.59, 380.82, 440.54, 464.48]$	$[38.00, 27.00, 19.00, 10.00, 6.00]$
	EEResnet20	$[19.57, 24.73, 29.89, 35.06, 43.99, 64.63, 85.28, 158.97, 239.48, 315.26]$	$[6.00, 0.00, 0.00, 33.00, 16.00, 18.00, 8.00, 9.00, 5.80, 4.20]$
CINIC-10	NACHOS	$[1.80 \pm 0.20, 2.45 \pm 0.50, 3.70 \pm 0.42, -, 5.87 \pm 0.10]$	$[69.75 \pm 26.03, 30.25 \pm 26.03, 0.00 \pm 0.00, -, -]$
	EDANAS	$[1.12 \pm 0.00, 1.48 \pm 0.04, 2.09 \pm 0.03, 4.06 \pm 0.05, 6.03 \pm 0.02]$	$[3.30 \pm 0.00, 25.74 \pm 9.86, 52.38 \pm 6.46, 15.8 \pm 6.79, 5.42 \pm 3.48]$
	EEAlexnet	$[87.27, 201.59, 380.82, 440.54, 464.48]$	$[16.00, 34.00, 13.00, 20.00, 17.00]$
	EEResnet20	$[19.57, 24.73, 29.89, 35.06, 43.99, 64.63, 85.28, 158.97, 239.48, 315.26]$	$[0.00, 0.00, 0.00, 0.00, 14.00, 24.00, 20.00, 0.40, 28.00, 10.00]$
Imagenette	NACHOS	$[-, 46.00 \pm 5.29, 59.81 \pm 3.63, 70.21 \pm 6.78, 104.09 \pm 1.09]$	$[4.67 \pm 3.33, 14.31 \pm 9.70, 26.02 \pm 14.98, -, 72.42 \pm 7.22]$
	EDANAS	$[59.58 \pm 1.90, 79.01 \pm 1.83, 92.86 \pm 1.80, 149.77 \pm 4.38, 231.34 \pm 4.49]$	$[0.37 \pm 0.35, 96.76 \pm 1.86, 0.70 \pm 0.47, 2.10 \pm 0.42, 1.48 \pm 0.91]$
	EEEfficientNet-B0	$[50.74, 103.16, 140.96, 201.13]$	$[0.00, 6.34, 20.19, 73.47]$

TABLE IV  
ECE SCORES FOR THE EXPERIMENTAL CAMPAIGN OF NACHOS ON CIFAR-10 DATASET.

Model	ECE scores (%)
NACHOS	$[18.23 \pm 0.00, 15.70 \pm 0.22, 4.53 \pm 0.15, -, 1.60 \pm 0.11]$
EEAlexnet	$[13.79, 40.55, 59.07, 71.44, 2.26]$
EEResnet20	$[16.51, 89.80, 77.80, 28.39, 62.35, 58.10, 46.60, 66.65, 62.25, 2.30]$

- the weights of the training loss function shown in Equation (14) are set as  $\{\omega_1 = 1.0, \omega_2 = 1.0, \omega_3 = 1.0\}$  and  $\omega_3 = 0.0$  in the step following the fine-tuning (i.e., no  $L_{peak}$  is applied);
- the initial number of samples populating the archive  $N_{start}$  is set to 100;
- the maximum number of iterations  $N_{iter}$  is 30 and the number of epochs to trigger the convergence criterion  $N_{tol}$  is 10;
- considering our constrained mobile setting, we chose the pretrained MobileNetV3-based supernet, being the lightest in terms of MACs. We highlight that our framework is flexible enough to be applied to other CNN-based types of supernets. The hyperparameters of  $\Omega_x$  refer to the depth,

the kernel size, and the expansion rate (the architecture search space and building blocks are detailed in [37]) while the resolution is fixed at 32 which is the size of the images of the datasets used in the experimental results;

- the number of exit points (i.e., the points of the network where EECs are possibly placed) is set to 4 since it is a reasonable arbitrary number of EECs in an EENN and is also chosen by EDANAS;
- the constraints are set as  $\bar{F}_A = 65\%$ ,  $\bar{F}_M = 2.7M$  on CIFAR-10;  $\bar{F}_A = 50\%$ , and  $\bar{F}_M = 2.7M$  on CINIC-10;  $\bar{F}_A = 75\%$ , and  $\bar{F}_M = 1.5M$  on SVHN;  $\bar{F}_A = 90\%$ , and  $\bar{F}_M = 93M$  on Imagenette. The constraints are chosen by considering the results of EDANAS on the different datasets to test whether the searched models perform better in a similar setting.
- the computational resources required: search time of 2 GPU days on a single GPU (NVIDIA A40) for both NACHOS and EDANAS. Their search times are the same since NACHOS takes more training time in the evaluation of the NNs but converges faster in terms of NAS epochs thanks to constraints as shown in Fig. 3.
- the results of all the tested NAS methods, NACHOS and EDANAS, reports the mean and the standard deviation

obtained by aggregating the results of the search over 5 different seeds following the best practice for NAS evaluations suggested in [57]. We omitted the mean and standard deviation of  $\gamma$ ,  $\epsilon$ , and  $U$  when the corresponding early exit has never been selected among the runs and placed a '-' instead.

Additionally, we report the hyperparameters of the genetic algorithm NSGA-II:

- population size: 40;
- initial population: non-dominated solutions of the previous NAS iteration;
- crossover operator: two-point crossover with crossover probability=0.9;
- mutation operator: polynomial mutation with distribution index=1.0;
- selection operator: tournament selection with crowding distance sorting;
- number of generations: 20;
- survival: only the eight best solutions are kept according to ranking (Pareto dominance) and crowding distance sorting.

This setup allows effective convergence toward optimal trade-offs while preserving the diversity of the solutions, reducing the risk of premature convergence.

$N_{start}$ ,  $N_{iter}$ , the hyperparameters of the genetic algorithm, and the training parameters such as batch size and learning rate follow the experimental setup of MSuNAS and EDANAS.

The output set of optimal architectures  $X^\circ$  of each search contains one EENN. This architecture is Pareto Optimal using the best trade-off between  $F_A(\tilde{x})$  and  $F_M(\tilde{x})$  as the selection criterion. We refer to these networks as NACHOS.

As a comparison, we considered the following state-of-the-art EENNs.

- 1) the EENN found by NACHOS without using the  $L_{cost}$  regularization by selecting the best trade-off. We refer to this network as NACHOS\_NC (NACHOS No Constraints).
- 2) the EENN found by NACHOS without using the  $L_{peak}$  regularization by selecting the best trade-off. We refer to this network as NACHOS\_NP (NACHOS No Peak).
- 3) the EENN found by EDANAS by selecting the best trade-off. We refer to these networks as EDANAS.
- 4) an EENN [58] with an auxiliary classifier consisting of two fully connected layers with 1024 and 10 neurons placed after the first max-pooling. We refer to this network as EECNN;
- 5) AlexNet, and ResNet-20 [59] modified to handle  $32 \times 32$  images with multiple EECs [35]. An EEC is placed after each layer for AlexNet, resulting in 5 EECs (including the last one), while the branches in ResNet20 are placed after each block, resulting in a total of 10 branches. Each branch is composed of a convolutional layer of 128 filters, followed by a max pooling operation, if the dimensions of the image are big enough. This block is followed by a ReLU activation function and




Fig. 3. Comparison between the Pareto fronts of runs of NACHOS enforcing and not enforcing constraints in the search on CIFAR-10 respectively, while setting the same seed.

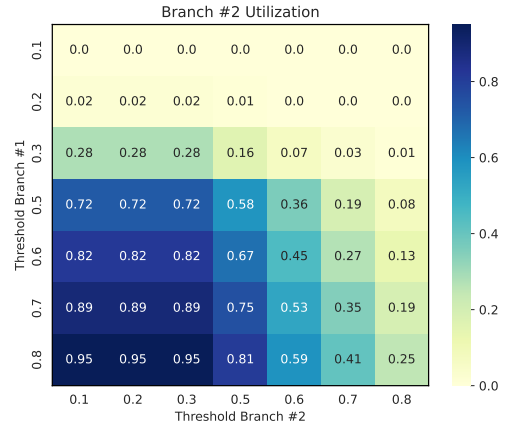


Fig. 4. Heatmap of the utilization of the second EEC of a model with three exits ( $B=3$ ) found by NACHOS\_NC varying the thresholds of the first two EECs.

a classification linear layer. We refer to these networks as EEAlexNet and EEResnet20, respectively.

- 6) EfficientNet-B0 [60] with 3 EECs placed by the computational-aware procedure explained in Algorithm 1. We refer to this network as EEEfficientNet-B0.

EDANAS has been run with the same configuration shown in [27] but differently, the resolution of the networks has been fixed to  $160 \times 160$  on Imagenette and  $32 \times 32$  on the other datasets.

The expert-designed networks have been trained with the following methodologies. EECNN has been trained by joint training. EEAlexNet, EEResnet20, and EEEfficientNet-B0 have been trained by the same differentiable training adopted in NACHOS presented in the work [35]. In particular, the backbone is initially trained from scratch (i.e., no fine-tuning) for  $\mu_1$  epochs, and then the standalone  $L_{acc}$  loss function is optimized for  $\mu_2 + \mu_3$  epochs during the training of the whole EENN. The thresholds are chosen with the same grid search on the same intervals used in NACHOS but considering no constraints.

We explain briefly the choice of the benchmarks. EDANAS is the only other framework in the literature optimizing jointly the backbone and the EECs. EECNN is a valid example of a CNN with a good trade-off between number of MACs and top1 accuracy. EEAlexNet, EEResnet20, and EEEfficientNet-B0 trained for the same number of epochs of NACHOS (differently from EECNN), are taken to ensure fair benchmarking with state-of-the-art EENNs. Additionally, NACHOS is tested without  $L_{peak}$  and  $L_{cost}$  regularizations to verify their effectiveness in the framework.

### C. Analysis of the results

The experimental results of NACHOS are presented in Table II. The column NAS specifies if the network is designed by a NAS procedure (yes) or not (no). The column B specifies the number of EECs (including the final exit). The column  $F_M(\tilde{x})$  shows the number of MACs in EE inference <sup>2</sup>. The column  $F_A(\tilde{x})$  shows the top1 accuracy in EE inference. The column  $F_A(x)$  shows the top1 accuracy in inference by using only the backbone of the EENN. The column  $\epsilon$  shows the thresholds of the EECs (apart from the last classifier not having a thresholding). The column Confidence shows the technique used to estimate the confidence in the prediction.

Additionally, the computational costs  $\gamma_i$ , i.e., the amounts of MACs to process an input at the  $i^{th}$  exit of an EENN, and the utilizations  $U_i$ , i.e., the exit ratios of the EECs, of the models presented in the aforementioned tables are reported in Table III. We highlight that the costs of EEAlexnet and EEResnet20 are the same for SVHN, CIFAR-10, and CINIC-10 since they refer to the same model.

We highlight that the high utilizations of some EECs in the presented experiments are due to the strong constraint on the computational demands.

Comparing NACHOS with the other state-of-the-art EENNs we show that NACHOS is effective in selecting a competitive network in terms of best trade-off between the two objectives of the search. In particular, NACHOS is able to find an EENN with similar computational demand to EDANAS with similar accuracy on CINIC-10 while improving the accuracy on CIFAR-10, SVHN, and Imagenette. Moreover, NACHOS is competitive with the expert-designed models on CIFAR-10, CINIC-10, and Imagenette while it performs worse on

<sup>2</sup>The number of Floating Point Operations (FLOPS), another popular metric regarding inference time, can be approximately computed by the number of MACs as  $FLOPS \approx 2 \times MACs$

SVHN. We highlight that the computational demand of the expert-designed models is significantly higher with respect to NACHOS.

Additionally, the Expected Calibration Error (ECE) scores for the experimental campaign of NACHOS on CIFAR-10 dataset are presented in Table IV. ECE is a popular metric to measure the calibration of the EECs [61]. The calibration is an important aspect of EENNs related to their capability of estimating properly the confidences on input samples.

Table IV shows that the ECE scores of the NAS model are good pointing out that well-calibrated EENNs are found by the NACHOS method. However, we expect calibration to play a major role in NACHOS when considering networks with higher computational demands since, for instance, deeper networks tend to be more miscalibrated [62] (e.g., with Resnet-based OFA). This aspect will be addressed in future works of the paper.

### D. Ablation studies

This section tests the effectiveness of the regularizations  $L_{cost}$  and  $L_{peak}$  introduced in Section IV-D.

The results of the models NACHOS, NACHOS\_NC and NACHOS\_NP found on the four datasets are reported in Table V. Additionally, the computational costs and the utilizations of the EECs of the models NACHOS, NACHOS\_NC, and NACHOS\_NP found on the four datasets are reported in Table VI.

In the NACHOS experiments, comparing NACHOS with NACHOS\_NC, we aim at testing the effectiveness of the  $L_{cost}$  regularization in selecting an EENN satisfying the given constraints. The results show that the constraints are satisfied in NACHOS while are not satisfied in NACHOS\_NC. Additionally, the model found shows even better accuracy on CIFAR-10 and SVHN. Additionally, Fig. 3 points out that NACHOS by using constraints and the  $L_{cost}$  regularization is able to reach faster the region of solutions satisfying the constraints, delimited by the red horizontal line and the purple vertical line in the figure, with respect to the unconstrained counterpart. Finally, we can argue that a NAS by enforcing constraints is able to find an admissible solution with a significant lower number of iterations. Indeed, NACHOS and EDANAS models are found by NACHOS in 10 iterations and by EDANAS in 30 iterations, respectively. Fig. 4 shows the utilization of the second EEC of a model with three exits (B=3) found by NACHOS\_NC on CIFAR-10 varying the thresholds of the first two EECs. This visualizes how the thresholds affect the behaviour of the Cumulative method in triggering the early exit in the EECs. In the NACHOS experiments, comparing NACHOS with NACHOS\_NP we aim at testing the effectiveness of the  $L_{peak}$  regularization in improving the accuracy of the EENN. An increase in accuracy is seen on CIFAR-10, SVHN, and Imagenette while the accuracy is approximately the same on CINIC-10 and maintains a similar computational complexity.

Lastly, we report in Fig. 5 the distributions of the confidences in NACHOS, NACHOS\_NP, and NACHOS\_NC mod-

TABLE V  
RESULTS FOR THE ABLATION STUDIES OF NACHOS ON ALL DATASETS. COLUMN LABELS ARE EXPLAINED IN TABLE I.

Dataset	Model	NAS	B	$F_M(\bar{x})$ (M)	$F_M(x)$ (M)	$F_A(\bar{x})$ (%)	$F_A(x)$ (%)	$\epsilon$	Confidence
SVHN	NACHOS	yes	$3.00 \pm 0.71$	$1.46 \pm 0.03$	$4.46 \pm 0.18$	$79.96 \pm 0.84$	$88.62 \pm 0.23$	$[0.26 \pm 0.05, 0.13 \pm 0.10, 0.10, -]$	Cumulative
	NACHOS_NC	yes	$2.80 \pm 0.84$	$1.95 \pm 0.10$	$4.74 \pm 0.08$	<b><math>84.59 \pm 0.47</math></b>	$88.24 \pm 0.27$	$[0.67 \pm 0.06, 0.16 \pm 0.05, 0.10 \pm 0.00, -]$	Cumulative
	NACHOS_NP	yes	$2.80 \pm 0.84$	<b><math>1.45 \pm 0.02</math></b>	$4.62 \pm 0.18$	$76.94 \pm 0.63$	$89.04 \pm 0.43$	$[0.18 \pm 0.08, 0.07 \pm 0.11, 0.10 \pm 0.00, -]$	Cumulative
CIFAR-10	NACHOS	yes	$2.80 \pm 0.84$	$2.44 \pm 0.02$	$5.78 \pm 0.19$	<b><math>72.65 \pm 0.49</math></b>	$71.68 \pm 0.31$	$[0.20 \pm 0.00, 0.24 \pm 0.05, 0.10 \pm 0.00, -]$	Cumulative
	NACHOS_NC	yes	$3.00 \pm 0.70$	$2.91 \pm 0.11$	$6.43 \pm 0.12$	$72.25 \pm 0.22$	$71.38 \pm 0.33$	$[0.65 \pm 0.06, 0.16 \pm 0.05]$	Cumulative
	NACHOS_NP	yes	$2.40 \pm 0.55$	<b><math>2.40 \pm 0.05</math></b>	$6.25 \pm 0.42$	$71.38 \pm 0.47$	$68.23 \pm 0.59$	$[-, 0.26 \pm 0.05, 0.15 \pm 0.07, -]$	Cumulative
CINIC-10	NACHOS	yes	$2.80 \pm 0.84$	$2.42 \pm 0.01$	$5.87 \pm 0.10$	<b><math>60.85 \pm 0.30</math></b>	$59.65 \pm 0.11$	$[0.20 \pm 0.04, 0.22 \pm 0.04, 0.10 \pm 0.00, -]$	Cumulative
	NACHOS_NC	yes	$3.00 \pm 0.70$	$3.74 \pm 0.13$	$6.85 \pm 0.11$	$59.96 \pm 0.12$	$59.51 \pm 0.29$	$[0.70 \pm 0.00, 0.60 \pm 0.17, 0.58 \pm 0.10, -]$	Cumulative
	NACHOS_NP	yes	$2.40 \pm 0.55$	<b><math>2.41 \pm 0.06</math></b>	$4.79 \pm 0.12$	$60.16 \pm 0.34$	$59.63 \pm 0.16$	$[0.20 \pm 0.00, 0.24 \pm 0.05, -]$	Cumulative
Imagenette	NACHOS	yes	$3.00 \pm 0.71$	$92.34 \pm 0.58$	$104.09 \pm 1.09$	<b><math>94.89 \pm 0.15</math></b>	$95.53 \pm 0.08$	$[-, 0.63 \pm 0.06, 0.68 \pm 0.04, 0.65 \pm 0.07]$	Cumulative
	NACHOS_NC	yes	$3.20 \pm 0.45$	<b><math>118.01 \pm 1.47</math></b>	$121.30 \pm 3.15$	<b><math>97.36 \pm 0.11</math></b>	$96.46 \pm 0.25$	$[-, 0.70 \pm 0.00, 0.68 \pm 0.04, 0.66 \pm 0.05]$	Cumulative
	NACHOS_NP	yes	$2.40 \pm 0.55$	$84.38 \pm 0.33$	$98.88 \pm 0.41$	$85.24 \pm 0.15$	$96.65 \pm 0.25$	$[-, 0.15 \pm 0.06, 0.17 \pm 0.06, -]$	Cumulative

TABLE VI  
COMPUTATIONAL COSTS ( $\gamma$ ) AND UTILIZATION (U) FOR THE ABLATION STUDIES OF NACHOS ON THE FOUR DATASETS.

Dataset	Model	Computational Costs $\gamma$ (M)	Utilization U (%)
SVHN	NACHOS	$[1.12 \pm 0.08, 1.85 \pm 0.11, 2.53 \pm 0.00, -, 4.46 \pm 0.18]$	$[58.78 \pm 21.89, 39.91 \pm 24.51, 0.00 \pm 0.00, -, 1.31 \pm 2.92]$
	NACHOS_NC	$[1.16 \pm 0.14, 1.86 \pm 0.18, 2.60 \pm 0.00, -, 4.74 \pm 0.08]$	$[22.65 \pm 2.61, 72.29 \pm 27.00, 49.58 \pm 0.00, -, 4.20 \pm 6.26]$
	NACHOS_NP	$[1.20 \pm 0.07, 2.03 \pm 0.07, 2.81 \pm 0.00, -, 4.62 \pm 0.18]$	$[78.29 \pm 12.81, 30.34 \pm 6.92, 0.00 \pm 0.00, -, 3.50 \pm 4.84]$
CIFAR-10	NACHOS	$[1.60 \pm 0.00, 2.18 \pm 0.31, 3.40 \pm 0.44, -, 5.78 \pm 0.19]$	$[41.31 \pm 30.37, 74.31 \pm 30.32, 0.00 \pm 0.00, -, 0.90 \pm 1.52]$
	NACHOS_NC	$[1.68 \pm 0.06, 2.57 \pm 0.38, 3.67 \pm 0.24, -, 6.43 \pm 0.12]$	$[44.54 \pm 22.45, 42.02 \pm 32.35, 17.07 \pm 2.23, -, 15.41 \pm 8.97]$
	NACHOS_NP	$[-, 2.01 \pm 0.27, 3.05 \pm 0.08, -, 6.25 \pm 0.42]$	$[55.2 \pm 1.91, 74.39 \pm 28.54, -, -, 3.53 \pm 3.27]$
CINIC-10	NACHOS	$[1.80 \pm 0.20, 2.45 \pm 0.50, 3.70 \pm 0.42, -, 5.87 \pm 0.10]$	$[69.75 \pm 26.03, 30.25 \pm 26.03, 0.00 \pm 0.00, -, -]$
	NACHOS_NC	$[1.50 \pm 0.00, 3.03 \pm 0.29, 3.77 \pm 0.12, -, 6.85 \pm 0.11]$	$[9.87 \pm 0.00, 39.70 \pm 27.63, 74.61 \pm 17.79, -, 6.58 \pm 11.93]$
	NACHOS_NP	$[1.25 \pm 0.07, 2.34 \pm 0.12, 2.95 \pm 0.21, -, 4.79 \pm 0.12]$	$[63.28 \pm 2.86, 60.85 \pm 44.39, 12.25 \pm 0.95, -, 8.94 \pm 3.21]$
Imagenette	NACHOS	$[-, 46.00 \pm 5.29, 59.81 \pm 3.63, 70.21 \pm 6.78, 104.09 \pm 1.09]$	$[4.67 \pm 3.33, 14.31 \pm 9.70, 26.02 \pm 14.98, -, 72.42 \pm 7.22]$
	NACHOS_NC	$[-, 52.35 \pm 0.00, 65.06 \pm 4.30, 92.06 \pm 3.08, 121.30 \pm 3.15]$	$[-, 0.0 \pm 0.0, 0.79 \pm 1.09, 8.77 \pm 9.54, 90.44 \pm 8.88]$
	NACHOS_NP	$[-, 51.00 \pm 2.95, 62.67 \pm 2.52, -, 98.88 \pm 0.41]$	$[-, 24.55 \pm 8.93, 24.16 \pm 12.65, -, 65.86 \pm 3.22]$

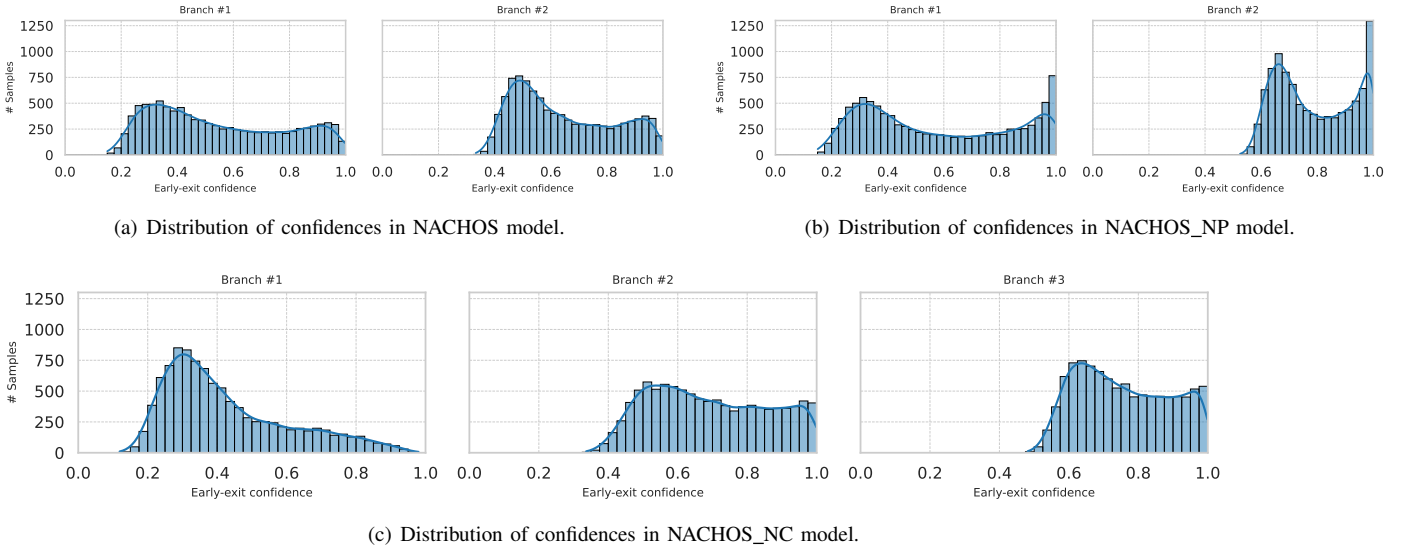


Fig. 5. Comparison between the distribution of confidences of NACHOS models trained by applying different regularization while setting the same seed.

els obtained with the same seed. We can point out that the  $L_{peak}$  is effective in aligning the confidences towards the operating point of the EECs. This is evident in the decrease of the number of samples for very high confidences for both EECs and the corresponding increase of samples towards the operating point in the NACHOS model with respect to the NACHOS\_NP model. Moreover, the  $L_{cost}$  is effective in favoring the training of the early fraction of the EENN. This is seen in the higher number of samples for the higher confidences in the NACHOS model in the first EEC with

respect to the NACHOS\_NC model.

### E. Statistical tests

To validate the statistical significance of our results, we performed the Friedman statistical test by using  $F_M(\bar{x})$  and  $F_A(\hat{x})$  of the multiple runs over five seeds of all our tested methods, i.e., NACHOS, NACHOS\_NP, NACHOS\_NC, and EDANAS, for each dataset and reported results in Table VII. The results shows that there are statistic differences among the methods, hence the competitive performance of NACHOS is statistically validated.

TABLE VII

FRIEDMAN TEST RESULTS FOR EACH DATASET AND OBJECTIVE. THE HIGHER  $\chi^2$ , THE MORE DIFFERENT THE RESULTS OF THE METHODS BASED ON THEIR RANKINGS AND THE SMALLER THE P-VALUE (TYPICALLY  $< 0.05$ ), THE MORE STATISTICALLY SIGNIFICANT THE  $\chi^2$  STATISTIC.

Dataset & Objective	Chi-squared ( $\chi^2$ )	p-value
SVHN ( $F_M(\tilde{x})$ )	9.9600	0.0189
SVHN ( $F_A(\tilde{x})$ )	15.0000	0.0018
CIFAR-10 ( $F_M(\tilde{x})$ )	14.0400	0.0029
CIFAR-10 ( $F_A(\tilde{x})$ )	13.5600	0.0036
CINIC-10 ( $F_M(\tilde{x})$ )	10.5625	0.0143
CINIC-10 ( $F_A(\tilde{x})$ )	13.0408	0.0045
Imagenette ( $F_M(\tilde{x})$ )	15.0000	0.0018
Imagenette ( $F_A(\tilde{x})$ )	15.0000	0.0018

## VI. CONCLUSIONS

In this work, we have presented Neural Architecture Search for Hardware Constrained Early Exit Neural Networks (NACHOS), the first NAS framework for the design of optimal EENNs satisfying a constraint on the number of Multiply and Accumulate (MAC) operations performed by the EENNs at inference time. The results show that NACHOS is able to select a set of Pareto Optimal Solutions competitive with the other state-of-the-art EENNs. Additionally, we investigated the effectiveness of the novel regularizations introduced in the training of the EENNs proposed in NACHOS. Future works will examine different OFAs, early-exit schemes, optimization techniques, application scenarios, and the combination of NACHOS with other AutoML techniques. In particular, we stress that one limitation of the current research is that it primarily focuses on mobile devices, optimizing neural networks with constraints in the range of millions of parameters and MACs. Future research directions could focus on further optimization of neural networks by incorporating hardware accelerator architectures and compiler mappings directly into the design space. This could lead to the development of models that are efficient enough to be deployed on MCUs and in general to scenarios with very severe constraints.

## VII. ACKNOWLEDGMENTS

This paper is supported by PNRR-PE-AI FAIR project funded by the NextGeneration EU program.

## REFERENCES

- [1] S. Teerapittayanon, B. McDanel, and H. T. Kung, "BranchyNet: Fast Inference via Early Exiting from Deep Neural Networks," Sep. 2017, arXiv:1709.01686 [cs]. [Online]. Available: <http://arxiv.org/abs/1709.01686>
- [2] S. Scardapane, M. Scarpiniti, E. Baccarelli, and A. Uncini, "Why Should We Add Early Exits to Neural Networks?" *Cognitive Computation*, vol. 12, no. 5, pp. 954–966, Sep. 2020. [Online]. Available: <https://link.springer.com/10.1007/s12559-020-09734-4>
- [3] X. Ying, "An Overview of Overfitting and its Solutions," *Journal of Physics: Conference Series*, vol. 1168, p. 022022, Feb. 2019. [Online]. Available: <https://iopscience.iop.org/article/10.1088/1742-6596/1168/2/022022>
- [4] X. Wang, Y. Luo, D. Crankshaw, A. Tumanov, F. Yu, and J. E. Gonzalez, "IDK Cascades: Fast Deep Learning by Learning not to Overthink," Jun. 2018, arXiv:1706.00885 [cs]. [Online]. Available: <http://arxiv.org/abs/1706.00885>
- [5] Y. Kaya, S. Hong, and T. Dumitras, "Shallow-Deep Networks: Understanding and Mitigating Network Overthinking," May 2019, arXiv:1810.07052 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1810.07052>
- [6] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training Recurrent Neural Networks," Feb. 2013, arXiv:1211.5063 [cs]. [Online]. Available: <http://arxiv.org/abs/1211.5063>
- [7] J. Frankle and M. Carbin, "The Lottery Ticket Hypothesis: Finding Sparse, Trainable Neural Networks," Mar. 2019, arXiv:1803.03635 [cs]. [Online]. Available: <http://arxiv.org/abs/1803.03635>
- [8] J. Deng, W. Dong, R. Socher, L.-J. Li, Kai Li, and Li Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*. Miami, FL: IEEE, Jun. 2009, pp. 248–255. [Online]. Available: <https://ieeexplore.ieee.org/document/5206848/>
- [9] S. Scardapane, D. Comminiello, M. Scarpiniti, E. Baccarelli, and A. Uncini, "Differentiable Branching In Deep Networks for Fast Inference," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Barcelona, Spain: IEEE, May 2020, pp. 4167–4171. [Online]. Available: <https://ieeexplore.ieee.org/document/9054209/>
- [10] C. Xu and J. McAuley, "A Survey on Dynamic Neural Networks for Natural Language Processing."
- [11] W. Zhu, P. Wang, Y. Ni, G. Xie, and X. Wang, "BADGE: Speeding Up BERT Inference after Deployment via Block-wise Bypasses and Divergence-based Early Exiting," in *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 5: Industry Track)*. Toronto, Canada: Association for Computational Linguistics, 2023, pp. 500–509. [Online]. Available: <https://aclanthology.org/2023.acl-industry.48>
- [12] E. Lattanzi, C. Contoli, and V. Freschi, "Do we need early exit networks in human activity recognition?" *Engineering Applications of Artificial Intelligence*, vol. 121, p. 106035, May 2023. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0952197623002191>
- [13] A. Sabet, J. Hare, B. Al-Hashimi, and G. V. Merrett, "Temporal Early Exits for Efficient Video Object Detection," Jun. 2021, arXiv:2106.11208 [cs]. [Online]. Available: <http://arxiv.org/abs/2106.11208>
- [14] Y. Rao, W. Zhao, B. Liu, J. Lu, J. Zhou, and C.-J. Hsieh, "DynamicViT: Efficient Vision Transformers with Dynamic Token Sparsification," Oct. 2021, arXiv:2106.02034 [cs]. [Online]. Available: <http://arxiv.org/abs/2106.02034>
- [15] E. Park, D. Kim, S. Kim, Y.-D. Kim, G. Kim, S. Yoon, and S. Yoo, "Big/little deep neural network for ultra low power inference," in *2015 International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS)*. Amsterdam, Netherlands: IEEE, Oct. 2015, pp. 124–132. [Online]. Available: <http://ieeexplore.ieee.org/document/7331375/>
- [16] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017, pp. 1321–1330.
- [17] P. Ren, Y. Xiao, X. Chang, P.-Y. Huang, Z. Li, X. Chen, and X. Wang, "A comprehensive survey of neural architecture search: Challenges and solutions," *arXiv preprint arXiv:2006.02903*, 2020.
- [18] D. Whitley, T. Starkweather, and C. Bogart, "Genetic algorithms and neural networks: optimizing connections and connectivity," *Parallel Computing*, vol. 14, no. 3, pp. 347–361, 1990. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0167819190900860>
- [19] V. K. Ojha, A. Abraham, and V. Snašal, "Metaheuristic design of feedforward neural networks: A review of two decades of research," *Engineering Applications of Artificial Intelligence*, vol. 60, pp. 97–116, 2017. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0952197617300234>
- [20] H. Benmezziane, K. El Maghraoui, H. Ouarnoughi, S. Niar, M. Wistuba, and N. Wang, "Hardware-Aware Neural Architecture Search: Survey and Taxonomy," in *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*. Montreal, Canada: International Joint Conferences on Artificial Intelligence Organization, Aug. 2021, pp. 4322–4329. [Online]. Available: <https://www.ijcai.org/proceedings/2021/592>
- [21] M. Gambella, A. Falcetta, and M. Roveri, "CNAS: Constrained Neural Architecture Search," in *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. Prague, Czech Republic:

- IEEE, Oct. 2022, pp. 2918–2923. [Online]. Available: <https://ieeexplore.ieee.org/document/9945080/>
- [22] T. King, Y. Zhou, T. Röddiger, and M. Beigl, “MicroNAS: Memory and Latency Constrained Hardware-Aware Neural Architecture Search for Time Series Classification on Microcontrollers,” Oct. 2023, arXiv:2310.18384 [cs]. [Online]. Available: <http://arxiv.org/abs/2310.18384>
- [23] H. Bouzidi, M. Odema, H. Ouarnoughi, M. A. Al Faruque, and S. Niar, “Hadas: Hardware-aware dynamic neural architecture search for edge performance scaling,” in *2023 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2023, pp. 1–6.
- [24] S. Li, Y. Sun, G. G. Yen, and M. Zhang, “Automatic Design of Convolutional Neural Network Architectures Under Resource Constraints,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 34, no. 8, pp. 3832–3846, Aug. 2023. [Online]. Available: <https://ieeexplore.ieee.org/document/9609007/>
- [25] M. Odema, N. Rashid, and M. A. A. Faruque, “EEExNAS: Early-Exit Neural Architecture Search Solutions for Low-Power Wearable Devices,” in *2021 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)*. Boston, MA, USA: IEEE, Jul. 2021, pp. 1–6. [Online]. Available: <https://ieeexplore.ieee.org/document/9502503/>
- [26] B. Zeinali, D. Zhuang, and J. M. Chang, “Esai: Efficient split artificial intelligence via early exiting using neural architecture search,” *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 9, no. 1, pp. 961–971, 2025.
- [27] M. Gambella and M. Roveri, “Edanas: Adaptive neural architecture search for early exit neural networks,” in *2023 International Joint Conference on Neural Networks (IJCNN)*, 2023, pp. 1–8.
- [28] H. Cai, C. Gan, T. Wang, Z. Zhang, and S. Han, “Once-for-All: Train One Network and Specialize it for Efficient Deployment,” Apr. 2020, arXiv:1908.09791 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1908.09791>
- [29] S. Sarti, E. Lomurno, A. Falanti, and M. Matteucci, “Enhancing once-for-all: A study on parallel blocks, skip connections and early exits,” in *2023 International Joint Conference on Neural Networks (IJCNN)*, 2023, pp. 1–9.
- [30] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, “A fast and elitist multiobjective genetic algorithm: NSGA-II,” *IEEE Transactions on Evolutionary Computation*, vol. 6, no. 2, pp. 182–197, Apr. 2002. [Online]. Available: <http://ieeexplore.ieee.org/document/996017/>
- [31] S. Jeon, Y. Choi, Y. Cho, and H. Cha, “Harvnet: Resource-optimized operation of multi-exit deep neural networks on energy harvesting devices,” in *Proceedings of the 21st Annual International Conference on Mobile Systems, Applications and Services*, ser. MobiSys ’23. New York, NY, USA: Association for Computing Machinery, 2023, p. 42–55. [Online]. Available: <https://doi.org/10.1145/3581791.3596845>
- [32] E. S. Marquez, J. S. Hare, and M. Niranjana, “Deep Cascade Learning,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 29, no. 11, pp. 5475–5485, Nov. 2018. [Online]. Available: <https://ieeexplore.ieee.org/document/8307262/>
- [33] S. Abbasi, M. Hajabdollahi, N. Karimi, and S. Samavi, “Modeling teacher-student techniques in deep neural networks for knowledge distillation,” *CoRR*, vol. abs/1912.13179, 2019. [Online]. Available: <http://arxiv.org/abs/1912.13179>
- [34] D. Stamoulis, T.-W. Chin, A. K. Prakash, H. Fang, S. Sajja, M. Bognar, and D. Marculescu, “Designing Adaptive Neural Networks for Energy-Constrained Image Classification,” Aug. 2018, arXiv:1808.01550 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1808.01550>
- [35] J. Pomponi, S. Scardapane, and A. Uncini, “A probabilistic re-interpretation of confidence scores in multi-exit models,” *Entropy*, vol. 24, no. 1, 2022. [Online]. Available: <https://www.mdpi.com/1099-4300/24/1/1>
- [36] M. Wistuba, A. Rawat, and T. Pedapati, “A Survey on Neural Architecture Search,” Jun. 2019, arXiv:1905.01392 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1905.01392>
- [37] A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, Q. V. Le, and H. Adam, “Searching for MobileNetV3,” Nov. 2019, arXiv:1905.02244 [cs]. [Online]. Available: <http://arxiv.org/abs/1905.02244>
- [38] R. Sutton and A. Barto, “Reinforcement learning: An introduction,” *IEEE Transactions on Neural Networks*, vol. 9, no. 5, pp. 1054–1054, 1998.
- [39] L. P. Kaelbling, M. L. Littman, and A. W. Moore, “Reinforcement learning: A survey,” 1996.
- [40] H. Liu, K. Simonyan, and Y. Yang, “DARTS: Differentiable Architecture Search,” Apr. 2019, arXiv:1806.09055 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1806.09055>
- [41] T. Back, D. B. Fogel, and Z. Michalewicz, *Handbook of Evolutionary Computation*, 1st ed. GBR: IOP Publishing Ltd., 1997.
- [42] D. Ashlock, *Evolutionary Computation for Modeling and Optimization*, 1st ed. Springer Publishing Company, Incorporated, 2010.
- [43] Y. Liu, Y. Sun, B. Xue, M. Zhang, G. G. Yen, and K. C. Tan, “A survey on evolutionary neural architecture search,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 34, no. 2, p. 550–570, Feb. 2023. [Online]. Available: <http://dx.doi.org/10.1109/TNNLS.2021.3100554>
- [44] M. S. Mahbub, “A comparative study on constraint handling techniques of nsgaii,” in *2020 International Conference on Electrical, Communication, and Computer Engineering (ICECCE)*, 2020, pp. 1–5.
- [45] C. A. Coello-Coello, “Theoretical and numerical constraint-handling techniques used with evolutionary algorithms: a survey of the state of the art,” *Computer Methods in Applied Mechanics and Engineering*, 2002. [Online]. Available: <https://api.semanticscholar.org/CorpusID:9235579>
- [46] C. Gong, D. Wang, M. Li, X. Chen, Z. Yan, Y. Tian, Q. Liu, and V. Chandra, “NASViT: Neural architecture search for efficient vision transformers with gradient conflict aware supernet training,” in *International Conference on Learning Representations (ICLR)*, 2022. [Online]. Available: <https://openreview.net/forum?id=Qaw16njk6L>
- [47] X. Dai, P. Zhang, B. Wu, H. Yin, F. Sun, Y. Wang, M. Dukhan, Y. Hu, Y. Wu, Y. Jia, P. Vajda, M. Uyttendaele, and N. K. Jha, “ChamNet: Towards Efficient Network Design through Platform-Aware Model Adaptation,” Dec. 2018, arXiv:1812.08934 [cs]. [Online]. Available: <http://arxiv.org/abs/1812.08934>
- [48] B. Baker, O. Gupta, R. Raskar, and N. Naik, “Accelerating neural architecture search using performance prediction,” *arXiv preprint arXiv:1705.10823*, 2017.
- [49] Z. Lu, K. Deb, E. Goodman, W. Banzhaf, and V. N. Boddeti, “NSGANetV2: Evolutionary Multi-Objective Surrogate-Assisted Neural Architecture Search,” Jul. 2020, arXiv:2007.10396 [cs]. [Online]. Available: <http://arxiv.org/abs/2007.10396>
- [50] M. G. Kendall, “A new measure of rank correlation,” *Biometrika*, vol. 30, no. 1/2, pp. 81–93, 1938. [Online]. Available: <http://www.jstor.org/stable/2332226>
- [51] T. Wei, C. Wang, Y. Rui, and C. W. Chen, “Network morphism,” 2016.
- [52] E. Baccarelli, S. Scardapane, M. Scarpiniti, A. Momenzadeh, and A. Uncini, “Optimized training and scalable implementation of Conditional Deep Neural Networks with early exits for Fog-supported IoT applications,” *Information Sciences*, vol. 521, pp. 107–143, Jun. 2020. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0020025520301249>
- [53] A. Krizhevsky, “Learning multiple layers of features from tiny images,” 2009.
- [54] L. N. Darlow, E. J. Crowley, A. Antoniou, and A. J. Storkey, “CINIC-10 is not ImageNet or CIFAR-10,” Oct. 2018, arXiv:1810.03505 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1810.03505>
- [55] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng, “Reading Digits in Natural Images with Unsupervised Feature Learning,” 2011.
- [56] J. Howard and S. Gugger, “fastai: A layered api for deep learning,” *Information*, vol. 11, no. 2, p. 108, 2020.
- [57] M. Lindauer and F. Hutter, “Best practices for scientific research on neural architecture search,” 2020. [Online]. Available: <https://arxiv.org/abs/1909.02453>
- [58] S. Disabato and M. Roveri, “Reducing the Computation Load of Convolutional Neural Networks through Gate Classification,” in *2018 International Joint Conference on Neural Networks (IJCNN)*. Rio de Janeiro: IEEE, Jul. 2018, pp. 1–8. [Online]. Available: <https://ieeexplore.ieee.org/document/8489276/>
- [59] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” *CoRR*, vol. abs/1512.03385, 2015. [Online]. Available: <http://arxiv.org/abs/1512.03385>
- [60] M. Tan and Q. V. Le, “Efficientnet: Rethinking model scaling for convolutional neural networks,” 2020. [Online]. Available: <https://arxiv.org/abs/1905.11946>

- [61] M. P. Naeini, G. Cooper, and M. Hauskrecht, "Obtaining well calibrated probabilities using bayesian binning," in *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- [62] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On Calibration of Modern Neural Networks," Aug. 2017, arXiv:1706.04599 [cs]. [Online]. Available: <http://arxiv.org/abs/1706.04599>