

Searching Neural Architectures for Sensor Nodes on IoT Gateways

Andrea Mattia Garavagno, *Graduate Student Member, IEEE*, Edoardo Ragusa, *Member, IEEE*, Antonio Frisoli, *Member, IEEE*, Paolo Gastaldo

Abstract—This paper presents an automatic method for the design of Neural Networks (NNs) at the edge, enabling Machine Learning (ML) access even in privacy-sensitive Internet of Things (IoT) applications. The proposed method runs on IoT gateways and designs NNs for connected sensor nodes without sharing the collected data outside the local network, keeping the data in the site of collection. This approach has the potential to enable ML for Healthcare Internet of Things (HIoT) and Industrial Internet of Things (IIoT), designing hardware-friendly and custom NNs at the edge for personalized healthcare and advanced industrial services such as quality control, predictive maintenance, or fault diagnosis. By preventing data from being disclosed to cloud services, this method safeguards sensitive information, including industrial secrets and personal data. The outcomes of a thorough experimental session confirm that –on the Visual Wake Words dataset– the proposed approach can achieve state-of-the-art results by exploiting a search procedure that runs in less than 10 hours on the Raspberry Pi Zero 2.

Index Terms—Neural Architecture Search, Edge AI, Healthcare Internet of Things, Industrial Internet of Things.

I. INTRODUCTION

Neural Networks (NNs) are widely used in Internet of Things (IoT) applications [1]. In this context, often the data collected by the available sensors are added to the training set with the purpose of improving generalization performances. On the other hand, in some cases, the data can be sensitive; healthcare data [2], industrial data [3] and biometric data [4] provide possible examples. Privacy concerns prevent some entities from accessing the benefits of machine learning (ML), as they may be unable or unwilling to share their data with cloud services that can train or even automatically design a custom neural network (NN) [5].

To overcome this problem and make NNs available even in privacy-sensitive applications, we provide a method to automatically design NNs for sensor nodes on IoT gateways using locally collected data. This in turn means that the whole process can be completed without any access to cloud services, therefore maintaining data privacy. In the case of

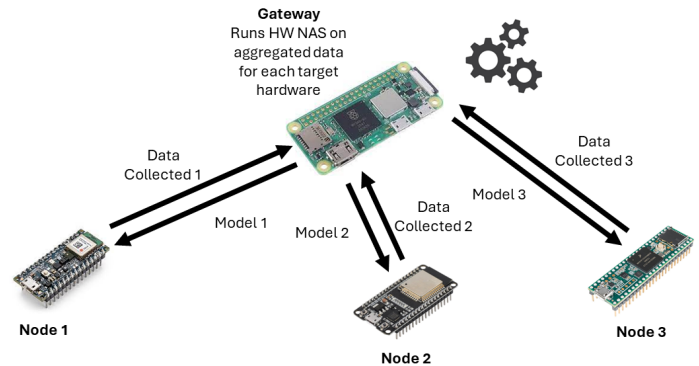


Fig. 1: Automatic designing of neural architectures at the edge, running HW-NAS on an IoT gateway, using the locally collected data.

Healthcare Internet of Things (HIoT), this approach can be used to provide personalized healthcare services with customized and hardware-friendly NNs for each patient [6], [7]. Industrial Internet of Things (IIoT) can also benefit from this approach. For example, one can provide hardware-friendly and customized NNs for intelligent fault diagnosis (IFD) of production machinery or quality control, without facing the risk of data leakage that could reveal industrial secrets [8].

A synopsis of the proposed scenario is provided in Fig. 1. First, the data collected by the sensor nodes are aggregated by the local gateway; then, the gateway adopts the proposed method to design a custom NN for each connected sensor node, as, in general, each node may differ from the other in terms of hardware resources. Finally, the designed NNs are deployed on the sensor nodes. The whole process can be completed without transferring data to any cloud infrastructure, thus preserving data privacy.

The proposed method exploits Hardware-Aware Neural Architecture Search (HW-NAS) to automatically design NNs for sensor nodes; such an approach is becoming popular in IoT applications [9]–[11]. Still, this work addresses a critical issue: the computational cost of the search procedure. The general goal of state-of-the-art HW-NAS is to generate NNs that fit the resources available on the hardware that will run the inference process. However, the resources available on the platform that runs the search procedure itself are usually considered unconstrained [12]. The latter assumption becomes critical when working with constrained devices such as IoT gateways [13], [14]. If one wants to avoid any data transfer

Manuscript received Month Day, Year; revised Month Day, Year; (Corresponding author: Andrea Mattia Garavagno.)

Andrea Mattia Garavagno, Edoardo Ragusa, and Paolo Gastaldo are with the Department of Electrical, Electronic, Telecommunications Engineering and Naval Architecture, University of Genoa, 16126 Genoa, Italy (e-mail: AndreaMattia.Garavagno@edu.unige.it, Edoardo.Ragusa@unige.it, Paolo.Gastaldo@unige.it).

Andrea Mattia Garavagno and Antonio Frisoli are with the Institute of Mechanical Intelligence, Scuola Superiore Sant’Anna, Pisa 56127, Italy (e-mail: AndreaMattia.Garavagno@santannapisa.it, Antonio.Frisoli@santannapisa.it).

to computing resources that are external to the IoT network, then the search procedure should be designed to run on the available IoT gateway.

The key challenges that arise when running HW-NAS on edge platforms are a) dealing with the low computational capabilities and b) dealing with the low memory availability of these devices. To address these challenges, we adopt a derivative-free search strategy combined with a regular search space. The regular search space eases the evaluation process by employing architectures that are hardware-friendly in the sense that they do not feature too many branches to be followed during the backpropagation phase, saving memory and computations. The derivative-free search strategy avoids the usage of derivatives and therefore requires less memory and less computations compared to derivative-based algorithms.

The present paper inherits and expands the work recently presented in [15], further developing the idea of running HW-NAS on constrained devices. The main contributions of this paper can be summarized as follows.

- A novel search strategy that obtains state-of-the-art results on the Visual Wake Words dataset while reducing search time from 4 days to 10 hours when running on the Raspberry Pi Zero 2.
- An adaptation mechanism which tunes the search space according to the energy and time available to the IoT gateway, enabling the execution of HW-NAS even in the case of a limited energy budget or with a limited amount of execution time at disposal.
- The added capability of working with time series, obtaining state-of-the-art results on the CWRU dataset, a dataset for Intelligent Fault Diagnosis, in just 2 hours and 52 minutes when running on a Raspberry Pi 4.
- The release of the proposed HW-NAS as open-source software at <https://github.com/AndreaMattiaGaravagno/GatewayNAS>; the software is designed for embedded devices based on Linux.

The experimental sessions involved three different single-board computers typically used for implementing IoT gateways and three different ultra-low-power microcontroller units (MCUs) as reliable instances of sensing nodes. The experiments showed that the proposed HW-NAS could successfully meet the constraints on execution time and energy consumption while generating architectures that scored state-of-the-art performances on the selected benchmarks.

The rest of the paper is organized as follows. Section II reviews related works. Section III briefly introduces the reader to the basics of HW-NAS and then focuses on the specific problem of running HW-NAS on resource-constrained devices. Section IV presents the proposed approach to the design of an HW-NAS that fits the requirements of the IoT environment. Section V reports on the experimental setup, while Section VI presents the outcomes of the experiments. Finally, concluding remarks are provided in Section VII.

II. RELATED WORKS

HW-NAS represents a valuable solution for automatically designing deep learning models whose inference phase should

run on IoT sensor nodes [16]–[18]. Wang et al. [11] applied HW-NAS to the design of NN that must be robust against adversarial attacks in 6G-enabled AI-enabled Internet-of-Things (AIoT) systems. In [19] an HW-NAS generates and explores multiple NN architectures for biosignal processing in wearable devices. Du et al. [20] exploited HW-NAS to design tiny deep learning models to analyze the emotional information of steady-state visual evoked potential signals in the IoT framework. Dissem et al. [21] used HW-NAS to design optimal autoencoders to perform anomaly detection in smart buildings. In [22], HW-NAS supports the development of IoT systems for vibration damage detection.

TABLE I: Target execution platform of state-of-the-art HW-NAS.

Work	GPU	CPU	Embedded	
			Static	Adaptive
Mnasnet [23]	x			
MCUNet [24]	x			
Micronets [25]	x			
ColabNAS [26]	x			
NanoNAS v1 [27]		x		
NanoNAS v2 [28]			x	
Proposal				x

Originally, HW-NAS was a resource-hungry procedure; as a result, one needed expensive computing facilities to run it. In the last years, though, efforts have been made to decrease the computational cost of HW-NAS procedures. From the 40,000 GPU hours required by Mnasnet [23], the execution time was reduced to 300 GPU hours by MCUNet [24]. In addition, custom techniques leveraging derivative-free optimization and hand-crafted search spaces have been introduced to meet the resource constraints of ultra-low-power hardware: in [26] the search cost was 4 GPU hours, while Garavagno et al. eventually removed the need for a GPU [27]. Indeed, two recent works showed that it can be feasible to run the HW-NAS itself on an embedded system [15], [28]. The main idea is to add a constraint on the memory available on the platform running the search procedure. Table I offers an overview of the state-of-the-art HW-NAS; the works are classified according to the computing resources needed to perform the search procedure.

The literature provides other approaches that can be adopted to adapt the architecture of a DNN to changes in the deployment environment. Dynamic Neural Networks enjoy favorable properties that are absent in common neural networks [29]: they can adapt their topology and parameters in response to varying input during inference; they can learn a set of diverse features, enhancing their generalization capabilities; they support interpretable ML, giving the ability to understand which part of the network has generated the output. However, the known gap between theoretical and practical efficiency makes them unsuitable for modern computing devices [30].

Domain adaptation [31] transfers knowledge from a source domain to a target domain having a different data distribution. It has been successfully applied to increase the generalization capability of NNs for IoT applications trained in a situation of data scarcity [32]–[35]. Nonetheless, the burden of designing

a neural architecture compatible with the hardware constraint of the sensor node remains on the user.

Hence, the present paper wants to go one step further in the design of HW-NAS that can run at the edge of IoT systems, without relying on external resources. The proposed adaptive HW-NAS can automatically adjust the search procedure according to the constraints on execution time and energy budget. This in turn means that it can run on generic gateways in an IoT network.

III. BACKGROUND

A. Hardware-Aware Neural Architecture Search

HW-NAS automatizes the process of finding the best neural architecture for solving a given task on a target hardware. HW-NAS is typically cast as an **optimization problem** constrained by the resources available on the target hardware. Its objective function defines the set of metrics adopted for selecting the best neural architecture among candidate solutions, i.e., the **evaluation process** [36]. Equation (1) shows an example of the typical optimization problem, where one wants to find the architecture A having the maximum validation accuracy, $\max \Gamma(A)$, among those meeting a set of constraints. In this example the constraints are the amount of RAM (ξ_{RAM}), the amount of Flash memory (ξ_{Flash}), and the multiply and accumulate (MAC) instructions (ξ_{MAC}); accordingly, $\phi_{RAM}(A)$, $\phi_{Flash}(A)$, and $\phi_{MAC}(A)$ refer, respectively, to the RAM usage, the Flash memory usage, and the MAC instructions of the candidate architecture A on the deployment device.

$$\left\{ \begin{array}{l} \max \Gamma(A) \\ \phi_{RAM}(A) \leq \xi_{RAM}, \phi_{Flash}(A) \leq \xi_{Flash} \\ \phi_{MAC}(A) \leq \xi_{MAC} \\ \xi_{RAM}, \xi_{Flash}, \xi_{MAC} > 0 \end{array} \right. \quad (1)$$

The set of candidate solutions, also known as **search space**, is defined by fixed architectural rules determining which NNs are allowed. Rules can be defined with different granularity: they can act on layers (layer-wise search space), on a group of layers, often referred as cells (cell-based search space), or both (hierarchical search space) [36].

The optimization technique adopted for solving the problem, i.e., the **search strategy**, defines how the search space is explored when looking for the best neural architecture [36]. In the first place, techniques based on Reinforcement Learning (RL) [23], Evolutionary Aging (EA) [37], Random Search [38], and Bayesian Optimization (BO) [39] were adopted given the non-smoothness of the HW-NAS problem. Then, a novel formulation of HW-NAS [40] allowed the usage of gradient-based techniques [41], [42] for solving the optimization problem. Nonetheless, the majority of search strategies are known to be resource intensive [36] and can require up to tens of thousands hours of execution on specialized hardware [24]. In the recent years, though, new approaches have been introduced that lower the search cost of HW-NAS by adopting custom search strategies [15], [27], [43].

B. Running HW-NAS on Resource Constrained Devices

Running HW-NAS on embedded devices poses major challenges, as limited resources are available. Two recent works [15] [28] have shown that the search cost can be significantly reduced by exploiting a custom derivative-free optimization technique and a cell-based search space; in both cases the HW-NAS was designed to generate NNs that should run on ultra-low-power computing platform.

The adopted search space [15] builds tiny Convolutional Neural Networks (CNNs) stacking four different types of cells: a pre-processing cell, a base cell, a building cell, and a classifier cell. The pre-processing cell applies min-max standardization to improve gradient descent convergence rate [44]. The base cell is a single convolutional layer with k kernels acting as the foundation on which building cells will be stacked. A single building cell first halves the resolution of the extracted features using the max pooling operator, then applies convolution followed by batch normalization to stabilize further and speed up the training [45]; then a rectified linear unit activation builds the new feature maps. The number of kernels n_c of a convolutional layer depends on k and on the number of cells previously stacked c , and can be computed as $n_c = n_{c-1} + 2^{1-c}n_{c-1}$, where $n_0 = k$ [15]. The latter formula is inspired by VGG16 [46], where instead of doubling the kernels after each block, we gradually decrease the increment to save resources.

Finally, the classifier cell reduces the extracted features by applying global average pooling to improve the model's generalization capability [47]; a fully connected layer produces the class predictions, using softmax activation and a number of neurons equal to the number of classes. All the kernels use a 3 by 3 receptive field.

Therefore, in the adopted search space, neural architectures are built stacking building cells upon the base cell. Every architecture starts with a pre-processing cell and ends with a classifier cell and can be conveniently represented by the tuple (k, c) , where k is the number of kernels used in the base cell and c is the number of cells stacked upon the base cell.

To explicitly take into account the limitations imposed by the device on which the HW-NAS should run, a custom constraint was added to the optimization problem in [15] [28]. Thus, a limit was added on the available memory for the training phase on the device hosting the search procedure. Accordingly, candidate architectures not fitting such constraint were considered unfeasible, even if they could be deployed on the target hardware.

IV. PROPOSAL

The core feature of the proposed HW-NAS is the capability to be "hardware aware" both at the edge (i.e., the device on which the NN should run) and at the gateway (i.e., the device on which the NAS should run). To this purpose, the HW-NAS involves six different constraints:

- at the edge: a candidate solution (i.e., an architecture $A(k, c)$) must meet the constraints on the amount of RAM (ξ_{RAM}), the amount of Flash memory (ξ_{Flash}), and the MAC instructions (ξ_{MAC}).

- At the gateway: the whole process to be completed to identify the best architecture must meet the constraints on the amount of memory (ξ_{MEM}), the execution time (ξ_{Time}), and the energy budget (ξ_{Energy}).

In practice, one wants to keep the search space as wide as possible (i.e., the number of candidate solutions), given the constraints at the gateway level. The present approach generates the set S_α of admissible solutions with a two-step procedure. First, one selects the architectures $A(k, c)$ that 1) can be deployed on the edge device and 2) can be trained on the gateway without violating the constraint on the memory usage ($\phi_{RAM}(A) \leq \xi_{RAM}$):

$$\begin{cases} \forall A \in S_\alpha \\ \phi_{RAM}(A) \leq \xi_{RAM}, \phi_{Flash}(A) \leq \xi_{Flash} \\ \phi_{MAC}(A) \leq \xi_{MAC}, \phi_{MEM}(A) \leq \xi_{MEM} \\ \xi_{RAM}, \xi_{Flash}, \xi_{MAC}, \xi_{MEM} > 0 \end{cases} \quad (2)$$

Then, a cropping process reduces the cardinality of S_α , if necessary. In fact, the search process must also meet the constraints on execution time and energy budget, which are evaluated by the functions ϕ_{Time} and ϕ_{Energy} , based on the candidate search space. Hence, the number of candidate solutions in S_α should be as large as possible, but, at the same time, it should be small enough to allow the gateway to finalize the search process without violating the constraints ξ_{Time} and ξ_{Energy} . Thus, the final set S_α is obtained as follows:

$$\begin{cases} \max |S_\alpha| \\ \phi_{Time}(S_\alpha) \leq \xi_{Time}, \phi_{Energy}(S_\alpha) \leq \xi_{Energy} \\ \xi_{Time}, \xi_{Energy} > 0 \end{cases} \quad (3)$$

Once the set S_α has been obtained, the search strategy can be applied to it. In the following, Sec. IV-A gives details about the algorithms that are adopted to generate S_α , while Section IV-B presents the search strategy.

A. Generating the Search Space On The Gateway

The actual search space S_α is obtained after a two-step procedure. In the first step, Algorithm 1 builds the set \hat{S}_α , i.e., a superset of S_α , in that it should include all admissible architectures $A(k, c)$ given the deployment target (as per (2)). It starts from the smallest admissible candidate, i.e., ($k = 1, c = 0$), and explores sequentially all the candidate architectures ($1, c$), adding them to \hat{S}_α until the first unfeasible architecture is reached (while loop at line 6 of the algorithm). Here, an architecture $A(k, c)$ is unfeasible when it does not meet the constraints set by (2). Then, k is incremented (repeat loop at line 4 of the algorithm) and again the exploration along the c parameter is completed. The algorithm continues accordingly until it reaches an architecture $A(k, 0)$ that is unfeasible (stopping criterion at line 11 of the algorithm).

Once \hat{S}_α has been computed, a cropping process can start. The issue to be tackled is that the gateway must meet the constraints on the execution time and the energy budget when running the NAS (as per (3)). Roughly, both quantities depend on the number of architectures included in the search space,

Algorithm 1 The algorithm that generates the extensive search space \hat{S}_α

Require: $\xi_{MEM}, \xi_{RAM}, \xi_{Flash}, \xi_{MAC}$

Ensure: \hat{S}_α

```

1: procedure SEARCH SPACE - CREATION
2:    $k \leftarrow 1$ 
3:    $\hat{S}_\alpha \leftarrow \emptyset$ 
4:   repeat ▷ Outer loop
5:      $c \leftarrow 0$ 
6:     while  $A(k, c)$  is feasible do ▷ Inner loop
7:        $\hat{S}_\alpha \leftarrow \hat{S}_\alpha \cup (k, c)$ 
8:        $c \leftarrow c + 1$ 
9:     end while
10:     $k \leftarrow k + 1$ 
11:  until  $(k, 0)$  is not feasible ▷ Stopping criterion
12:  return  $\hat{S}_\alpha$ 
13: end procedure

```

Algorithm 2 The algorithm that obtains the final search space S_α by cropping \hat{S}_α .

Require: $\hat{S}_\alpha, \xi_{Time}, \xi_{Energy}, \bar{w}$

Ensure: S_α

```

1: procedure SEARCH SPACE - CROPPING
2:    $A^* \leftarrow (A \in \hat{S}_\alpha | max\_params)$ 
3:    $\bar{t} \leftarrow eval\_time(A^*)$  ▷ Time upper bound
4:    $\bar{e} \leftarrow \bar{t}\bar{w}$  ▷ Energy upper bound
5:    $k \leftarrow 1$ 
6:    $S_\alpha \leftarrow \emptyset$ 
7:    $stop \leftarrow 0$ 
8:   repeat
9:      $c \leftarrow 0$ 
10:    while  $(k, c) \in \hat{S}_\alpha$  do
11:      if  $|S_\alpha| \bar{t} \leq \xi_{Time}, |S_\alpha| \bar{e} \leq \xi_{Energy}$  then
12:         $S_\alpha \leftarrow S_\alpha \cup A(k, c)$ 
13:      else
14:         $stop \leftarrow 1$ 
15:      end if
16:       $c \leftarrow c + 1$ 
17:    end while
18:     $k \leftarrow k + 1$ 
19:  until  $stop = 0$ 
20:  return  $S_\alpha$ 
21: end procedure

```

as the performance of each architecture should be assessed to find the best CNN among the candidate solutions. This in turn means that each architecture should be trained. Therefore, one cannot assume that the search space can be as large as \hat{S}_α . At the start of the cropping process S_α coincides with \hat{S}_α . Then, the cropping strategy formalized in the algorithm 2 basically reduces the cardinality of S_α to the purpose of obtaining a search space that allows the gateway to fit the constraints.

First, an upper bound \bar{t} is set for the time needed to evaluate a candidate architecture. To estimate such upper bound the algorithm selects the architecture with the highest number of parameters to be updated during training among those

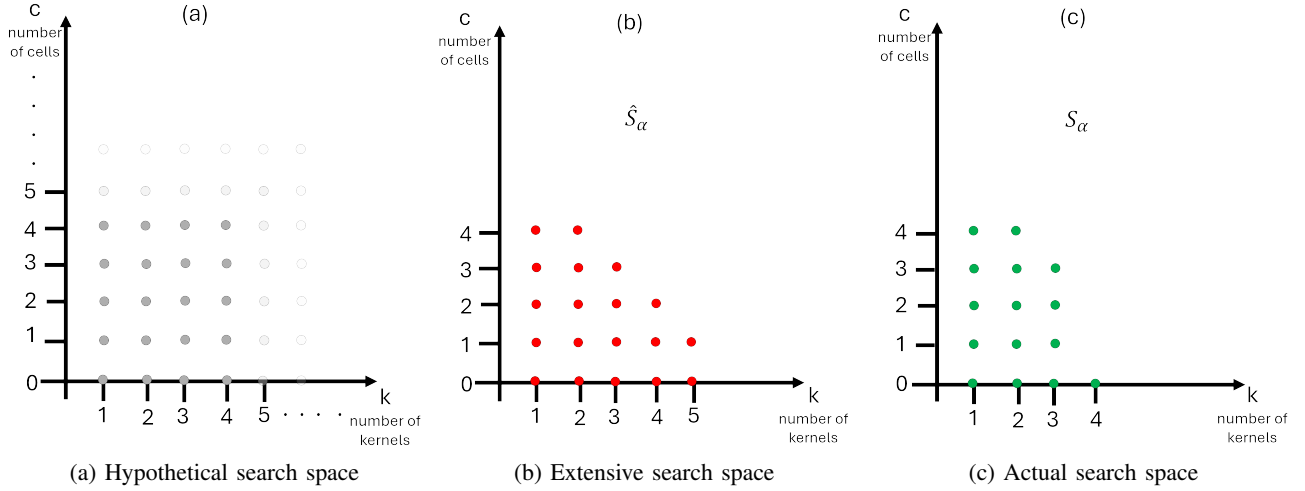


Fig. 2: Example of the process leading to the search space S_α

included in \hat{S}_α . Accordingly, the value of \bar{t} is set empirically by completing the evaluation process of such architecture (i.e., training and computation of the validation accuracy) on the gateway. The upper bound on the energy consumption, \bar{e} , is obtained by multiplying \bar{t} by the maximum power consumption of the gateway when running the HW-NAS, \bar{w} . The latter quantity is an input to the algorithm.

Then, starting from the smallest solution included in \hat{S}_α , ($k = 1, c = 0$), the algorithm adds one architecture at a time to S_α following the order in which the architectures have been added to \hat{S}_α . The process ends as soon as the cardinality of S_α reaches the threshold, i.e., when it is estimated that the gateway would not be able to complete the NAS by meeting the constraints on ξ_{Time} and ξ_{Energy} .

Figure 2 provides an example of the process leading to S_α . The starting point is the hypothetical search space (Fig. 2a), with all the possible architectures $A(k, c)$. The algorithm 1 generates an actual, yet temporary, search space \hat{S}_α (Fig. 2b). Finally, the algorithm 2 yields as output S_α by cropping \hat{S}_α (Fig. 2c).

B. Search Strategy

The proposed search strategy is a bi-level optimization technique that does not use derivatives. It searches for the best neural architecture within the search space described in section III-B to solve the given classification task on the chosen microcontroller, with a preference towards smaller solutions. The inner level of the technique (Alg. 3, while loop in row 7) searches the best number of building cells to stake given the number of kernels used in the base cell, acting on the parameter c , while the outer level (Fig. 3, repeat until loop in row 5) changes the parameter k , to find the best combination of k and c .

In detail, the inner level finds c^* , i.e., the best value of c given a fixed \bar{k} by adding building cells, one at a time, until the solution remains feasible. Instead, the outer level proposes a new candidate \bar{k} , based on the previously confirmed one k , obtained by adding to it a variable increment $\lfloor k2^{-\beta} \rfloor$, with beta initialized to zero, and calls the inner level. If the new

Algorithm 3 Proposed search strategy.

```

1: procedure PROPOSED OPTIMIZATION TECHNIQUE
2:    $k \leftarrow 1, c \leftarrow 0$  ▷ Starting point
3:    $\beta \leftarrow 0, \gamma \leftarrow 0$ 
4:    $\bar{k} \leftarrow k$ 
5:   repeat ▷ Outer loop
6:      $\bar{c} \leftarrow 0, c^* \leftarrow 0$ 
7:     while  $(\bar{k}, \bar{c} + 1) \in S_\alpha$  do ▷ Inner loop
8:       if  $f(\bar{k}, \bar{c} + 1) > f(\bar{k}, c^*)$  then
9:          $c^* \leftarrow \bar{c} + 1$  ▷ Update of best c
10:      end if
11:       $\bar{c} \leftarrow \bar{c} + 1$ 
12:    end while
13:    if  $f(\bar{k}, c^*) > f(k, c)$  then
14:       $k \leftarrow \bar{k}, c \leftarrow c^*$  ▷ Candidate confirmation
15:    else
16:       $\gamma \leftarrow 1$  ▷ Start k increment reduction
17:    end if
18:     $\beta \leftarrow \beta + \gamma$  ▷ Optional k increment reduction
19:     $\bar{k} \leftarrow k + \lfloor k2^{-\beta} \rfloor$  ▷ Variable increment of k
20:  until  $\lfloor k2^{-\beta} \rfloor = 0$  ▷ Stopping criterion
21: end procedure

```

candidate solution found is better than the previously confirmed one, the new candidate is confirmed, and the algorithm naturally proceeds. Otherwise, the beta starts to grow by one unit, decreasing the increment of the variable, and then a new \bar{k} is proposed. Eventually, the increment becomes zero, and the new proposed value coincides with the last confirmed one. When the latter condition is met, the algorithm stops, returning the last confirmed candidate. The evaluation metric adopted is the generalization capability, i.e. the validation accuracy, computed over a validation set.

Figure 3 proposes a possible run of the search strategy plotted on the search plane. In this figure, every parameter change that results in a better candidate is marked with a green arrow, whereas changes that lead to a worse candidate are marked with red arrows.

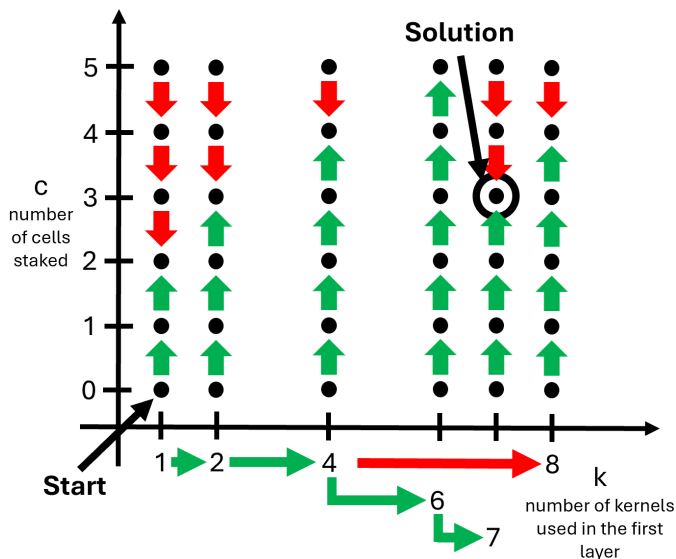


Fig. 3: A possible run of the search strategy plotted on the search plane.

The search starts from the smallest admissible solution, i.e. $A_0 = A(k \leftarrow 1, c \leftarrow 0)$ and explores the corresponding c -axis, finding the first candidate $A_1 = A(\bar{k} \leftarrow 1, c^* \leftarrow 2)$. Since A_1 performs better than A_0 , A_1 is confirmed, and the search proceeds by exploring the k -axis, doubling the confirmed k value, obtaining $\bar{k} = 2$. Then the c -axis is again explored finding the second candidate $A_2 = A(\bar{k} \leftarrow 2, c^* \leftarrow 3)$. The new candidate is confirmed since it performs better than the previous one, i.e. A_1 , and for this reason the change on the k -axis is marked with a green horizontal arrow.

The algorithm continues in the same way, until the performance on the k -axis stops improving, in this case with $\bar{k} = 8$, which is indeed pointed by a red arrow. The correspondent candidate $A_4 = A(\bar{k} \leftarrow 8, c^* \leftarrow 4)$ is therefore not confirmed leaving $A_3 = A(k \leftarrow 4, c \leftarrow 4)$ in the position of the confirmed candidate.

The stop of performance improvement triggers the k increment reduction, setting $\gamma \leftarrow 1$. Therefore a new k value $\bar{k} \leftarrow 6$ is proposed. The corresponding candidate found by exploring the c -axis, i.e., $A_5 = A(\bar{k} \leftarrow 6, c^* \leftarrow 5)$ performs better than A_3 , the previous confirmed candidate, and therefore it is confirmed being marked with a horizontal green arrow pointing to it.

The next candidate found $A_6 = A(k \leftarrow 7, c \leftarrow 3)$ improves the latest confirmed solution and therefore is marked with a horizontal green arrow. The last k value proposed is 7, which is the same as the last confirmed solution because the variable increment reached the zero value. Therefore the algorithm terminates returning A_6 .

V. EXPERIMENTAL SETUP

This paper presents the outcomes of four different experiments:

- The first experiment demonstrates that the proposed HW-NAS is capable of generating compact neural network ar-

chitectures tailored for ultra-low-power microcontrollers with varying hardware capabilities.

- A second experiment compares the proposed HW-NAS with state-of-the-art methods, showing that it can deliver cutting-edge tiny CNNs optimized for ultra-low-power computing platforms.
- A third experiment demonstrates that the proposed HW-NAS can successfully adapt the cost of the search procedure to the constraints on execution time and energy budget set by an embedded device.
- A fourth experiment proves that the proposed HW-NAS can run on different embedded devices.

a) *Gateways*: The proposed HW-NAS has been executed on three different single-board computers that can serve as IoT gateways: the Raspberry Pi 4 (RPi 4), Raspberry Pi 3 (RPi 3), and Raspberry Pi Zero 2 (RPiZ 2). Table II summarizes the key features of the boards: SoC information, available RAM, and measured peak power consumption. To evaluate the maximum power consumption of each device, the “stress-ng” utility from Debian was used. The following command line forced full utilization of all CPU cores and of the main memory of the device: “stress-ng -cpu 0 -cpu-method all -vm 0 -vm-bytes 100% -vm-method all -verify -t 1m -v”. A power meter, directly connected to the power supply, was used to measure the maximum power consumption of each device. The same power meter was also used to monitor the energy consumption and execution time of the devices during the HW-NAS process.

TABLE II: Hardware features of the embedded devices targeted in this work.

Embedded Device	SoC	RAM	Power Consumption
Raspberry Pi	Broadcom	[GiB]	[W]
4 Model B	BCM2711	4	5.6
3 Model B	BCM2837	1	4.3
Zero 2	BCM2710A1	0.5	2.8

b) *Sensor Nodes*: Three different microcontrollers from the ultra-low-power production line of ST Microelectronics—namely, the L010RBT6, U083RCT6, and L412KBU3—have been selected as computing units for the sensor nodes. Table III shows the available RAM, Flash memory, and CoreMark score for each microcontroller.

TABLE III: Hardware features of the MCUs targeted in this work.

MCU	RAM	Flash	CoreMark
STM32	[kiB]	[kiB]	score
ultra-low-power line			
L010RBT6	20	128	75
U083RCT6	32	256	134
L412KBU3	40	128	273

c) *NAS*: The proposed HW-NAS indeed implemented a stand-alone process to monitor the execution time and to estimate the energy consumption in real-time. Such a backup process prevents the HW-NAS to run when either the bound on the execution time or the bound on the energy budget is reached. In fact, when dealing with an IoT gateway, non-real-time Operating Systems (OS), such as Debian or Windows,

introduce unpredictable overloads for maintaining the system, which cannot be easily estimated. As a result, the search space S_α returned by Alg. 2 might result too large if the goal is to meet the constraints ξ_{Time}, ξ_{Energy} . Therefore, the software implementation of the HW-NAS is set to generate as output an architecture even when is not possible to fully explore the search space S_α as expected.

d) Hyper-parameters: The maximum validation accuracy of a candidate architecture is evaluated after three epochs of training, with a validation split of 30% of the training data, a batch size of 16, and a learning rate of 1×10^{-3} . The Adam optimizer [48] is adopted for this process. In the case of time series the architectures were evaluated after fifty epochs of training.

e) Deployment: Both Quantization Aware Training (QAT) and Post Training Quantization (PTQ) are utilized to obtain 8-bit quantized models. The memory features, such as Flash and RAM occupancy, have been measured using the STM utility called "stm32tflm", which is included in the X-CUBE-AI software package. This utility provides specific measurements for running the obtained architecture on STM32 microcontrollers using the TFLite Micro runtime. Additionally, the proposed method computes MAC operations and architectural details, denoted as (k, c) .

The architectures selected by the proposed strategy are finally trained for 100 epochs, 500 in the case of time series, with a validation split of 10%, a batch size of 128, and a learning rate of $1e^{-2}$, using the Adam optimizer and implementing validation-based checkpointing. The training data is augmented by applying random horizontal flips and random rotations within the range $[-\frac{2}{5}\pi; \frac{2}{5}\pi]$.

The reported latency is measured on the NUCLEO-L432KC using the 'validation on target' feature provided by X-CUBE-AI. X-CUBE-AI version 8.1.0, along with STM32CubeIDE version 1.17.0, was used to program the boards.

f) Datasets: Four different datasets have been considered for the experiments:

- the Visual Wake Words dataset [49], which is a tinyML benchmark [50] containing 123,000 images that identify whether a person is present in the image;
- the CIFAR-10 dataset, which is a tinyML benchmark [50] consisting of 60,000 32x32 color images organized into 10 classes, with 6,000 images per class;
- the Melanoma Skin Cancer dataset [51], as a case for IIoT applications, which consists of 10,000 images used to identify whether a sample is malignant or not;
- the Case Western Reserve University (CWRU) dataset [52], as a case for IIoT application, which consists of accelerometer data coming from rolling bearings with different injected faults.

The test accuracy of the resulting architectures is evaluated on the public test set of the dataset used. In the case of the Visual Wake Words dataset, only one-tenth of the training split was used to run the HW-NAS, while the entire training split was utilized to train the resulting architecture.

VI. RESULTS

A. Adaptability to ultra-low-power microcontrollers

The first experiment aims to show that the proposed HW-NAS can generate custom architectures for different ultra-low-power computing platforms, which may play the role of sensor nodes in IoT applications. The assumption in this case is that the HW-NAS can run on a standard computing platform, i.e., without constraints on the available resources.

Table IV presents the results, which are organized into three subtables, one for each dataset. Each subtable gives - for a target microcontroller- the configuration of the selected architecture, the hardware requirements of the deployed model (specifically RAM, Flash memory, and MAC operations), the test accuracy, and latency for a single prediction.

TABLE IV: Performance of the proposed HW-NAS when targeting ultra-low-power computing platforms.

Target MCU	Arch. (k,c)	RAM [kiB]	Flash [kiB]	MAC [MM]	Test Acc. [%]	Lat. [ms]
Visual Wake Words ($vi = 50 \times 50$)						
L010RBT6	(3,4)	19	10.8	0.4	71	42
U083RCT6	(5,5)	24.5	22.7	0.9	75.2	63.2
L412KBU3	(8,3)	31	18.8	2	78.3	79.1
Melanoma Skin Cancer ($vi = 50 \times 50$)						
L010RBT6	(3,5)	18.5	8.1	0.4	84.2	43.2
U083RCT6	(6,4)	26.5	20.4	1.3	88.5	67.3
L412KBU3	(9,4)	34	35.7	2.6	90	129.8
CIFAR10 ($vi = 32 \times 32$)						
L010RBT6	(6,5)	14	28.7	0.5	63.3	31.6
U083RCT6	(9,4)	17	36	1.1	67.1	60
L412KBU3	(13,4)	21.5	65.2	2.1	70.9	97.5

Experimental outcomes prove that the features of the selected architectures scale with the available resources on the target platforms. As resources increase, the HW-NAS selects an architecture that requires more RAM, Flash memory, and MACs. The test accuracy improves accordingly, just as latency increases.

B. State of the art comparison

The proposed method is compared with four state-of-the-art NAS techniques: MCUNet [24], Micronets [25], ColabNAS [26], and NanoNAS [28]. The first two studies focused on high-performance MCUs, while the last two targeted ultra-low-power MCUs. The comparison is carried out using the Visual Wake Words dataset [49], a standard benchmark for tiny machine learning visual applications [53].

To ensure a fair comparison, Table V lists the smallest models selected by NAS targeting high-performance MCUs [24] [25] and the largest models selected by NAS targeting ultra-low-power MCUs [26] [28]. For the method proposed in this work, the largest architecture among those reported in Section VI-A has been selected. For each NAS, the table gives test accuracy, RAM, Flash memory, and MAC operations.

The results show that MCUNet produces the network with the highest test accuracy, albeit with the largest architecture in terms of RAM, Flash memory, and MAC operations. In contrast, the proposed method achieves the second-best test accuracy while maintaining the lowest resource usage

in terms of Flash and the second-lowest in terms of RAM and MAC operations. ColabNAS ranks third in test accuracy, while featuring larger RAM and Flash. NanoNAS is fourth in test accuracy while featuring the lowest RAM occupancy and MAC operations. Micronets occupy the last position, delivering the lowest test accuracy but being only slightly smaller than MCUNet in RAM and Flash usage.

TABLE V: Comparison with existing HW-NAS on VWW dataset.

Work	Acc [%]	RAM [kiB]	Flash [kiB]	MAC [MM]
MCUNet [24]	87.4	168.5	530.5	6
Micronets [25]	76.8	70.5	273.8	3.3
ColabNAS [26]	77.6	31.5	20.83	2
NanoNAS [28]	77	28.5	23.7	1.3
Proposal	78.3	31	18.8	2

C. Running HW-NAS on energy and time constrained devices

This experiment aims to demonstrate that the proposed HW-NAS exploits an adaptive strategy, which can fit the available resource budget. A Raspberry Pi Zero 2 has been selected as the gateway for this experiment, while STM32L412KBU3 has been used as sensor node.

First, the extensive search space \hat{S}_α for the STM32L412KBU3 device has been generated according to Algorithm 1. Then, the energy budget and the time budget required by Raspberry Pi Zero 2 to get $S_\alpha = \hat{S}_\alpha$ have been empirically estimated. Hence, in practice, it was assessed that with an energy budget of 16.5 Wh and a time budget of 9:51 (expressed hereafter as hh:mm) one can avoid any cropping of \hat{S}_α (according to Algorithm 2).

Table VI reports the results of three experiments in which the proposed HW-NAS run on the selected gateway -Raspberry Pi Zero 2- with as many different energy / time budgets. The first row of the table shows the outcomes of the first experiment, in which the available budget was 16.5 Wh for energy and 9:51 for time. In this table, the second column gives the actual cost of the search procedure when running on the gateway; the third column gives the actual size of S_α after cropping, expressed as percentage with respect to \hat{S}_α ; the fourth column gives the percentage of S_α that has been covered by implementing the search strategy formalized in Algorithm 3. The second part of the table provides details about the architecture generated by the NAS: RAM usage, Flash Memory usage, MAC operations, test accuracy on the VWW dataset, and latency measured on the STM32L412KBU3 device. The second row of the table reports the outcomes of the experiment in which the available budget was cut down to 11.0 Wh for energy and 6:34 for time, i.e., two-third of the best possible budget. Analogously, the third row of the table reports the results of the experiment in which the available budget further reduced to 5.50 Wh for energy and 3:17 for time, i.e., one-third of the best possible budget.

The table shows that, as expected, only when the available budget was, respectively, 16.5 Wh for energy and 9:51 for time, the search procedure could work on the entire search

TABLE VI: Performance of the proposed HW-NAS when running on Raspberry Pi Zero 2 with different energy/time budgets, with STM32L412KBU3 as a target.

NAS Details				Resulting Architecture				
Budget [Wh][time]	Cost [Wh][time]	S_α [%]	Expl. [%]	RAM [kiB]	Flash [kiB]	MAC [MM]	Acc. [%]	Lat. [ms]
16.5 - 9:51	16.5 - 9:51	100	51	28.5	16.1	1.6	77.8	87.4
11.0 - 6:34	11.0 - 6:30	33	98	21.5	13.1	0.7	73.1	43.3
5.50 - 3:17	5.41 - 3:17	15	95	18.5	7.2	0.1	66	24.6

space \hat{S}_α without any cropping: in fact, S_α was 100% of \hat{S}_α . Indeed, the search strategy managed to explore only 51% of S_α to find the best architecture. In the other two experiments, the exploration of the search space S_α was halted due to resource depletion. In the second experiment, the search procedure was stopped after the threshold on the available energy was reached; four minutes were still available in terms of time budget. S_α was 33% of \hat{S}_α and, notably, the search strategy was able to explore 98% of S_α due to its small size. In the last experiment, conversely, the search procedure was stopped after the threshold on the available time was reached. S_α was 15% of \hat{S}_α and the search strategy was able to explore 95% of S_α due to its very small size.

Obviously, the three experiments led to as many different architectures to be deployed on the target device. The footprint and the latency of the selected architecture reduce as the available budget for running the HW-NAS drops. In fact, generalization performance also deteriorates. Overall, the results obtained with this experimental session demonstrate that the proposed HW-NAS can adapt to the available budget by adjusting the portion of the search space to be explored.

D. Effect of different platform on HW-NAS execution

This experiment aims to demonstrate that the proposed HW-NAS can adapt its search strategy to the characteristics of the specific gateway. To this end, three different devices that can be used as gateways have been involved in the experiment: Raspberry Pi Zero 2, Raspberry Pi 3 Model B and the Raspberry Pi 4 Model B. Again, STM32L412KBU3 has been used as sensor node.

Table VII reports on the result obtained by running the HW-NAS on the three selected gateways always with the same energy/time budget: 16.5 Wh for energy and 9:51 for time. As stated in the previous section, this budget allows one to avoid any cropping of the search space \hat{S}_α defined by the STM32L412KBU3 microcontroller when working on a Raspberry Pi Zero 2.

The table gives, for each experiment: the gateway; the actual cost of the search procedure on that gateway, the actual size of S_α after cropping, expressed as percentage with respect to \hat{S}_α ; the percentage of S_α that has been covered by implementing the search strategy formalized in Algorithm 3. In addition, the second part of the table provides details about the architecture generated by the NAS: RAM usage, Flash Memory usage, MAC operations, test accuracy on the VWW dataset, and latency measured on the STM32L412KBU3 device.

The results show that when the Raspberry Pi 3 was adopted as a gateway, the search process stopped after 5 hours and 4

minutes due to energy depletion. In fact, Raspberry Pi 3 is generally less efficient than Raspberry Pi 0 in terms of energy consumption. Consequently, the proposed HW-NAS procedure had to cut \hat{S}_α : eventually, S_α was 38% of the extensive search space. Conversely, the HW-NAS was able to run on the the Raspberry Pi 4 without depleting the available budget. In this case, S_α was 100% of the extensive search space; therefore, no cropping was needed. When running on the Raspberry Pi 4, the proposed HW-NAS managed to explore 57% of the search space \hat{S}_α in 3 hours and 25 minutes, while 9 hours and 51 minutes were required on the Raspberry Pi Zero to explore 51% of \hat{S}_α .

These results demonstrate that the proposed HW-NAS is effective in adapting to different execution platforms. The HW-NAS can set the size of S_α according to the available budget and the features of the gateway. In addition, the search strategy (Algorithm 3) can adequately explore both small search spaces and large search spaces.

TABLE VII: Performance of the proposed HW-NAS with STM32L412KBU3 as a target platform, maintaining a fixed resource budget while varying the execution platforms

RPI	NAS Details				Resulting Architecture				
	Cost [Wh]	S_α [time]	Expl. [%]	Expl. [%]	RAM [kiB]	Flash [kiB]	MAC [MM]	Acc. [%]	Lat. [ms]
Z2	16.5 - 9:51	100	51	28.5	16.1	1.6	77.8	87.4	
3	16.5 - 5:04	38	98	21	9.5	0.6	72.7	44.5	
4	14.2 - 3:25	100	57	35	52.1	2.6	75.9	125.7	

E. Time Series

This experiment aims to show that the proposed technique can also be applied to time series, leading to state-of-the-art results on the CWRU dataset [52]. A Raspberry Pi 4 has been adopted as gateway and a STM32-L010RBT6 as target sensor node. The algorithm has been run without energy and time constraints. The resulting network has been compared with the state-of-the-art work published by Chen et al. [54].

TABLE VIII: Resulting Architecture VS Reference Architecture

Work	Search Cost		(k, c)	Resulting Architecture				
	Energy [Wh]	Time [hh]:[mm]		RAM [kiB]	Flash [kiB]	MAC [MM]	Test Acc. [%]	Latency [ms]
Proposal [54]	6.4	1:52	(6, 4)	13.5	12.9	0.6	99.5	34
	n/a	n/a	n/a	66.5	163.4	0.2	99.3	38.2

Table VIII reports on the results of this experiment. The table compares the neural network generated by the proposed HW-NAS with that adopted in [54]. The table gives RAM usage, Flash Memory usage, MAC operations, test accuracy on the CWRU dataset, and latency measured on the STM32-L010RBT6 device.

The HW-NAS proved capable of producing a tiny CNN that needs 4.9 times less RAM and 12.7 times less Flash than the one adopted in [54]. Notably, such CNN scored higher test accuracy while decreasing latency. The resulting architecture can run on an STM32-L010RBT6, i.e., an ultra-low-power MCU, and is therefore suitable for IIoT sensor nodes.

The search procedure was carried out on a Raspberry Pi 4 for 1 hour and 52 minutes; the energy consumption was 6.4 Wh. The average CPU usage per core was 67.9% and the average RAM usage was 593 MB (out of 4 GB available). This result makes the proposed method compatible with existing services running in deployed gateways of IoT networks.

VII. CONCLUSION

Using the proposed method, NNs can be automatically designed at the edge, without the data leaving the site of collection. This enables the usage of NNs even in privacy-sensitive applications: customized and hardware-friendly NNs can be designed for personalized healthcare in the IIoT framework, maintaining the data inside the hospital; or for IFD and quality control in production plants in the IIoT framework, preserving industrial secrets that could be disclosed by the data.

A thorough experimental campaign proved that the proposed HW-NAS a) can run on embedded devices that frequently serve as gateway nodes in IoT networks and b) achieves state-of-the-art results when ultra-low-power MCUs are selected as computing source for sensor nodes.

To the best of the authors' knowledge the literature does not provide approaches to the design of HW-NAS that could allow the search process to run on an embedded device. The only exception is a recent work [28] published by the authors of this paper; however, the HW-NAS presented in this work improves on [28] in that it can dynamically adapt to the features of the gateway, thus enabling its use in IoT networks. The proposed approach may further simulate the development of lightweight HW-NAS that can run on resource-constrained devices, thus opening new fields of applications. Future works may focus, for example, on training-free evaluation techniques to further reduce the search cost.

VIII. ACKNOWLEDGMENT

Project funded under the National Recovery and Resilience Plan (NRRP), Mission 4 Component 2 Investment 1.1 - Call for tender No. 1409 published on Sept 14, 2022 by the Italian Ministry of University and Research (MUR) funded by the European Union - NextGenerationEU - Project Title "LEARN - muLtimodal Edge computing-bAsed weaRable exoskeletons for assistance in daily life" - CUP: D53D23016200001, J53D23014090001 - Grant Assignment Decree No. 1383 adopted on September 01, 2023 by the Italian Ministry of University and Research (MUR).



REFERENCES

- [1] G. Dong, M. Tang, Z. Wang, J. Gao, S. Guo, L. Cai, R. Gutierrez, B. Campbel, L. E. Barnes, and M. Boukhechba, "Graph neural networks in iot: A survey," *ACM Transactions on Sensor Networks*, vol. 19, no. 2, pp. 1–50, 2023.
- [2] M. N. Bhuiyan, M. M. Rahman, M. M. Billah, and D. Saha, "Internet of things (iot): A review of its enabling technologies in healthcare applications, standards protocols, security, and market opportunities," *IEEE Internet of Things Journal*, vol. 8, no. 13, pp. 10474–10498, 2021.
- [3] J. Chen, K. Wen, J. Xia, R. Huang, Z. Chen, and W. Li, "Knowledge embedded autoencoder network for harmonic drive fault diagnosis under few-shot industrial scenarios," *IEEE Internet of Things Journal*, 2024.
- [4] S.-C. Wu, P.-L. Hung, and A. L. Swindlehurst, "Ecg biometric recognition: unlinkability, irreversibility, and security," *IEEE Internet of Things Journal*, vol. 8, no. 1, pp. 487–500, 2020.

- [5] C. Li and B. Palanisamy, "Privacy in internet of things: From principles to technologies," *IEEE Internet of Things Journal*, vol. 6, no. 1, pp. 488–505, 2018.
- [6] H. Habibzadeh, K. Dinesh, O. R. Shishvan, A. Boggio-Dandry, G. Sharma, and T. Soyata, "A survey of healthcare internet of things (hiot): A clinical perspective," *IEEE Internet of Things Journal*, vol. 7, no. 1, pp. 53–71, 2019.
- [7] G. Aceto, V. Persico, and A. Pescapé, "Industry 4.0 and health: Internet of things, big data, and cloud computing for healthcare 4.0," *Journal of Industrial Information Integration*, vol. 18, p. 100129, 2020.
- [8] M. Serror, S. Hack, M. Henze, M. Schuba, and K. Wehrle, "Challenges and opportunities in securing the industrial internet of things," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 5, pp. 2985–2996, 2020.
- [9] 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," *ACM Computing Surveys (CSUR)*, vol. 54, no. 4, pp. 1–34, 2021.
- [10] C. Zhang, X. Yuan, Q. Zhang, G. Zhu, L. Cheng, and N. Zhang, "Toward tailored models on private aiot devices: Federated direct neural architecture search," *IEEE Internet of Things Journal*, vol. 9, no. 18, pp. 17 309–17 322, 2022.
- [11] K. Wang, P. Xu, C.-M. Chen, S. Kumari, M. Shojafar, and M. Alazab, "Neural architecture search for robust networks in 6g-enabled massive iot domain," *IEEE internet of things journal*, vol. 8, no. 7, pp. 5332–5339, 2020.
- [12] K. T. Chitty-Venkata and A. K. Somani, "Neural architecture search survey: A hardware perspective," *ACM Computing Surveys*, vol. 55, no. 4, pp. 1–36, 2022.
- [13] K. Adhinugraha, W. Rahayu, T. Hara, and D. Taniar, "On internet-of-things (iot) gateway coverage expansion," *Future Generation Computer Systems*, vol. 107, pp. 578–587, 2020.
- [14] H. Lu, M. Du, X. He, K. Qian, J. Chen, Y. Sun, and K. Wang, "An adaptive neural architecture search design for collaborative edge-cloud computing," *IEEE Network*, vol. 35, no. 5, pp. 83–89, 2021.
- [15] A. M. Garavagno, E. Ragusa, A. Frisoli, and P. Gastaldo, "Running hardware-aware neural architecture search on embedded devices under 512mb of ram," in *2024 IEEE International Conference on Consumer Electronics (ICCE)*. IEEE, 2024, pp. 1–2.
- [16] J. Zhang and D. Tao, "Empowering things with intelligence: a survey of the progress, challenges, and opportunities in artificial intelligence of things," *IEEE Internet of Things Journal*, vol. 8, no. 10, pp. 7789–7817, 2020.
- [17] B. Guo, S. He, M. Shi, K. Yu, J. Chen, and X. Shen, "Semantic-darts: Elevating semantic learning for mobile differentiable architecture search," *IEEE Internet of Things Journal*, 2024.
- [18] C. Lin, P. Yang, Q. Wang, Y. Guo, and Z. Wang, "Multi-hardware adaptive latency prediction for neural architecture search," *IEEE Internet of Things Journal*, 2024.
- [19] B. S. Prabakaran, A. Akhtar, S. Rehman, O. Hasan, and M. Shafique, "Bionetexplorer: Architecture-space exploration of biosignal processing deep neural networks for wearables," *IEEE Internet of Things Journal*, vol. 8, no. 17, pp. 13 251–13 265, 2021.
- [20] Y. Du, J. Liu, X. Wang, and P. Wang, "Ssvp-based emotion recognition for iot via multiobjective neural architecture search," *IEEE Internet of Things Journal*, vol. 9, no. 21, pp. 21 432–21 443, 2022.
- [21] M. Dissem, M. Amayri, and N. Bouguila, "Neural architecture search for anomaly detection in time series data of smart buildings: A reinforcement learning approach for optimal autoencoder design," *IEEE Internet of Things Journal*, 2024.
- [22] E. Ragusa, F. Zonzini, L. De Marchi, and R. Zunino, "Compression–accuracy co-optimization through hardware-aware neural architecture search for vibration damage detection," *IEEE Internet of Things Journal*, vol. 11, no. 19, pp. 31 745–31 757, 2024.
- [23] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 2820–2828.
- [24] J. Lin, W.-M. Chen, Y. Lin, C. Gan, S. Han *et al.*, "Mcnunet: Tiny deep learning on iot devices," *Advances in Neural Information Processing Systems*, vol. 33, pp. 11 711–11 722, 2020.
- [25] C. Banbury, C. Zhou, I. Fedorov, R. Matas, U. Thakker, D. Gope, V. Janapa Reddi, M. Mattina, and P. Whatmough, "Micronets: Neural network architectures for deploying tinymml applications on commodity microcontrollers," *Proceedings of Machine Learning and Systems*, vol. 3, pp. 517–532, 2021.
- [26] A. M. Garavagno, D. Leonardis, and A. Frisoli, "Colabnas: Obtaining lightweight task-specific convolutional neural networks following oc-cam's razor," *Future Generation Computer Systems*, vol. 152, pp. 152–159, 2024.
- [27] A. M. Garavagno, E. Ragusa, A. Frisoli, and P. Gastaldo, "A hardware-aware neural architecture search algorithm targeting low-end microcontrollers," in *2023 18th Conference on Ph. D Research in Microelectronics and Electronics (PRIME)*. IEEE, 2023, pp. 281–284.
- [28] —, "An affordable hardware-aware neural architecture search for deploying convolutional neural networks on ultra-low-power computing platforms," *IEEE Sensors Letters*, 2024.
- [29] J. Skarding, B. Gabrys, and K. Musial, "Foundations and modeling of dynamic networks using dynamic graph neural networks: A survey," *IEEE Access*, vol. 9, pp. 79 143–79 168, 2021.
- [30] Y. Han, G. Huang, S. Song, L. Yang, H. Wang, and Y. Wang, "Dynamic neural networks: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 11, pp. 7436–7456, 2021.
- [31] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on knowledge and data engineering*, vol. 22, no. 10, pp. 1345–1359, 2009.
- [32] J. Yang, H. Zou, Y. Zhou, and L. Xie, "Learning gestures from wifi: A siamese recurrent convolutional architecture," *IEEE Internet of Things Journal*, vol. 6, no. 6, pp. 10 763–10 772, 2019.
- [33] H. Zou, J. Yang, Y. Zhou, L. Xie, and C. J. Spanos, "Robust wifi-enabled device-free gesture recognition via unsupervised adversarial domain adaptation," in *2018 27th International Conference on Computer Communication and Networks (ICCCN)*. IEEE, 2018, pp. 1–8.
- [34] X. Yu, F. Dong, B. Xia, S. Yang, E. Ding, and W. Yu, "An intelligent fault diagnosis scheme for rotating machinery based on supervised domain adaptation with manifold embedding," *IEEE Internet of Things Journal*, vol. 10, no. 1, pp. 953–972, 2022.
- [35] G. Santamato, A. M. Garavagno, M. Solazzi, and A. Frisoli, "Leveraging systems' non-linearity to tackle the scarcity of data in the design of intelligent fault diagnosis systems," *Nonlinear Dynamics*, vol. 112, no. 18, pp. 16 153–16 166, 2024.
- [36] H. Benmeziame, K. E. Maghraoui, H. Ouarnoughi, S. Niar, M. Wistuba, and N. Wang, "A comprehensive survey on hardware-aware neural architecture search," *arXiv preprint arXiv:2101.09336*, 2021.
- [37] H. Cai, C. Gan, T. Wang, Z. Zhang, and S. Han, "Once-for-all: Train one network and specialize it for efficient deployment," *arXiv preprint arXiv:1908.09791*, 2019.
- [38] L. Li and A. Talwalkar, "Random search and reproducibility for neural architecture search," in *Uncertainty in artificial intelligence*. PMLR, 2020, pp. 367–377.
- [39] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le, "Learning transferable architectures for scalable image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 8697–8710.
- [40] H. Liu, K. Simonyan, and Y. Yang, "Darts: Differentiable architecture search," *arXiv preprint arXiv:1806.09055*, 2018.
- [41] H. Cai, L. Zhu, and S. Han, "Proxylessnas: Direct neural architecture search on target task and hardware," *arXiv preprint arXiv:1812.00332*, 2018.
- [42] B. Wu, X. Dai, P. Zhang, Y. Wang, F. Sun, Y. Wu, Y. Tian, P. Vajda, Y. Jia, and K. Keutzer, "Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 10 734–10 742.
- [43] L. L. Zhang, Y. Yang, Y. Jiang, W. Zhu, and Y. Liu, "Fast hardware-aware neural architecture search," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 692–693.
- [44] M. Shanker, M. Y. Hu, and M. S. Hung, "Effect of data standardization on neural network training," *Omega*, vol. 24, no. 4, pp. 385–397, 1996.
- [45] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International conference on machine learning*. pmlr, 2015, pp. 448–456.
- [46] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [47] M. Lin, Q. Chen, and S. Yan, "Network in network," *arXiv preprint arXiv:1312.4400*, 2013.
- [48] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [49] A. Chowdhery, P. Warden, J. Shlens, A. Howard, and R. Rhodes, "Visual wake words dataset," *arXiv preprint arXiv:1906.05721*, 2019.

- [50] C. Banbury, V. J. Reddi, P. Torelli, J. Holleman, N. Jeffries, C. Kiraly, P. Montino, D. Kanter, S. Ahmed, D. Pau *et al.*, "Mlperf tiny benchmark," *arXiv preprint arXiv:2106.07597*, 2021.
- [51] Melanoma skin cancer dataset. [Online]. Available: <https://www.kaggle.com/datasets/hasnainjaved/melanoma-skin-cancer-dataset-of-10000-images>
- [52] W. A. Smith and R. B. Randall, "Rolling element bearing diagnostics using the case western reserve university data: A benchmark study," *Mechanical systems and signal processing*, vol. 64, pp. 100–131, 2015.
- [53] C. R. Banbury, V. J. Reddi, M. Lam, W. Fu, A. Fazel, J. Holleman, X. Huang, R. Hurtado, D. Kanter, A. Lokhmotov *et al.*, "Benchmarking tinyml systems: Challenges and direction," *arXiv preprint arXiv:2003.04821*, 2020.
- [54] C.-C. Chen, Z. Liu, G. Yang, C.-C. Wu, and Q. Ye, "An improved fault diagnosis using 1d-convolutional neural network model," *Electronics*, vol. 10, no. 1, p. 59, 2020.