arXiv:2408.16535v1 [cs.LG] 29 Aug 2024

TinyTNAS: GPU-Free, Time-Bound, Hardware-Aware Neural Architecture Search for TinyML Time Series Classification

Bidyut Saha, Riya Samanta, Soumya K. Ghosh, and Ram Babu Roy

Indian Institute of Technology Kharagpur, India
bidyutsaha@kgpian.iitkgp.ac.in, riya.samanta@iitkgp.ac.in,
skg@cse.iitkgp.ac.in, rambabu@see.iitkgp.ac.in

Abstract. In this work, we present TinyTNAS, a novel hardware-aware multi-objective Neural Architecture Search (NAS) tool specifically designed for TinyML time series classification. Unlike traditional NAS methods that rely on GPU capabilities, TinyTNAS operates efficiently on CPUs, making it accessible for a broader range of applications. Users can define constraints on RAM, FLASH, and MAC operations to discover optimal neural network architectures within these parameters. Additionally, the tool allows for time-bound searches, ensuring the best possible model is found within a user-specified duration. By experimenting with benchmark datasets—UCI HAR, PAMAP2, WISDM, MIT-BIH, and PTB Diagnostic ECG Database—TinyTNAS demonstrates state-of-the-art accuracy with significant reductions in RAM, FLASH, MAC usage, and latency. For example, on the UCI HAR dataset, TinyTNAS achieves a 12x reduction in RAM usage, 144x reduction in MAC operations, and 78x reduction in FLASH memory while maintaining superior accuracy and reducing latency by 149x. Similarly, on the PAMAP2 and WISDM datasets, it achieves a 6x reduction in RAM usage, 40x reduction in MAC operations, an 83x reduction in FLASH, and a 67x reduction in latency, all while maintaining superior accuracy. Notably, the search process completes within 10 minutes in a CPU environment. These results highlight TinyTNAS's capability to optimize neural network architectures effectively for resource-constrained TinyML applications, ensuring both efficiency and high performance. The code for TinyTNAS is available at the GitHub repository and can be accessed through `https://github.com/BidyutSaha/TinyTNAS.git`.

Keywords: Neural Architecture Search · TinyML · Hardware-Aware Neural Networks · Neural Network Optimization · Time Series Data

1 Introduction

Neural Architecture Search (NAS) offers significant advantages, including automating the design of neural network architectures, thereby saving considerable time and expert effort required in manual design [6]. It also enhances performance by exploring and optimizing a wide range of architectures that might not have been feasible through manual methods, leading to more efficient and effective neural networks [29]. Traditionally, NAS relies on extensive computational resources, predominantly GPUs, to explore

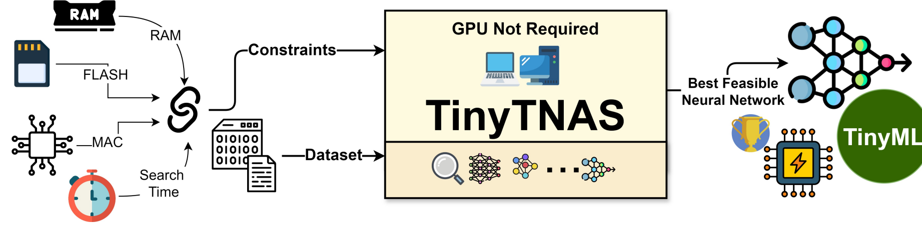

Fig. 1. Concept Diagram of TinyTNAS

vast search spaces and identify optimal configurations. This involves determining the best combination of layers, operations, and hyperparameters to maximize model performance for specific tasks [31]. While effective, these methods are resource-intensive and often impractical for applications requiring rapid deployment or operating within hardware constraints. In response to the limitations of traditional NAS, *hardware-aware NAS* methods have been developed. These approaches integrate hardware parameters such as memory usage, computational complexity (measured in FLOPs or MACs), and power consumption into the architecture search process [5]. The hardware-aware NAS ensures that the resulting models are not only high-performing but also suitable for deployment on resource-constraint devices, such as microcontrollers (MCUs) and edge devices. This is particularly important for applications in the Internet of Things (IoT), where computational resources are limited. *Multi-objective NAS* further extends this concept by simultaneously optimizing for multiple criteria, such as accuracy, latency, and resource usage [8]. This approach balances trade-offs between different objectives, enabling the discovery of architectures that meet diverse and sometimes conflicting requirements. Multi-objective optimization is crucial for developing models that are both efficient and effective across various deployment scenarios.

As the demand for intelligent edge devices grows, there is an increasing need to deploy machine learning models on MCU and other resource-constraint hardware. This is where TinyML comes into play. TinyML represents the intersection of machine learning and embedded systems, focusing on deploying machine learning models on MCUs and other resource-constraint devices [25,21]. The primary motivation for TinyML is to bring intelligence to edge devices, enabling real-time decision-making and reducing reliance on cloud-based processing. This shift offers several advantages, including lower latency, reduced bandwidth usage, and enhanced privacy since data can be processed locally [12]. However, developing TinyML solutions presents significant challenges. The limited computational power, memory, and storage available on MCUs necessitate highly efficient models [7]. Traditional approaches to model optimization for TinyML often involve manual tuning and simplification of pre-designed architectures, which is both time-consuming and sub-optimal. Moreover, balancing performance with resource constraints requires a deep understanding of both machine learning and embedded system design. Despite these challenges, there have been notable successes in the TinyML domain. For instance, models optimized for voice recognition and keyword spotting have been successfully deployed on MCUs, enabling applications like always-on voice

assistants [2]. Similarly, accelerometer-based activity recognition systems have been implemented on wearable devices, demonstrating the potential of TinyML for fitness tracking, health monitoring and human computer interactions [21,23,22,24].

Integrating NAS with TinyML offers a promising solution to overcome the challenges of model optimization for resource-constraint environments. By leveraging NAS, it is possible to automate the design of efficient neural network architectures tailored to the specific hardware constraints of TinyML devices. This approach can significantly accelerate the development process and ensure optimal performance [10]. Despite the advancements in NAS and TinyML, there are still significant challenges in achieving state-of-the-art (SOTA) performance while adhering to strict resource constraints. Traditional NAS methods are typically impractical due to their reliance on GPUs and extensive computational requirements. Therefore, there is a critical need for NAS methods that can operate efficiently on CPUs and deliver high-quality models within a reasonable time frame.

In this context, we introduce **TinyTNAS**[1], a novel *hardware-aware multi-objective NAS tool* specifically designed for *TinyML time series classification* (see Figure 1). Unlike existing methods, TinyTNAS operates efficiently on CPUs, making it accessible and practical for a broader range of applications. Users can define constraints on RAM, FLASH, and MAC operations to discover optimal neural network architectures that meet these parameters. Additionally, TinyTNAS allows for *time-bound* searches, ensuring the best possible model is found within a user-specified duration. By experimenting with benchmark datasets—UCIHAR [20], PAMAP2 [19], WISDM [26], MIT-BIH [16], and PTB Diagnostic ECG Database [3]—TinyTNAS demonstrates state-of-the-art accuracy with significant reductions in RAM, FLASH, and MAC usage.

Our objective is to provide a comprehensive solution that bridges the gap between NAS and TinyML, enabling the efficient deployment of high-accuracy models on resource-constraint devices. TinyTNAS represents a substantial advancement in the TinyML domain, offering a practical tool for optimizing neural networks within stringent hardware constraints and delivering superior performance in time-sensitive environments.

2 Related Work

Neural Architecture Search (NAS) has revolutionized the automated design of neural networks, significantly enhancing performance and reducing the need for manual intervention. Traditional NAS methods, such as [31] and [18], have primarily relied on extensive GPU resources to explore vast search spaces, resulting in high-performing but computationally expensive models. These methods often overlook the constraints of deploying models on resource-limited devices like microcontrollers (MCUs) used in TinyML applications. Recent advancements in hardware-aware NAS (HW NAS), exemplified by works like [27]and [4], integrate hardware metrics such as memory usage and computational complexity into the search process, producing models that are more suitable for deployment on edge devices. However, these approaches still predominantly rely on GPUs for model optimization, limiting their accessibility and practicality for broader applications.

[1] The code repository is available at `https://github.com/BidyutSaha/TinyTNAS.git`.

In the domain of TinyML, recent works like MCUNet [13] and MicroNets [2] have made significant strides. MCUNet introduces a joint design framework that optimizes both the neural architecture and the inference engine to fit the resource constraints of MCUs, achieving impressive efficiency and accuracy. However, MCUNet and similar tools typically require heavy GPU resources for their computation, making them less feasible for deployment on CPU-only environments. Similarly, MicroNets focuses on building extremely compact and efficient neural networks suitable for MCUs by leveraging advanced pruning and quantization techniques. Notably, MicroNets also faces challenges with heavy GPU requirements, limiting its practicality in CPU-centric environments. Despite these advancements, a significant research gap remains in the specific optimization of NAS for time series classification on resource-constraint devices. Existing works often focus on image classification tasks, leaving time series data underexplored in CPU-centric environments.

Moreover, recent research [9] presented an HW NAS approach that can run on CPUs, producing tiny convolutional neural networks (CNNs) targeting low-end microcontrollers. However, their approach has notable limitations, such as using standard CNN layers when more computationally efficient layers could be employed. Additionally, their method relies on traditional grid search, which is computationally expensive, and focuses on image classification rather than time series data, without incorporating time-bound searches. This restricts its applicability for tasks requiring efficient temporal data processing and adherence to specified search durations.

Our work addresses these gaps by introducing TinyTNAS, a CPU-efficient, hardware-aware NAS tool specifically designed for TinyML time series classification. TinyTNAS optimizes neural network architectures within RAM, FLASH, and MAC constraints and ensures rapid deployment by completing searches within a user defined time window on CPUs, setting a new standard for efficiency in the TinyML domain. We compare the models generated by TinyTNAS against state-of-the-art (SOTA) models for benchmark datasets—UCI HAR [20], PAMAP2 [19], WISDM [26], MIT-BIH [16], and PTB Diagnostic ECG Database [3]—using existing methods known for achieving SOTA accuracy: CNN (1D) [30], DeepConvLSTM (1D) [17], and LSTM-CNN (1D) [28] for HAR, and CNN (1D) [11] for MIT-BIH and PTB-DED.

3 TinyTNAS

3.1 Search Space

In Neural Architecture Search (NAS), the *search space* refers to the complete set of possible neural network architectures explored by the algorithm. NAS encompasses three primary search spaces: layer-wise, cell-wise, and hierarchical. *Layer-wise NAS* focuses on optimizing individual layer configurations such as convolutional and pooling layers [31]. *Cell-wise NAS* identifies optimal repeating cell structures, enhancing scalability and efficiency [32,15]. *Hierarchical NAS* integrates layer-wise and cell-wise strategies across multiple abstraction levels for adaptive model design [14,5].

TinyTNAS utilizes a cell-wise search space. Each generated neural network architecture follows the template, depicted in Figure 2.

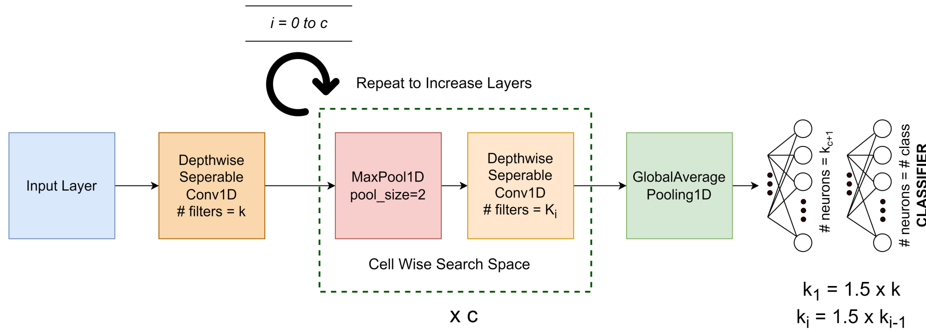

Fig. 2. Graphical representation of candidate architectures created by TinyTNAS, where k denotes the number of filters in the first depthwise separable 1D convolutional layer and c represents the number of repeating blocks. The kernel size of each depthwise separable 1D CNN is 3 with a stride of 1 and ReLU activation. The final dense layer uses a softmax activation function, while the preceding dense layer uses a ReLU activation function.

- Candidate architecture begins with a Depthwise Separable 1d Convolutional layer with stride=1, kernel size=3, and k number of kernels with ReLU activation.
- Followed by a block that may repeat 0 to c times, where c is the maximum number of possible MaxPooling1D layers of size=2 adaptable to the input shape from the dataset.
- Next a GlobalAveragePooling1D layer.
- Then a Dense layer with ReLU activation and k_{c+1} neurons.
- Finally, a classifier layer with Dense neurons corresponding to the number of classes and a softmax function.

The repeating block structure starts with MaxPooling1D size=2 followed by Depthwise Separable 1D Convolution with kernel size=3, stride=1, ReLU activation, and filter size k_i, where k_i follows a growth pattern. Specifically, $k_i = 1.5 \times k_{i-1}$, starting with $k_1 = 1.5 \times k$. The search space is constraint by user-defined hardware limits for RAM, FLASH, and MAC. The generated model is optimised by TensorFlow Lite using *integer quantization*, and its resourse requirement is estimated by MLTK profiler.

3.2 Search Algorithm

Our focus is on developing the TinyTNAS tool, which operates efficiently on CPUs, eliminating the need for a GPU, and generates architectures within a feasible time frame. Therefore, we avoid reinforcement learning or evolutionary algorithms and instead use a variant of grid search methods. Users can define constraints on RAM, FLASH, and MAC operations to discover optimal neural network architectures that meet these parameters. Additionally, TinyTNAS allows for time-bound searches, ensuring the best possible model is found within a user-specified duration.

Our proposed search algorithm, described in the Algorithm section, optimizes two dimensions to find the optimal architecture: k and c. Here, k represents the number of

filters in the first depthwise 1d convolutional layer, and c is the number of repeating blocks. The structure of the neural architecture based on k and c is detailed in the previous section and illustrated in Figure 3.

Unlike traditional grid search, which explores every point in the grid for k and c, our method reduces computational demand by generally doubling k and keeping c constant unless the accuracy drops in consecutive explored architectures. For a detailed workflow of our search algorithm, readers are referred to Algorithm 1.

It is important to note that, unlike many traditional methods, our candidate architectures are trained directly on the full target dataset for four epochs. Users can adjust this value: increasing it will require more search time but potentially yield better decisions, while decreasing it will reduce search time but may lead to suboptimal decisions. We determined four epochs through empirical analysis to balance this trade-off effectively.

Algorithm 1 TinyTNAS Main Module

Input: RAM_{max}, MAC_{max}, $Flash_{max}$, *search_time*, dataset *ds*
Output: First layer 1d CNN filter size K, Repeated block count C,
Initialization: Initialize $max_acc_found \leftarrow 0$, $pendings \leftarrow []$, $epoch \leftarrow 4$, $K \leftarrow 4$, $k \leftarrow 4$, $C \leftarrow 3$, $c \leftarrow 3$

```
1:  Start
2:  while True do
3:      if has_search_time then
4:          model, ram, flash, mac ← buildModel(k, c, ds)
5:          acc ← 0
6:          if checkFeasibility(ram, mac, flash, RAM_max, MAC_max, Flash_max) then
7:              acc ← trainModel(model, ds, epoch) {TensorFlow training}
8:              if max_acc_found < acc then
9:                  max_acc_found ← acc
10:                 K, k, C, pendings ← updateStatus(k, c, pendings)
11:                 continue
12:             end if
13:         end if
14:         is_continueable, k, c, K, C ←exploreDepth(acc, k, c, RAM_max, MAC_max, Flash_max, pendings, search_time, K, C)
15:         if is_continueable then
16:             continue
17:         else
18:             return K, C
19:         end if
20:     else
21:         return K, C
22:     end if
23: end while
```

4 Experiments

4.1 Datasets

In our experiments, we utilize five benchmark datasets spanning the domains of lifestyle, healthcare, and human-computer interaction. Human activity recognition directly impacts lifestyle and human-computer interaction domains, and indirectly contributes to healthcare. For this purpose, we employ UCIHAR, PAMAP2, and WISDM datasets. Additionally, in healthcare applications, we leverage ECG benchmark datasets such as MIT-BIH and PTB Diagnostic ECG Database. Details for each dataset are provided

Algorithm 2 updateStatus

Input: First layer 1d CNN filter size k, Repeated block count c, *pendings*
Output: K, k, C, *pendings*
Initialization: Initialize $multiplier \leftarrow 2$, $divider \leftarrow 4$
1: Start
2: $delta \leftarrow k \times multiplier - k$
3: $incr \leftarrow \lfloor \frac{delta}{divider} \rfloor$
4: $pendings \leftarrow [\,]$
5: **if** $incr \geq 1$ **then**
6: **for** $i \leftarrow 1$ to $divider$ **do**
7: Push $(k + i \times incr, c)$ to *pendings*
8: **end for**
9: **end if**
10: $K \leftarrow k, C \leftarrow c$
11: $k \leftarrow k \times multiplier$
12: **return** K, k, C, *pendings*

Algorithm 3 checkFeasibility

Input: ram, mac, $flash$, RAM_{max}, MAC_{max}, $Flash_{max}$
Output: *bool*
1: Start
2: **return** $(ram \leq RAM_{max}) \wedge (flash \leq Flash_{max}) \wedge (mac \leq MAC_{max})$

Algorithm 4 buildModel

Input: First layer 1d CNN filter size k, Repeated block count c, Dataset ds
Output: $model, ram, flash, mac$
1: Start
2: $model \leftarrow$ Generate the model based on k, c, ds using the method shown in Figure 2.
3: $model' \leftarrow$ Optimise the *model* using Integer Quantization {using TensorFlow Lite }
4: $ram, flash, mac \leftarrow$ Profile ($model'$) {Using MLTK library}
5: **return** $model, ram, flash, mac$

below. TinyTNAS is designed to generalize across various types of time-series datasets beyond those specifically mentioned.

– **UCIHAR :** The UCI Human Activity Recognition [20] dataset captures daily activities using a waist-mounted smartphone with inertial sensors. It includes recordings of 6 activities (WALKING, WALKING_UPSTAIRS, WALKING_DOWNSTAIRS, SITTING, STANDING, LAYING) from 30 subjects aged 19-48. The dataset includes 3-axial linear acceleration and 3-axial angular velocity data captured at 50Hz and processed into fixed-width sliding windows of 2.56 seconds with 50% overlap. Features were extracted from both time and frequency domains after noise filtering and gravitational separation using a low-pass filter with a cutoff frequency of 0.3 Hz.

– **PAMAP2 :** The PAMAP2 [19] dataset captures data from 18 physical activities performed by 9 subjects using 3 IMU sensors and a heart rate monitor. For simplicity, we consider 6 activities: WALKING, RUNNING, CYCLING, COMPUTER_WORK, CAR_DRIVING, and ROPE_JUMPING. Data from the IMU attached to the wrist is utilized, with sensors sampling at 100Hz originally. We resample this data to 20Hz and apply a sliding window of 2 seconds with 50% overlap for data processing, focusing on accelerometer and gyroscope readings. This dataset is suitable

Algorithm 5 exploreDepth

Input: Accuracy *acc*, First layer 1d CNN filter size *k*, Repeated block count *c*, RAM_{max}, MAC_{max}, $Flash_{max}$, *pendings*, *search_time*, K, C
Output: *is_continueable*, k, c,
Initialization: Initialize $accs \leftarrow [0]$, $cs \leftarrow [0]$, *is_continueable* ← **False**, $n \leftarrow$ *maximum_possible_repeatable_blocks*

```
1:  Start
2:  for i ← 0 to n do
3:      if has_search_time then
4:          model, ram, flash, mac ← buildModel(k, i, ds)
5:          if checkFeasibility(ram, mac, flash, RAM_max, MAC_max, Flash_max) then
6:              acc ← trainModel(model, ds, epoch)
7:              Append acc to accs
8:              Append i to cs
9:          else
10:             break
11:         end if
12:     else
13:         break
14:     end if
15: end for
16: indx ← index of maximum value in accs
17: if accs[indx] > acc then
18:     c ← cs[indx]
19:     K, k, C, pendings ← updateStatus(k, c, pendings)
20:     is_continueable ← True
21: else if pendings then
22:     k, c ← pop from pendings
23:     is_continueable ← True
24: end if
25: return is_continueable, k, c, K, C
```

for developing algorithms in data processing, segmentation, feature extraction, and activity classification.

- **WISDM :** The WISDM [26] dataset contains accelerometer and gyroscope time-series sensor data collected from a smartphone and smartwatch as 51 test subjects perform 18 activities for 3 minutes each. The raw sensor data, sampled at 20Hz, is gathered from both devices. Specifically, data from the smartwatch's accelerometer and gyroscope are used. Instead of using the recommended 10-second time window, we opt for a 2-second window with 50% overlap to generate sliding window features, aiming to reduce computational overhead. For simplicity, the analysis focuses on six activities: WALKING, JOGGING, TYPING, WRITING, STAIRS, and BRUSHING_TEETH.
- **MIT-BIH :** The MIT-BIH Arrhythmia Database [16] contains 48 half-hour excerpts of two-channel ambulatory ECG recordings from 47 subjects studied between 1975 and 1979. These recordings were digitized at 360 samples per second per channel with 11-bit resolution over a 10 mV range. The database includes approximately 110,000 beats annotated by multiple cardiologists, covering both common and less frequent clinically significant arrhythmias. In our study, we used ECG lead II data resampled to a sampling frequency of 125Hz as input, following the approach described by [11]. Each beat in the dataset is annotated by at least two cardiologists and classified into five categories (N, S, V, F, Q) following the Association for the Advancement of Medical Instrumentation (AAMI) EC57 standard [1]. For preprocessing the ECG beats, we employed the pipeline described by [11].

– **PTB Diagnostic ECG Database :** The PTB Diagnostics dataset [3] (PTB-DED) comprises ECG records from 290 subjects, including 148 diagnosed with MI, 52 healthy controls, and others diagnosed with 7 different diseases. Each record includes ECG signals from 12 leads sampled at 1000Hz. In our study, we focused solely on ECG lead II and analyzed the MI and healthy control categories. For preprocessing the ECG beats, we followed the pipeline detailed by [11].

4.2 Specifications of Generated Architectures

This paper introduces TinyTNAS, a novel hardware-aware multiobjective Neural Architecture Search (NAS) tool tailored for generating architectures under stringent resource constraints. For our experiments, we configured TinyTNAS to operate within the limitations of **20 KB RAM**, **64 KB FLASH** memory, **60K Multiply-Accumulate Operations (MAC)**, and a maximum search time of **10 minutes** per dataset.

The tool was deployed on a desktop system equipped with an **AMD Ryzen 5 Pro 4650G processor, 32 GB of RAM, and a 256 GB SSD**, leveraging **CPU computation** with the absence of a GPU. We uniformly applied these constraints across the aforementioned five distinct datasets, utilizing TinyTNAS to discover optimal architectures. Figure 3 in our study visualizes the architectures identified along with corresponding search times.

4.3 Comparision with State-Of-The-Art (SOTA)

To evaluate the architecture generated by TinyTNAS under specific constraints with the dataset, we conduct full training using the *Adam optimizer* with a learning rate of 0.001. During training, we utilize the *ReduceLROnPlateau* callback with monitoring of maximum validation accuracy, ensuring improved model convergence and efficiency. The best model, based on *maximum validation accuracy*, is saved periodically in a .h5 file for future use every 200 epochs.

For the Human Activity Recognition (HAR) dataset, we benchmark the performance of our optimized architecture against SOTA methods including *CNN-based* [30], *Deep-Conv-LSTM-based* [17], and *LSTM-CNN-based* [28]. Similarly, for ECG-related datasets, comparisons are made with *CNN-based models* [11].

Furthermore, we assess the hardware requirements for potential deployment on MCU environments. To facilitate fair comparisons, we re-implement these SOTA models, training them on the datasets used in this study. Additionally, we optimize these models using *TensorFlow Lite* and profile their resource requirements using the *MLTK* library to ensure suitability for deployment in resource-constraint environments.

Each dataset, along with the architectures produced by TinyTNAS and the re-implemented SOTA models, undergoes latency evaluation after optimization on two widely used low-cost, low-power IoT-enabled MCUs: *ESP32* and *Nano BLE Sense*. Here are their specifications:

ESP32: Dual-core Xtensa 32-bit LX6 microprocessors, 520 KB SRAM, 4 MB flash memory, priced at approximately $5.

Nano BLE Sense: Nordic nRF52840 microcontroller, ARM Cortex-M4F processor, 256 KB RAM, 1 MB flash memory, priced at approximately $30.

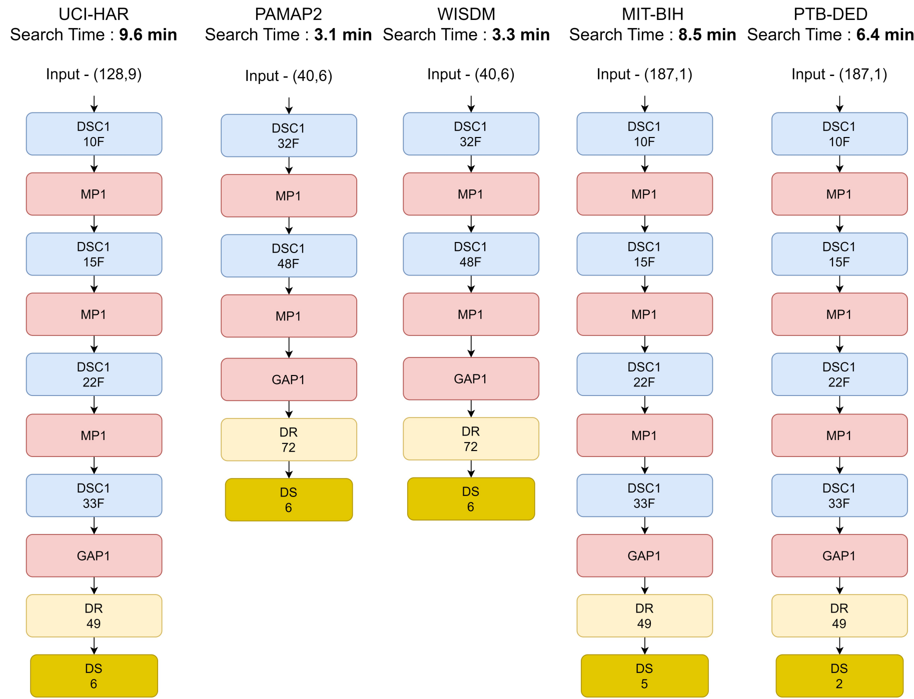

Fig. 3. Architectures Generated by TinyTNAS under Specific Constraints on Various Datasets. Constraints include maximum RAM of 20 KB, maximum FLASH of 64 KB, maximum MAC of 60K, and a maximum search time of 10 minutes. DSC1 denotes Depthwise Separable 1D Convolution with a kernel size of 3 and ReLU activation. MP1 represents Max Pooling 1D with a size of 2. GAP1 indicates Global Average Pooling 1D. DR refers to a Dense Layer with ReLU activation, and DS denotes a Dense Layer with Softmax activation.

We measure the inference latency on each MCU, summarizing the results in Table 1 for UCI HAR, Table 2 for PAMAP2, Table 3 for WISDM, Table 4 for MIT-BIH, and Table 5 for the PTB Diagnostics dataset. It is important to note that the inference time for *Deep-Conv-LSTM-based* models [17] is not available due to high flash memory requirements, rendering the optimized models undeployable on the MCUs. These instances are marked as ND (Not Deployable) in the tables.

4.4 Discussion

Our evaluation highlights TinyTNAS's achievement of state-of-the-art (SOTA) accuracy while significantly reducing resource requirements—RAM, FLASH memory, latency and MAC operations—which are crucial for efficient time series classification tasks on TinyML domain. Given constraints of 20 KB RAM, 64 KB flash memory, 60K MAC operations, and a 10-minute search time on CPU, TinyTNAS sucessfuly generate architectures for given benchmark datasets.

Table 1. Accuracy and Resource Requirements for Deploying Models Generated by TinyTNAS in TinyML Environments: RAM, FLASH, MAC, Inference Time on Select MCUs, and Comparison with State-of-the-Art Methods for Dataset UCI-HAR

Methods	Structure	RAM (kB)	MAC	FLASH (bytes)	Accuracy %	Latency (ms) ESP32	Latency (ms) Nano 33 BLE
[30]	CNN (1d)	53.6	1M	623.9K	92.93	697	2525
[17]	DeepConvLSTM (1d)	124.4	7.6M	1.5M	92.61	ND	ND
[28]	LSTM-CNN (1d)	38.9	1.3M	204.9K	93.21	1338	4617
TinyTNAS	Depthwise Seperable CNN (1d)	**10.8**	**53.1K**	**19.3K**	**93.4**	**13**	**31**

Table 2. Accuracy and Resource Requirements for Deploying Models Generated by TinyTNAS in TinyML Environments: RAM, FLASH, MAC, Inference Time on Select MCUs, and Comparison with State-of-the-Art Methods for Dataset PAMAP2

Methods	Structure	RAM (kB)	MAC	FLASH (bytes)	Accuracy %	Latency (ms) ESP32	Latency (ms) Nano 33 BLE
[30]	CNN (1d)	18.4	217.2K	140.9K	95.66	115	564
[17]	DeepConvLSTM (1d)	34.3	1.8 M	1.3 M	96.44	ND	ND
[28]	LSTM-CNN (1d)	16.4	357.1K	203.3K	95.97	337	1273
TinyTNAS	Depthwise Seperable CNN (1d)	**5.9**	**44.9K**	**15.8K**	**96.7**	**8**	**19**

Table 3. Accuracy and Resource Requirements for Deploying Models Generated by TinyTNAS in TinyML Environments: RAM, FLASH, MAC, Inference Time on Select MCUs, and Comparison with SOTA Methods for Dataset WISDM

Algorithms	Structure	RAM (kB)	MAC	FLASH (bytes)	Accuracy %	Latency (ms) ESP32	Latency (ms) Nano 33 BLE
[30]	CNN (1d)	18.4	217.2K	141.2K	94.91	115	564
[17]	DeepConvLSTM (1d)	34.3	1.8 M	1.3 M	96.1	ND	ND
[28]	LSTM-CNN (1d)	16.4	357.1K	203.3K	96.13	337	1273
TinyTNAS	Depthwise Seperable CNN (1d)	**5.9**	**44.9K**	**15.8K**	**96.5**	**8**	**19**

Table 4. Accuracy and Resource Requirements for Deploying Models Generated by TinyTNAS in TinyML Environments: RAM, FLASH, MAC, Inference Time on Select MCUs, and Comparison with State-of-the-Art Methods for Dataset MIT-BIH

Algorithms	Structure	RAM (kB)	MAC	FLASH (bytes)	Accuracy %	Latency (ms) ESP32	Latency (ms) Nano 33 BLE
[11]	CNN (1d)	80.7	3.6M	230.5K	**98.36**	2393	9710
TinyTNAS	Depthwise Seperable CNN (1d)	**9**	**56.5K**	**19K**	97.4	**13**	**33**

Notably, TinyTNAS sets itself apart by employing a novel optimized grid search methodology instead of traditional reinforcement learning or evolutionary algorithms, optimizing architecture effectively on CPU-based but GPU-free systems within spec-

Table 5. Accuracy and Resource Requirements for Deploying Models Generated by TinyTNAS in TinyML Environments: RAM, FLASH, MAC, Inference Time on Select MCUs, and Comparison with State-of-the-Art Methods for Dataset PTB Diagnostic ECG Database

Algorithms	Structure	RAM (kB)	MAC	FLASH (bytes)	Accuracy %	Latency (ms) ESP32	Latency (ms) Nano 33 BLE
[11]	CNN (1d)	80.7	3.6M	230.5K	**99.24**	2393	9710
TinyTNAS	Depthwise Seperable CNN (1d)	**9**	**56.5K**	**18.8K**	95	**13**	**33**

ified constraints. Detailed comparisons with alternative methods and comprehensive resource profiles are presented in Tables [1,2,3,4,5].

In the case of the UCI HAR dataset, TinyTNAS-generated architectures meet the constraints and outperform other SOTA methods. It achieves a 12*x* reduction in RAM usage, 144*x* reduction in MAC operations, 78*x* reduction in flash memory, and a 149*x* reduction in latency. For the PAMAP2 and WISDM datasets, TinyTNAS achieves a 6*x* reduction in RAM usage, 40*x* reduction in MAC operations, 83*x* flash memory, and 67*x* reduction in latency, while maintaining superior accuracy compared to alternatives.

Moreover, for the PTB dataset, TinyTNAS achieves significant reductions: 9*x* less RAM usage, 64*x* fewer MAC operations, 13*x* lower flash memory requirement, and a 295*x* decrease in latency, with accuracy slightly below 5%. Similarly, for the MIT BIH dataset, TinyTNAS shows impressive reductions: 9*x* less RAM usage, 64*x* fewer MAC operations, 13*x* lower flash memory requirement, and a 295*x* decrease in latency, although its accuracy is slightly below 1%, indicating a minimal trade-off compared to other methods.

5 Conclusion

TinyTNAS represents a pioneering effort in bridging Neural Architecture Search (NAS) with TinyML, specifically targeting time series classification on resource-constrained devices. In this context, we introduce TinyTNAS, a novel hardware-aware multi-objective NAS tool designed specifically for TinyML time series classification. Unlike existing methods, TinyTNAS operates efficiently on CPUs without the need for GPUs, making it accessible and practical for a broader range of applications.

Users can define constraints on RAM, FLASH, and MAC operations to discover optimal neural network architectures that meet these parameters for a given dataset. Additionally, TinyTNAS allows for time-bound searches, ensuring the best possible model among the explored models is found within a user-specified duration. By leveraging an optimized grid search methodology that skips cells in a geometric progression in its explorable dimensions, TinyTNAS achieves state-of-the-art accuracy while drastically reducing RAM, FLASH memory, latency, and MAC operations, all within a very short search time on CPU.

Our comprehensive evaluations across benchmark datasets—UCI HAR, PAMAP2, WISDM, MIT-BIH, and PTB—demonstrate significant performance improvements compared to existing methods. These results underscore the effectiveness of integrating

hardware-aware, multi-objective NAS approaches to efficiently meet stringent TinyML constraints.

Importantly, to the best of our knowledge, this work represents the first comprehensive effort in developing a NAS tool specifically tailored for TinyML time series classification. It incorporates hardware-aware multi-objective optimization with constraints on RAM, MAC operations, FLASH memory, and time-bound searches, operating efficiently on CPUs without the need for GPUs. This sets a new standard in optimizing neural network architectures for AIoT and low-cost, low-power embedded AI applications.

References

1. Association for the Advancement of Medical Instrumentation, et al.: Testing and Reporting Performance Results of Cardiac Rhythm and ST Segment Measurement Algorithms, ANSI/AAMI EC38, vol. 1998 (1998)
2. Banbury, C., Zhou, C., Fedorov, I., Matas, R., Thakker, U., Gope, D., Janapa Reddi, V., Mattina, M., Whatmough, P.: Micronets: Neural network architectures for deploying tinyml applications on commodity microcontrollers. Proceedings of machine learning and systems **3**, 517–532 (2021)
3. Bousseljot, R., Kreiseler, D., Schnabel, A.: Nutzung der ekg-signaldatenbank cardiodat der ptb über das internet (1995)
4. Cai, H., Gan, C., Wang, T., Zhang, Z., Han, S.: Once-for-all: Train one network and specialize it for efficient deployment. arXiv preprint arXiv:1908.09791 (2019)
5. Cai, H., Zhu, L., Han, S.: Proxylessnas: Direct neural architecture search on target task and hardware. arXiv preprint arXiv:1812.00332 (2018)
6. Chauhan, A., Bhattacharyya, S., Vadivel, S.: Dqnas: Neural architecture search using reinforcement learning. arXiv preprint arXiv:2301.06687 (2023)
7. David, R., Duke, J., Jain, A., Janapa Reddi, V., Jeffries, N., Li, J., Kreeger, N., Nappier, I., Natraj, M., Wang, T., et al.: Tensorflow lite micro: Embedded machine learning for tinyml systems. Proceedings of Machine Learning and Systems **3**, 800–811 (2021)
8. Dong, J.D., Cheng, A.C., Juan, D.C., Wei, W., Sun, M.: Dpp-net: Device-aware progressive search for pareto-optimal neural architectures. In: Proceedings of the European conference on computer vision (ECCV). pp. 517–531 (2018)
9. Garavagno, A.M., Ragusa, E., Frisoli, A., Gastaldo, P.: A hardware-aware neural architecture search algorithm targeting low-end microcontrollers. In: 2023 18th Conference on Ph. D Research in Microelectronics and Electronics (PRIME). pp. 281–284. IEEE (2023)
10. He, X., Zhao, K., Chu, X.: Automl: A survey of the state-of-the-art. Knowledge-based systems **212**, 106622 (2021)
11. Kachuee, M., Fazeli, S., Sarrafzadeh, M.: Ecg heartbeat classification: A deep transferable representation. In: 2018 IEEE international conference on healthcare informatics (ICHI). pp. 443–444. IEEE (2018)
12. Lane, N.D., Bhattacharya, S., Georgiev, P., Forlivesi, C., Jiao, L., Qendro, L., Kawsar, F.: Deepx: A software accelerator for low-power deep learning inference on mobile devices. In: 2016 15th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN). pp. 1–12. IEEE (2016)
13. Lin, J., Chen, W.M., Lin, Y., Gan, C., Han, S., et al.: Mcunet: Tiny deep learning on iot devices. Advances in neural information processing systems **33**, 11711–11722 (2020)

14. Liu, C., Zoph, B., Neumann, M., Shlens, J., Hua, W., Li, L.J., Fei-Fei, L., Yuille, A., Huang, J., Murphy, K.: Progressive neural architecture search. In: Proceedings of the European conference on computer vision (ECCV). pp. 19–34 (2018)
15. Liu, H., Simonyan, K., Yang, Y.: Darts: Differentiable architecture search. arXiv preprint arXiv:1806.09055 (2018)
16. Moody, G.B., Mark, R.G.: The impact of the mit-bih arrhythmia database. IEEE engineering in medicine and biology magazine **20**(3), 45–50 (2001)
17. Ordóñez, F.J., Roggen, D.: Deep convolutional and lstm recurrent neural networks for multimodal wearable activity recognition. Sensors **16**(1), 115 (2016)
18. Real, E., Aggarwal, A., Huang, Y., Le, Q.V.: Regularized evolution for image classifier architecture search. In: Proceedings of the aaai conference on artificial intelligence. vol. 33, pp. 4780–4789 (2019)
19. Reiss, A.: PAMAP2 Physical Activity Monitoring. UCI Machine Learning Repository (2012), DOI: https://doi.org/10.24432/C5NW2H
20. Reyes-Ortiz, J., Anguita, D., Ghio, A., Oneto, L., Parra, X.: Human Activity Recognition Using Smartphones. UCI Machine Learning Repository (2012), DOI: https://doi.org/10.24432/C54S4K
21. Saha, B., Samanta, R., Ghosh, S., Roy, R.B.: Bandx: An intelligent iot-band for human activity recognition based on tinyml. In: Proceedings of the 24th International Conference on Distributed Computing and Networking. pp. 284–285 (2023)
22. Saha, B., Samanta, R., Ghosh, S.K., Roy, R.B.: From wrist to world: Harnessing wearable imu sensors and tinyml to enable smart environment interactions. In: Proceedings of the Third International Conference on AI-ML Systems. pp. 1–3 (2023)
23. Saha, B., Samanta, R., Ghosh, S.K., Roy, R.B.: Tinyml-driven on-device personalized human activity recognition and auto-deployment to smart bands. In: Proceedings of the Third International Conference on AI-ML Systems. pp. 1–9 (2023)
24. Saha, B., Samanta, R., Roy, R.B., Chakraborty, C., Ghosh, S.K.: Personalized human activity recognition: Real-time on-device training and inference. IEEE Consumer Electronics Magazine (2024)
25. Warden, P., Situnayake, D.: Tinyml: Machine learning with tensorflow lite on arduino and ultra-low-power microcontrollers. O'Reilly Media (2019)
26. Weiss, G.: WISDM Smartphone and Smartwatch Activity and Biometrics Dataset . UCI Machine Learning Repository (2019), DOI: https://doi.org/10.24432/C5HK59
27. Wu, B., Dai, X., Zhang, P., Wang, Y., Sun, F., Wu, Y., Tian, Y., Vajda, P., Jia, Y., Keutzer, K.: Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10734–10742 (2019)
28. Xia, K., Huang, J., Wang, H.: Lstm-cnn architecture for human activity recognition. IEEE Access **8**, 56855–56866 (2020)
29. Xie, X., Sun, Y., Liu, Y., Zhang, M., Tan, K.C.: Architecture augmentation for performance predictor via graph isomorphism. IEEE Transactions on Cybernetics **54**(3), 1828–1840 (2023)
30. Yang, J., Nguyen, M.N., San, P.P., Li, X., Krishnaswamy, S.: Deep convolutional neural networks on multichannel time series for human activity recognition. In: Ijcai. vol. 15, pp. 3995–4001. Buenos Aires, Argentina (2015)
31. Zoph, B., Le, Q.V.: Neural architecture search with reinforcement learning. arXiv preprint arXiv:1611.01578 (2016)
32. Zoph, B., Vasudevan, V., Shlens, J., Le, Q.V.: Learning transferable architectures for scalable image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8697–8710 (2018)