

REDS: Resource-Efficient Deep Subnetworks for Dynamic Resource Constraints

Francesco Corti¹, Balz Maag², Joachim Schauer³, Ulrich Pferschy⁴, and Olga Saukh^{1,5}

¹Graz University of Technology, Austria

²ABB Research, Switzerland

³University of Applied Sciences, FH Joanneum, Austria

⁴University of Graz, Austria

⁵CSH Vienna, Austria

{francesco.corti,saukh}@tugraz.at, balz.maag@ch.abb.com,
joachim.schauer@fh-joanneum.at, ulrich.pferschy@uni-graz.at

Abstract

Deep models deployed on edge devices frequently encounter resource variability, which arises from fluctuating energy levels, timing constraints, or prioritization of other critical tasks within the system. State-of-the-art machine learning pipelines generate resource-agnostic models, not capable to adapt at runtime. In this work we introduce Resource-Efficient Deep Subnetworks (REDS) to tackle model adaptation to variable resources. In contrast to the state-of-the-art, REDS use structured sparsity constructively by exploiting permutation invariance of neurons, which allows for hardware-specific optimizations. Specifically, REDS achieve computational efficiency by (1) skipping sequential computational blocks identified by a novel iterative knapsack optimizer, and (2) leveraging simple math to re-arrange the order of operations in REDS computational graph to take advantage of the data cache. REDS support conventional deep networks frequently deployed on the edge and provide computational benefits even for small and simple networks. We evaluate REDS on six benchmark architectures trained on the GOOGLE SPEECH COMMANDS, FMNIST and CIFAR10 datasets, and test on four off-the-shelf mobile and embedded hardware platforms. We provide a theoretical result and empirical evidence for REDS outstanding performance in terms of submodels' test set accuracy, and demonstrate an adaptation time in response to dynamic resource constraints of under $40\mu\text{s}$, utilizing a 2-layer fully-connected network on Arduino Nano 33 BLE Sense.

1 Introduction

Data processing pipelines in edge devices increasingly rely on deep learning models to identify patterns and extract insights from multimodal IoT data. Examples include predictive maintenance in industrial automation, object identification and tracking in smart camera systems (Qu et al., 2022), activity and healthcare trackers in mobile (Ravi et al., 2016), wearable and hearable applications (Maag et al., 2017; Sabry et al., 2022). In all these systems, deep models are deployed and run along with other tasks, under constraints and priorities dictated by the current context and available resources, including storage, CPU time, energy and bandwidth.

A large body of work explore different techniques to optimize and compress deep models without hurting accuracy and generalization abilities, while accelerating their execution in software (Boehm et al., 2018) and in hardware (Jouppi et al., 2017). Network pruning and quantization (Han et al., 2016) have become part of standard deep learning deployment pipelines, *e.g.*, TFLMicro (David et al., 2021) and




Figure 1: An example of layer slicing in a toy dense model. The weight tensor dimensions are shown in brackets (activation functions omitted). We slice off two neurons in the first hidden layer of the toy network with 4 neurons, 3 inputs and 2 outputs. The matrices W_1 and W_2 store the weights along all connections, and the respective columns and rows in W_1 and W_2 get eliminated by layer slicing. This breaks the contiguous memory layout and the memory arrangement of W_1 , yet not W_2 . A simple transpose of W_1 and a change of the multiply order preserve contiguous memory of W_1 after slicing.

TensorRT (Vanhoder, 2016), to enable deep learning on severely constrained embedded hardware operated by low-power microcontrollers with only a few kB or RAM. One drawback of compile-time optimizations is that the resulting models are *resource-agnostic*. They thus yield suboptimal performance in many interesting applications where resource availability depends on different factors such as available energy, task priority and timing constraints. Another drawback of one-shot model compression techniques, is that these are applied to the *whole model*, making exploration of different options for a resource-aware on-device model reconfiguration challenging. Both problems are described in detail below.

Dynamic resource constraints. Many interesting applications can make use of resource-aware deep models, *i.e.*, models that can adapt their execution to available computational resources and time constraints. For example, camera image processing by a drone or a car may depend on the respective speed (Qu et al., 2022). Processing interesting and relevant data can justify using more energy and computational time than when running regular environment scans (Gherman et al., 2021b). A naive solution to address dynamic resource constraints is to store several independent deep models and switch between them as resource availability and task priorities change. The drawback of this approach is both the increased memory consumption to store these independent models which does not scale, and the overhead of switching between models at runtime, *e.g.*, by loading these from flash to RAM and reallocating the necessary tensors.

Several approaches have been proposed to adapt a model to dynamic resource constraints. These can be classified into the methods that implement early exit predictions at a cost of a reduced accuracy, *e.g.*, BudgetRNN (Kannan and Hoffmann, 2021) and ASVN (Bambusi et al., 2021), and methods that maintain a subnetwork structure using weight sharing, with each subnetwork being more efficient yet possibly less accurate, *e.g.*, DRESS (Qu et al., 2022) and NestDNN (Fang et al., 2018). This work falls into the latter category, yet makes use of structured sparsity constructively to allow for additional flexibility and hardware support on typical IoT devices.

Deep learning support on IoT devices. Modern IoT hardware often provides support for running deep learning models. This includes support for floating point instructions, dedicated instructions for frequent operations, such as MLA, FMA, and their vector versions. Deep model frameworks for a specific platform make use of the available hardware-specific features to provide maximum speedup. However, the methods that introduce subnetwork structures to a model can not rely on the toolchain support to optimize the execution of each subnetwork. For examples, fine-grained weight pruning may lead to accelerated execution due to an abundance of zeros in the weight matrices and sparse filters, yet unconstrained locations of these zeros present a difficulty to make use of unstructured sparsity. Moreover, there are specific algorithms designed to best handle different sparsity levels (Hoefler et al., 2021), that may or may not enjoy the available hardware

support.

To overcome the issues, NestDNN (Fang et al., 2018) prunes convolutional filters; these have to be paged in or out when switching from one multi-capacity model to another. DRESS (Qu et al., 2022) relies on a specialized hardware to ensure the applied sparsity pattern translates into computation efficiency. Hardware accelerators may, however, appear too power-hungry and expensive for battery-powered IoT devices. Specialized hardware is also inflexible as the deep learning technology advances fast and hardware modernizations are costly.

Contributions. We present the design of Resource-Efficient Deep Subnetworks (REDS) featuring a nested submodel structure to address dynamic resource constraints on mobile and IoT devices. Similarly to DRESS and NestDNN, REDS uses structured sparsity. Differently though, it leverages permutation invariance of neurons to keep subnetwork weight tensors in contiguous memory regions. This makes the description of the subnetwork structure embarrassingly simple, reduces REDS adaptation time down to nearly zero, and allows for further hardware-specific optimizations. Furthermore, we introduce a novel method to design the nested subnetwork structure by formulating and solving an iterative knapsack problem, provide theoretical guarantees and empirical evidence of outstanding solution performance. Our contributions are:

- We present REDS that adapt to dynamically changing resource constraints at near-zero time. REDS make use of the permutation invariance of neurons to enable hardware-specific optimizations (Sec. 3).
- We present a novel hardware-aware method to convert a pre-trained model to REDS. We formulate the optimization problem as iterative knapsack, present theoretical analysis and provide empirical evidence of the solution effectiveness, especially in the low-data regime (Sec. 3 and Sec. 4).
- For resource-constrained devices that use data caches, REDS provide a compile-time optimization to ensure all subnetwork weights reside in contiguous memory (Sec. 5).
- REDS are evaluated on six benchmark architectures from (Zhang et al., 2017) trained on GOOGLE SPEECH COMMANDS (Warden, 2018), FMNIST (Xiao et al., 2017) and CIFAR10 (Krizhevsky et al., 2009) and tested on four mobile and IoT devices (Sec. 6). Our code is publicly available.¹

Sec. 7 summarizes related works and Sec. 8 concludes this paper. In the next section we explain REDS on a sample toy deep network.

2 A Toy Example

We first illustrate advantages of REDS on a toy fully-connected network depicted in Fig. 1 and featuring an input layer with 3 neurons, two hidden layers with 4 hidden units each, and an output layer with 2 neurons. The three dense weight matrices $W_1^{3 \times 4}$, $W_2^{4 \times 4}$ and $W_3^{4 \times 2}$ connect neurons in consecutive layers. These matrices are stored in memory as unfolded one-dimensional arrays. We assume the row-major format is used to store and access model weights in memory, which is the standard choice on most hardware architectures. Given a trained and optimized model, weight matrices are mapped to continuous memory regions, allowing for a straightforward use of vector instructions to speed up on-device inference.

REDS organizes a model into a set of *nested submodels*, *i.e.*, the active weights of a child subnetwork are fully contained in its parent subnetwork. To enable this structure, we slice each parent subnetwork into an *active* and an *inactive* part, where the active part shapes the child subnetwork. REDS makes use of structured sparsity, *i.e.*, model slicing occurs at the level of individual neurons and convolutional filters. Slicing off two neurons in the first hidden layer in Fig. 1 leads to the removal of two columns and two rows in the weight matrices W_1 and W_2 respectively.

First, notice that pruning a neuron in a layer has a *local scope* and affects only the number of adjacent connections, *i.e.*, the orange part of the network conveyed by matrices W_1 and W_2 in our example. This

¹<https://github.com/FraCorti/REDS.git>

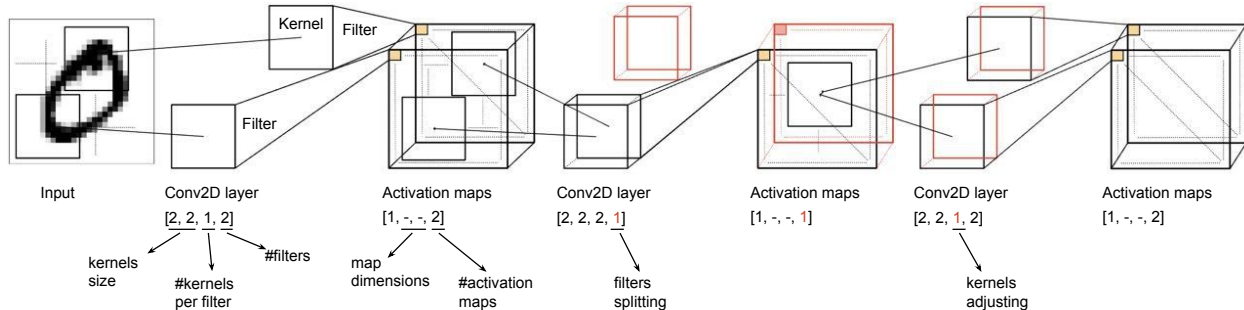


Figure 2: Filter removal in a convolution layer; the dimensions of the weight tensors are shown in brackets (pooling and activation functions are omitted for brevity). The dimensions of the sliced filters, the sequence of the reorganized kernels, and the adjustments to the weight tensors are highlighted in red.

also holds for the convolutional network exemplified in Fig. 2 and in (Fang et al., 2018) Fig. 2, although not discussed in the text.

Secondly, even though any neuron can be removed from a layer, we claim we can always re-order neurons to have a group of active neurons followed by inactive neurons. This is possible due to the *permutation invariance* phenomenon of neural networks (Entezari et al., 2021; Jordan et al., 2022). For instance in Fig. 1, the first and the last neurons in the first layer can be permuted by swapping both the first and last columns in W_1 and rows in W_2 , respectively. The permutation does not change the function of the network, *i.e.*, network predictions remain identical. This is an important property which allows optimizing the memory layout to keep subnetwork weights in contiguous memory. For the same reason, considering contiguous groups of active and inactive neurons or filters does not limit the expressiveness of trained networks in the REDS format. The above observations make the subnetwork layout very simple. We discuss this in Sec. 3.

Thirdly, REDS achieve *near-zero adaptation overhead*: Only the widths of the layers have to be updated to switch to a different model. No overhead is incurred at inference time.

Finally, pruning a neuron in one layer may remove more multiply-accumulate operations (MACs) than in another layer. For example, removing one neuron in the first hidden layer yields more MAC reduction than removing a neuron in the second hidden layer in our toy example due to a different number of connections. Also the contribution of these neurons to model accuracy may be different. Related research suggests several importance measures for individual weights, neurons, and convolutional filters (Frantar and Alistarh, 2022; Hoeffler et al., 2021; Molchanov et al., 2019; Shen et al., 2022). These can be used to optimize a network for performance while keeping essential elements. Our problem is different: We build a nested structure. This requires iterating over the subnetworks. DRESS starts from the largest network. NestDNN implements a top-down pruning and a bottom-up freeze-and-regrow algorithm. We formulate the optimization problem for a parent-child subnetwork pair as a variant of the knapsack problem, similarly to Shen et al. (2022). To extend the solution to multiple nested subnetworks, we generalize the approach to an iterative knapsack problem. In contrast to the previous works, our theoretical findings (Appendix A.2) show that a knapsack method applied to the smallest subnetwork first and letting it grow, *i.e.*, the *bottom-up approach*, is more effective than starting from the largest subnetwork and pruning it down to the smallest one. We empirically show this in Sec. 4.

In the following sections we focus on specific contributions and design choices for REDS on the way to a fully-functioning framework. Sec. 3 discusses a pipeline to convert a model to REDS and presents a hardware-aware iterative knapsack method to choose the model slicing points. Sec. 5 revisits the memory layout optimization for cache and provides an empirical evaluation of the achieved gain on embedded hardware.

3 Resource-Efficient Deep Subnets

Before describing the pipeline how we build REDS, we first argue why layer slicing by chopping off a *contiguous* sequence of computational units, such as neurons or convolutional filters, doesn't constraint the expressiveness of REDS in any way, *i.e.*, these computational units in all but the last layer of the network can be pruned arbitrary, but can be reordered to form a contiguous sequence without changing the network function. We then describe REDS construction pipeline. In the heart of it is the novel method to decide on the choice of the subnetwork structure of REDS based on the knapsack problem. Finally, we discuss how to fine-tune the models in REDS using few-shot learning to recover lost accuracy due to the applied pruning of computational units.

3.1 Permutation invariance

The optimization landscape of neural networks contains an abundance of minima and saddle points as a result of numerous symmetries. One important class are permutation symmetries, *i.e.*, the neurons or channels in the same layer can be randomly permuted and, if the incoming and outgoing connections are adjusted accordingly, form structurally different, yet functionally identical networks. An example is given in Sec. 2. A layer with n neurons has $n!$ permutations. A deep network with l such layers has $\prod_l n!$ functionally identical solutions. For example, Ainsworth et al. (2022) calculated that ResNet50 contains $10^{55'109}$ permutation symmetries.

We leverage permutation invariance of neural networks to permute the units of a layer in descending order based on their importance scores. This operation doesn't change the subnetwork function, but ensures that the nested subnetworks keep the most important units of each layer as part of their architecture, thus simplifying each subnetwork architecture search. To preserve the correct feature extraction process, the permutation operation is performed layerwise and the structural elements to permute depend on the layer type. In a standard convolution layer, for example, the weights are stored as a four-dimensional tensor (see Fig. 2). When the convolutional filters are permuted, the kernels in the subsequent layer are similarly reordered to maintain consistent input channel and kernel sequence during the forward pass. This permutation process involves manipulating the original weight tensors, initially unfolding the layer's tensor into a series of three-dimensional tensors—representing the convolutional filters—and subsequently unrolling each filter's tensor into a set of two-dimensional tensors, corresponding to the filter's kernels.

3.2 Importance and cost measures

How do we identify key neurons or convolutional filters essential for constructing a high-precision subnetwork? Several studies employ performance scores to assess the significance of weights, neurons, and channels within a neural network. NestDNN (Fang et al., 2018) ranks each convolutional filter by calculating the L2 distance between feature maps of the same class and those of different classes. This method uses all the training data, which is often unavailable due to resource constraints when the model is used in production. Magnitude-based pruning methods, that use the magnitude of the weights to estimate their importance (Cai et al., 2020; Han et al., 2015; Zhu and Gupta, 2017), provide a computationally efficient way to rank the model's computational units. However, numerous studies (Molchanov et al., 2019; Mozer and Smolensky, 1988) have reported that the magnitude does not necessarily reflect the importance of a unit in terms of its contribution to the loss reduction.

In REDS, we adopt an efficient approach to score each model's computational units. This approach can be used on hardware platforms, which support on-device training through gradient descent. We leverage the works by Molchanov et al. (2019) and Shen et al. (2022) to compute for each of the encoder's weight i the importance as $I_i = |g_i \gamma_i|$, where g_i is the sum of the accumulated gradients computed from backpropagation (Rumelhart et al., 1986) and γ_i is the scalar value of weight i . The I metric approximates the squared difference of the loss when a weight is removed (Molchanov et al., 2019). In REDS the gradients are computed on a predefined number of minibatches of the training data, which equals 100. This number was determined through a grid search exploration. For each computational unit c , such as a convolutional filter or a fully-connected neuron,

let \mathcal{W}_c be the set of its weights. Then I_c is computed as the grouped sum of the importance scores over all these weights:

$$I_c = \sum_{i \in \mathcal{W}_c} I_i = \sum_{i \in \mathcal{W}_c} |g_i \gamma_i|. \quad (1)$$

Each computational unit is characterized by an importance score, denoted as I_c , and a computational cost is defined in terms of model latency. We use the number of multiply-accumulate operations (MACs) as the latency predictor for each unit on an MCUs. This choice is empirically validated by [Liberis et al. \(2021\)](#) for typical IoT devices.

3.3 Iterative knapsack problem

This section describes the novel method to design REDS subnetwork structure by formulating and solving a variant of an iterative knapsack problem. The discussion in main paper is limited to the simplest case, while the generalization for a more complex model architecture and the theoretical analysis of the problem are moved to the appendix ([Appendix A.1](#) and [Appendix A.2](#)) for presentation clarity.

3.3.1 Problem formulation

Given a pre-trained model, REDS identifies how to best slice the weight tensors by formulating the problem as an iterative knapsack problem with k stages: the items included in a knapsack with capacity c have to be included in all later stages, *i.e.*, knapsacks with larger capacities. This resembles the setting of the incremental knapsack problem treated *e.g.*, in ([Della Croce et al., 2019](#); [Faenza et al., 2023](#)), although with a different objective function. In REDS we solve an iterative knapsack problem with as many stages as the number of subnetworks we wish to have. The items correspond to all the computational units that compose the model encoder architecture, and each layer of the architecture induces a set of items that behave similarly with respect to their computational cost. Given a list of MAC values obtained by multiplying the computational cost of the full model (100% MACs) by some predefined percentages (in our experiments 75%, 50% and 25%), we give two heuristic algorithms that solve the corresponding iterative knapsack problem and hence find subnetwork architectures, *i.e.*, slicing points for each layer and each subnetwork, that satisfy these capacity constraints while maximizing I . Both heuristics are based on solving several knapsack problems. These are formulated as integer programs which we then solve by using a mixed integer programming solver. In our experiments Google OR-Tools ([Perron and Furnon](#)) is used for modelling the problems and Gurobi ([Gurobi Optimization, LLC, 2023](#)) for solving them. Let C_{MACs} denote the maximum number of MACs in a subnetwork. Then the knapsack problem is formulated as follows:

$$\begin{aligned} \max \quad & \sum_{l=1}^L \sum_{i=1}^{u_l} x_{il} \cdot I_{il} \\ \text{s.t.} \quad & \sum_{l=1}^L \sum_{i=1}^{u_l} x_{il} \cdot MACs_{il} \leq C_{MACs}, \\ & x_{il} \in \{0, 1\}, \quad \forall l \in \{1, \dots, L\}, \quad i \in \{1, \dots, u_l\}, \\ & I_{i1} \geq I_{i2} \geq \dots \geq I_{il}, \end{aligned} \quad (2)$$

where x_{il} is a binary decision variable taking value 1 if an item (*i.e.*, a computational unit) i from layer l is selected and 0 otherwise; I_{il} is the importance score of item i from layer l (denoted as I_c in [Eq.1](#)); $MACs_{il}$ is the number of MACs of item i from layer l ; L is the number of layers in the model and u_l is the number of items in layer l . The above knapsack problem formulation is for simple dense and convolutional networks. In [Appendix A.1](#), we present a non-trivial generalization of our knapsack framework to depth-wise convolutions (DS-CNN), part of modern architectures such as multiple versions of MobileNet ([Howard et al., 2017](#); [Sandler et al., 2018](#)).




Figure 3: REDS fine-tuning is inspired by DRESS (Qu et al., 2022). The parameters $\pi_{i=1}^N$ ensure the contribution of individual models to the loss aligns with the fraction of the shared weights.

3.3.2 Bottom-up (BU) and top-down (TD) heuristics

REDS examines two heuristics for solving the iterative knapsack problem named bottom-up (BU) and top-down (TD). The former iteratively calculates the subnetwork architectures by considering the tightest constraints for the smallest subnetwork first. Once a solution is found by the knapsack solver, these units are frozen, *i.e.*, they are now part of all nested subnetworks, and the second smallest subnetwork is now being computed by the solver. The latter top-down method determines the solutions by considering the weakest constraints first and then iteratively searching the architectures for increasingly smaller subnetworks. Related literature, including DRESS and NestDNN use a variant of the top-down approach. Our theoretical analysis of the iterative knapsack problem (based on the classical 0-1 knapsack problem) in Appendix A.2 shows that the bottom-up approach promises a better worst-case performance. In particular, we prove that the solution found by the two-stage bottom-up iterative knapsack heuristic is not worse than $\frac{2}{3} \cdot Opt$, where Opt is the optimal solution of the knapsack with larger capacity, and that this bound is tight. We then show that the tight bound for the top-down iterative knapsack is $\frac{1}{2} \cdot Opt$, where Opt in this case is the optimal solution of the knapsack with smaller capacity. Since our generalized problem suited for DS-CNN architectures has the classical 0-1 knapsack as its core problem, we believe that a similar result is valid for this case as well.

3.4 REDS fine-tuning

Fine-tuning models after pruning is common practice to recover model accuracy. The REDS subnetworks found by the BU heuristic algorithm are now fine-tuned to recover their full accuracy. The fine-tuning process draws inspiration from the subnetwork training procedure of Once-for-All (Cai et al., 2020) and DRESS (Qu et al., 2022). In the former, a pre-trained model is iteratively fine-tuned along its structured dimensions to enforce and maintain high accuracies in all of its 10^{19} subnetworks. After fine-tuning, a hardware-specialized static subnetwork is selected and deployed in a real-world environment. In the latter, multiple subnetworks are sampled based on the magnitude of the weights. They are fine-tuned in parallel with a weighted loss function to enforce high accuracy by applying for each subnetwork a binary mask to the inputs.

Similarly to DRESS, all REDS subnetworks are fine-tuned simultaneously. Differently, for each subnetwork we slice each layer to construct a subnetwork architecture by storing only the ID of the last active neuron on the child subnetwork called the slicing point. During fine-tuning, sketched in Fig. 3, the weights of each subnetwork i are reused by all lower-capacity subnetworks $\{i+k\}_{k=1}^N$ by dynamically creating a tensor slice for each layer, *i.e.*, a tensor object that points to the original weight tensor with a specific predefined number of its computational units. Similarly to DRESS, we balance the contribution of the loss of each individual model with parameters $\{\pi_i\}_{i=1}^N$ equal to the percentage of weights used in the encoder part of the model by the subnetwork i .

MACs	Small (S) - Accuracy 94.96				Large (L) - Accuracy 95.1			
	Scratch	Knapsack BU	Knapsack TD	REDS training	Scratch	Knapsack BU	Knapsack TD	REDS training
100%	93.83±0.22	93.38±0.45	93.34±0.21	93.19 ±0.18	94.87 ±0.33	94.46±0.08	94.35 ±0.21	94.05 ±0.14
75%	93.37±0.34	93.18±0.21	93.03±0.26	92.85 ±0.47	94.27 ±0.08	94.32±0.13	94.18 ±0.17	93.76 ±0.02
50%	93.41±0.67	92.12±0.24	92.26±0.50	91.50 ±0.10	94.11 ±0.26	94.17±0.01	94.00 ±0.25	93.08 ±0.06
25%	91.46±0.80	89.64±0.76	88.59±1.69	85.14 ±1.07	93.80 ±0.14	93.20±0.19	93.17 ±0.73	92.11 ±0.42

Table 1: Test set accuracy [%] of training S and L depth-wise separable convolutional architectures (DS-CNN) from (Zhang et al., 2017): training a network of each size from scratch ("Scratch"), conversion from a pre-trained network using two knapsack versions ("Knapsack BU" and "Knapsack TD"), and training REDS structure from scratch ("REDS training"). Reported results from three independent runs. The accuracy of each 100% network reported by Zhang et al. (2017) is listed in the header row.

4 REDS Performance

4.1 Evaluation setup

We evaluate the performance of REDS on three datasets: GOOGLE SPEECH COMMANDS (Warden, 2018), FMNIST (Xiao et al., 2017) and CIFAR10 (Krizhevsky et al., 2009). In the GOOGLE SPEECH COMMANDS dataset for keyword spotting, models are trained to classify 10 most common keywords, along with "silence" and "unknown". During data preprocessing, 36k samples undergo conversion to a set of frequency-domain spectral coefficients using the MFCC feature extractor. Both CIFAR10 and FMNIST are image classification datasets, each comprising 60k samples across 10 classes. In CIFAR10, the samples are RGB images, whereas in FMNIST, they are grayscale images. Hyper-parameter values and training details are summarized in Appendix B. The main paper presents selected results for DS-CNN architectures (S and L) on GOOGLE SPEECH COMMANDS. Additional results, also for other architectures and datasets, are presented in Appendix C.

REDS is evaluated on three pre-trained GOOGLE SPEECH COMMANDS architectures (DNN, CNN and DS-CNN) of two sizes (S and L) each introduced in (Zhang et al., 2017). For CIFAR10 and FMNIST REDS is tested on two pre-trained DS-CNN size S architectures. DNN is a 2-layer fully-connected architecture with layer width 144 (S) and 436 (L). CNN is a simple architecture with two convolutional layers followed by 3 fully-connected layers. The widths of convolutional layers are 28 and 30 for the S architecture and 60 and 76 for the L architecture, respectively. DS-CNN is composed of a standard convolutional layer followed by several blocks of depth-wise and point-wise convolutional layers. There are 4 blocks for the network of size S, and 5 blocks for the network of size L with a layer width of 64 and 276 respectively. Each convolutional layer in the CNN and DS-CNN architectures is followed by a batchnorm layer (Ioffe and Szegedy, 2015).

4.2 REDS empirical evaluation

The REDS structure used in the main paper comprises four nested subnetworks with 25%, 50%, 75% and 100% of MACs. The results for a higher number of nested models are reported in Appendix C.2. Having a larger number of subnetworks does not degrade the accuracy of REDS submodels. REDS also support different data domains (FMNIST and CIFAR10) without degrading accuracy compared to training from scratch. In contrast to the state-of-the-art optimizers for microcontrollers, such as μ NAS (Liberis et al., 2021), REDS demonstrate a faster solution search time, taking only a few minutes as opposed to 3 days by μ NAS.

The effectiveness of REDS is shown in Table 1. Training from scratch shows the accuracy achieved by S and L DS-CNN architectures when trained in isolation, *i.e.*, not as part of the REDS nested structure. The next two columns show the resulting REDS performance when the solution is found by the BU and TD heuristics, respectively. Finally, the fourth column reports REDS test set accuracy achieved by each REDS subnetwork trained from scratch, yet following the structure obtained by the BU heuristics.

Fig. 5 visualizes differences between REDS structures obtained with the BU and TD heuristics when using the same initially pre-trained model. Colorful bars show the number of filters kept by each nested model in each layer by BU. Overlaid black bars correspond to the number of kept TD filters for the same

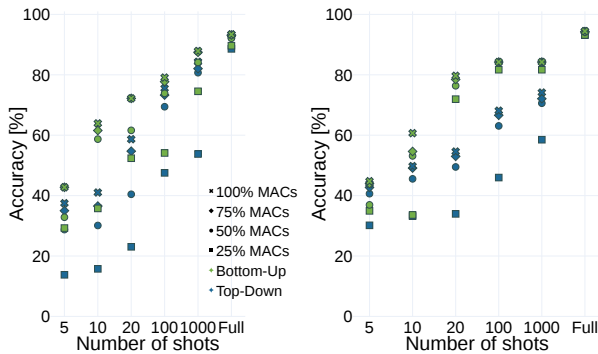


Figure 4: DS-CNN S and L few-shot fine-tuning on GOOGLE SPEECH COMMANDS (Warden, 2018). The subnetworks obtained from the BU Knapsack formulation exhibit faster recovery in accuracy compared to the TD Knapsack ones.

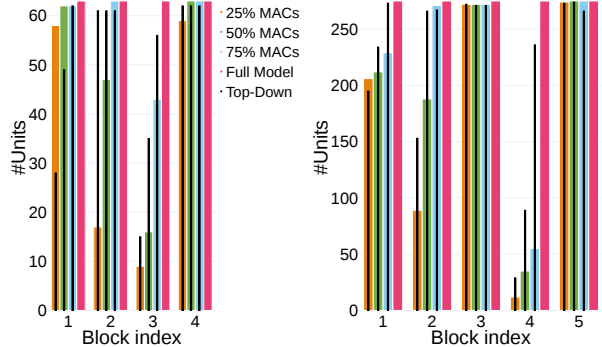


Figure 5: Analysis of the structured sparsity of subnetworks obtained by the knapsack BU and TD heuristics (Sec. 3.3). Left to right: DS-CNN S and DS-CNN L. The slicing point of each subnetwork is visualized with a different color. The results of the TD heuristic are reported in black.

layer and number of MACs. The plot reveals consistent differences in each layer as we move from parent to child submodels. The distribution of computational units across subnetworks in each layer depends on the model architecture and dataset.

We observe that the accuracies of each submodel for the same percentage of MACs are similar, even for small S architectures with only 25% of MACs. The accuracies of the full models, highlighted in bold in the top row are computed using pre-trained models from (Zhang et al., 2017) and closely match our experimental results. The performance difference between BU and TD heuristics are largely insignificant due to the fine-tuning step with extensive training data. Even though the models are overparameterized and in all cases achieve high performance, low-capacity subnetworks S sliced by the bottom-up approach yield minor performance gains.

However in the low-data regime, where only few samples per class are available for fine-tuning, the performance differences between BU and TD heuristics are remarkable. Fig. 4 shows the performance when few-shot learning is applied to fine-tune REDS structures. Our empirical findings confirm a consistently superior performance of the BU heuristic compared to the TD alternative. The differences diminish as more samples are used for fine-tuning, similarly to the effects observed by Entezari et al. (2023). There is barely any difference if full-finetuning is applied using all available training data.

5 REDS Optimization for Caches

Embedded ML frameworks like Tensorflow Lite typically store model weight matrices in row-major order. This means, that each row of a weight matrix is stored contiguously in memory. Without the use of sub-networks these weights are also accessed in a contiguous fashion. However, when using subnetworks for model inference, some neurons and their corresponding weights are omitted, resulting in non-contiguous memory access. This effect is illustrated in Figure 1. We now show how by simply adapting the computational graph at compile time, we are able to optimize the computation of REDS subnetworks for devices with a cache memory architecture.

5.1 Row-major and column-major stores

The calculation of a fully-connected neural network layer during the forward pass is a matrix multiplication $H = \sigma(X_{[m \times b]}^T \cdot W_{[m \times n]})$, where x is an input matrix of shape $m \times b$ (input samples \times batch size), W the

	float	uint8	uint16	uint32
Speed-up	12%	58.6%	54.5%	47.4%
Cache-Hit-Rate Baseline	99.1%	90.9%	91.7%	90.9%
Cache-Hit-Rate Optimized	99.7%	99.4%	98.9%	97.6%

Table 2: Speed-up and cache-hit rate for different matrix data types. There is no notable difference across the different splits 25, 50, 75 and 100%.

layer weights matrix of shape $m \times n$ (input samples \times number of neurons) and σ the activation function. For simplicity we omit the typically used bias term. During matrix multiplication, like it is for instance implemented in Tensorflow Lite, an inner loop multiplies and sums each element of row x_i of X^T with column w_j of W . This column-wise access of row-major ordered weights $w \in W$ may lead to consecutive reads from memory addresses, which are not located close-by depending on the size of the matrix and the use of the REDS subnetworks. However, this can easily be circumvented by changing the matrix multiplication involving a transposed W , *i.e.*, $\sigma(X^T \cdot W) = \sigma((W^T \cdot X)^T)$, which ensures contiguous memory accesses of the weights. Note that the additional transposition of the resulting matrix introduces no additional computations if the input matrix is a single sample, *i.e.*, X has size $[m \times 1]$. Both $X_{[m \times 1]}$ and $X_{[m \times 1]}^T$ are stored equally in memory and, thus, the transposition is redundant in this case. Let us consider the simple example of $X_{[2 \times 2]}$ and $W_{[2 \times 3]}$ both stored in row-major mode. In the basic matrix multiplication cases $X^T \cdot W$, the order of the relative memory access locations of weights w is $0 \rightarrow 3 \rightarrow 1 \rightarrow 4 \rightarrow 2 \rightarrow 5 \rightarrow 0 \rightarrow 3 \rightarrow \dots$. In the optimized cases $(W^T \cdot X)^T$ where W^T is also stored in column-major form, the order of weights accesses is $0 \rightarrow 1 \rightarrow 0 \rightarrow 1 \rightarrow 0 \rightarrow 1 \rightarrow 2 \rightarrow 3 \rightarrow \dots$. Depending on the memory and cache setup of a device, this highlights the potential for exploiting cache systems by simply changing the matrix multiplications to the optimized transposed form.

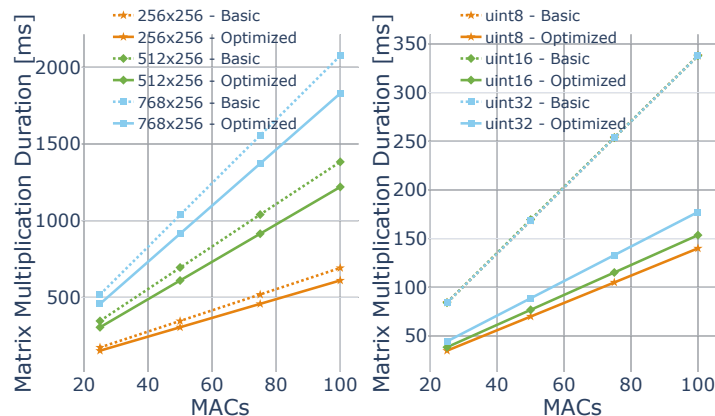


Figure 6: Duration of the matrix multiplication using differently sized weight matrices of floats (left) and different unsigned integer precisions for a 512 x 256 weight matrix (right) at different weight matrix slicing points. The cache optimization shows a clear computational time reduction for all the different test scenarios.

5.2 Optimizing REDS computational graph

The optimization proposed above can be implemented by assuring that each matrix multiplication within a computational graph of each submodel follows the for cache optimized access pattern and the corresponding weight matrices are stored in column-major order in memory. We now show the effect of using this optimized computational graph by benchmarking different matrix multiplications. To this end, we use a Raspberry Pi Pico board, which features a RP2040 chip based on a dual-core Arm Cortex-M0+ processor architecture with

264kB RAM and 16MB off-chip flash memory (RP2040, 2023). The chip also features execute-in-place (XIP) support, such that the flash memory can be treated like internal memory. These accesses to flash are cached with a 16kB, two way set-associative cache with 8 byte wide cache lines.

Fig. 6 shows the cache effect on the duration of matrix multiplications when we use the optimized computation graph compared to the basic one, *i.e.*, the default way where W is stored in row-major mode and we calculate $X^T \cdot W$. We benchmark the matrix multiplication at different splits and use 25%, 50%, 75% and 100% of the neurons in W . In Fig. 6 (left) we show the duration of a matrix multiplication with differently shaped weights W times an input matrix X of shape 256×4 , where the weights and inputs are 32 Bit wide floating points. Similarly, Fig. 6 (right) shows the duration of a weights matrix multiplication of size 512×256 with the same X , but now the weights and inputs are unsigned integers with different widths, *i.e.*, 8, 16 and 32 Bit. We observe in all cases a clear speed-up in matrix multiplication when comparing the optimized and the basic computation approaches, which is also evident from Table 2. For all test cases the cache-hit rate is above 97% when using optimized matrix multiplications. Additionally, we see a notable difference between floating point and integer based matrices. This difference is due to the fact that the RP2040 does not feature any floating point arithmetic support and, thus, the cache effect is diminished due to longer duration of basic multiplications. Finally, we also observe that there is no notable difference in speed-up and cache-hit rate between the different splits.

We can therefore conclude that a simple change in the computational graph at compile time, *i.e.*, using column-major weight matrices, leads to increased matrix multiplication speeds on devices featuring a cached memory. Cache optimization is thus essential to speed-up inference of nested submodels part of REDS.

5.3 REDS layout format and adaptation cost

To store information about each subnetwork architecture, it is sufficient to store one integer value denoting the *slicing point* for each layer in each subnetwork, which corresponds to the number of active computational units. This is possible since active and inactive units build continuous grouped in memory. Activating a particular subnetwork means changing the length of the dimension of the weight tensor in each layer corresponding to its number of units. If optimization for caches is used by a layer, an additional bit should be stored to indicate that a flipped operation order is applied.

Switching from one REDS submodel to another requires adjusting only the sizes of the layers. There is no need to account for gaps in a layer due to the above optimization to ensure the weights are stored in a contiguous memory, even if the platform does not have a data cache. Due to a local scope of a slice, the modification affects only the incoming and outgoing connections. Inactive computational units do not participate in inference, although the values of the actual weights are never overwritten, since these are used by larger subnetworks. There is also no copying of weights. At inference time, the computational cost of running inference using a subnetwork is not affected by the presence of other REDS networks, in strong contrasts to the approaches that use masks to select active neurons or channels.

6 REDS on Mobile and IoT

We evaluate REDS models on two mobile phones, specifically Xiaomi Redmi Note 9 Pro and Google Pixel 6, and on two IoT devices, namely Arduino Nano 33 BLE Sense and Infineon CY8CKIT-062S2. For the mobile phone evaluation, we employed the official Google Tensorflow Lite benchmarking tool, which measures the model’s inference time on Android-based mobile devices. On IoT devices we deployed and evaluated the models using the cloud-based platform Edge Impulse (Hymel et al., 2022). We report the number of parameters, accuracy, and latency on the mobiles and on the IoT devices for the full architectures (100% MACs) and the subnetworks architecture (75%, 50% and 25% MACs) found by the BU knapsack solvers. The obtained results are reported in Fig. 7 for the DNN, CNN and DS-CNN architectures S (top row) and L (bottom row) on GOOGLE SPEECH COMMANDS. All results are averages over three runs. Tensorflow Lite and TFLMicro lack support for runtime adaptation of model weight tensors. We extended the TFLMicro

framework to support REDS out of the box² and measure on-device inference and submodel switching times.

The first left-most column of plots in Fig. 7 shows the relationship between the subnetworks’ MACs and the number of model parameters. All curves are close to linear, even though the parameters of these linear relationships are architecture-specific. The DS-CNN models have the lowest number of parameters, thanks to their more sophisticated architecture, which is more efficient on our dataset compared to the standard convolutional and dense networks. DS-CNN models also yield better accuracy, even for only 25% of MACs, while DNN models perform worst, as can be observed in the second column of plots. This result can be attributed to the successful generalization of the iterative knapsack problem to support depth-wise convolutions (Appendix A.1). We also compare REDS to several baselines. Recent works (Tan and Le, 2019) consider model optimization by pruning an equal share of computational blocks in each layer. We use the block selection strategy based on the L1 norm (used by Cai et al. (2020)) and a random selection. REDS knapsack solution outperforms both baselines achieving higher accuracy for the same percentage of MACs. In addition, the red star in the first two figures shows the performance of the constrained neural architecture search for microcontrollers μ NAS as reported by Liberis et al. (2021). The network architecture uses depth convolutions, yet is more complex than our DS-CNN. μ NAS shows 3% performance improvement over DS-CNN at the cost of 32% higher model size. Moreover, μ NAS took 39 GPU-days to compute the solution, in contrast to minutes taken by our knapsack solver.

The last two plots on the right show REDS performance on mobile and IoT devices. We observe a difference of three orders of magnitude in the inference times between the two categories. DNN models perform best thanks to the more optimized algorithm and libraries for matrix-matrix multiplication (Goto and Geijn, 2008). All models show a linear relationship between the percentage of MACs and inference time. This empirical evidence validates MACs as a robust predictor of model latency, extending its applicability to mobile devices.

We also compare REDS to binary masks, where instead of slicing a model (Sec. 2) we select the number of units for each layer by applying a binary mask to the layers’ inputs. Our tests on mobile devices in Fig. 7, 3rd column of plots show a constant overhead of masking for each subnetwork. On Arduino Nano 33 BLE Sense, TFLMicro framework extended to support REDS yields $38 \pm 1 \mu\text{s}$ model adaptation time for a 2-layer fully-connected network, while the model inference times for the same network with 25% and 50% of MACs are $2'131 \pm 27 \mu\text{s}$ and $4'548 \pm 13 \mu\text{s}$ respectively. This highlights the efficiency of permutation-based approach adopted by REDS. The energy consumption of REDS inference as measured with the Power Profiler Kit (PPK2) on Nordic nRF52840 for DS-CNN varies between 20mJ and 61mJ, whereas switching takes $<0.01\text{mJ}$. Results for other architectures are in the extended abstract.

7 Related Work

We structure related work by first giving a concise overview of model compression for resource-constrained devices that produce static, optimized, resource-agnostic models. We then survey existing machine learning frameworks for resource-constrained devices and their capabilities, followed by the discussion of the hardware support for deep learning. Finally, we summarize related works that support model adaptation at runtime.

Deep model compression. Deep networks are known to be highly redundant and their size can be reduced without hurting model accuracy. This motivated many model compression techniques to search for efficient subnetworks. Quantization and binarization rely on weights with discrete values, *e.g.*, used by Wu et al. (2016) to quantize filters and dense layers of a CNN. Decomposition and factorization explore low-rank basis of filters (Golubeva et al., 2020) to reduce model size and achieve inference speed-up. For example, Jaderberg et al. (2014) represent a filter of a CNN as a combination of low-rank filters leading to a considerable speed-up. Network pruning covers a set of methods that reduce the model size by eliminating unimportant operations (weights, neurons, kernels) in the model (Dai et al., 2018; Li et al., 2017). It can be combined with other techniques like quantization to further reduce the precision and overhead of operations (Han et al., 2016).

²<https://github.com/FraCorti/REDS#tensorflow-lite-for-microcontrollers-analysis>

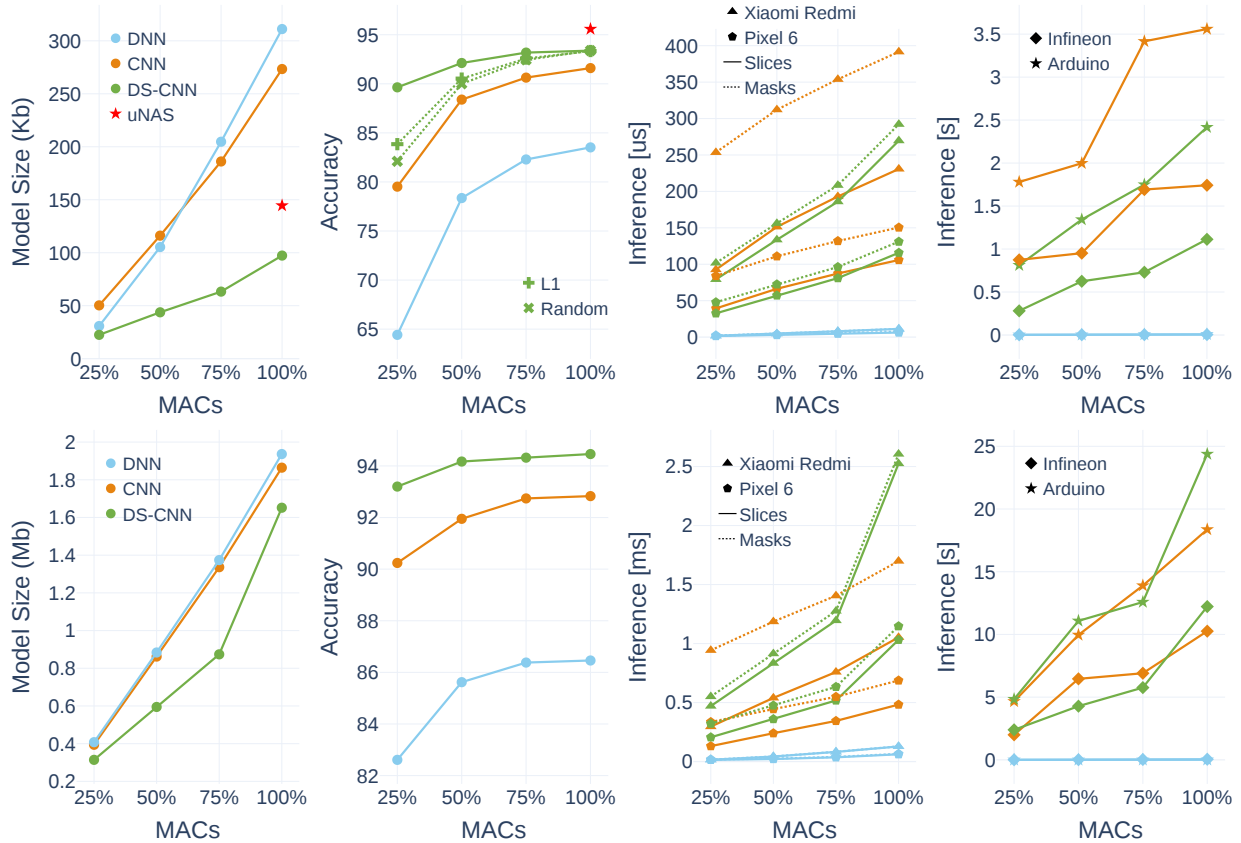


Figure 7: REDS size S (top row) and L (bottom row) architecture analysis. The two plots on the left show the number of model parameters and model accuracy as a function of MAC percentage in each REDS submodel. The two plots on the right evaluate model inference time on two classes of devices: more powerful platforms comprising Xiaomi Redmi Note 9 Pro (Qualcomm Snapdragon 720G, ARM Cortex-A76), Pixel 6 (Octa-core 2x2.8 GHz Cortex-X1, 2x2.25 GHz Cortex-A76, 4x1.8 GHz Cortex-A55); and low-power IoT platforms including Arduino Nano 33 BLE Sense (nRF52840, ARM Cortex-M4) and Infineon CY8CKIT-062S2.

Pruning methods date back to the optimal brain damage (LeCun and *et al.*, 1990) and the optimal brain surgeon (Hassibi and Stork, 1993), which suggest to prune the weights based on the Hessians of the loss function. Recent methods, *e.g.*, (Corti *et al.*, 2022; Entezari and Saukh, 2019; Han *et al.*, 2015; Timpl *et al.*, 2022), propose to prune network weights with a low magnitude or a low magnitude increase. Li *et al.* (2017) suggest to prune channels in CNNs based on a filter weight norm, while Hu *et al.* (2016) use the average percentage of zeros in the output to prune unimportant channels. Neural architecture search is used to build optimized models for a given hardware platform (Cai *et al.*, 2019; Tan *et al.*, 2019). Cai *et al.* (2020) train a once-for-all network that can be deployed under diverse architectural configurations to address heterogeneity of IoT hardware. An overview of model compression is given in Hoefler *et al.* (2021).

Optimization frameworks. Several libraries, frameworks and tools support machine learning on resource-constrained platforms. Bonsai (Kumar *et al.*, 2017) is a framework for resource-efficient decision trees which reduces the model size by learning a single sparse shallow tree. CMSIS-NN (Lai *et al.*, 2018) is a library developed for Cortex-M processor cores used by TFLMicro for hardware-accelerated execution of machine learning models. FastGRNN and FastRNN (Kusupati *et al.*, 2019) implement RNNs and gated RNNs on IoT

devices by leveraging residual connections, low-rank, sparse and quantized matrices. TensorFlow Lite (David et al., 2021) supports mobile and resource-scarce platforms. TFLMicro (David et al., 2021) allows hardware vendors to optimize kernels for their devices and improve their hardware performance. NCNN (NCN, 2019) is a framework for mobile deployment to support accelerated execution.

Hardware support for deep learning. In the past few years new low-power processors with optimized machine learning capabilities appeared on the market. In 2017 Microchip implemented the first processor with a high-performance 2D GPU (Mic, 2017). In 2019 ARM has launched the Helium technology (Dirvin, 2019) to be integrated into the Cortex-M processors and provide better machine learning capabilities with M-Profile Vector Extension. Modern sensors make use of embedded machine learning (Gherman et al., 2021a,b; Warden et al., 2022). Hardware accelerators, including power-optimized TPUs, NVIDIA Jetson, NVIDIA A100 to support 2:4 sparsity (Qu et al., 2022) frequently appear in edge systems.

Model adaptation to dynamic resources. Deep model adaptation is mainly explored in the context of input transformation and data distribution shifts (Saukh et al., 2023; Xu et al., 2018). State-of-the-art deep learning pipelines and post-training compression techniques build resource-agnostic models (Han et al., 2022). There are no standardized tools to adjust the computational graph over time to skip or include specific computational blocks. A resource-agnostic model does not account for changing resource availability, such as inference latency, memory usage and energy consumption, at runtime. Dynamic neural networks (Han et al., 2022) is a novel paradigm that allows neural networks to change the architecture and parameters at inference time. For example, Yu et al. (2019) propose a technique to train CNNs to support multiple sub-networks with different widths. In (Cai et al., 2020), the set of sub-networks was constructed by progressive shrinking, which enforces training from bigger sub-networks to the smaller ones. It allows the model to change the number of layers, widths and kernel size. In (Qu et al., 2022) the sub-networks are sampled by the magnitude of the rows' weights, and then optimized in parallel with a weighted loss function.

8 Conclusion, Limitations, Outlook

This paper presents Resource-Efficient Deep Subnetworks (REDS), an approach to adapting deep neural networks to dynamical resource constraints typical of mobile and IoT devices. REDS employ structured sparsity and neuron permutation invariance to craft nested subnetworks, facilitating adaptation to variable resource availability without compromising the models' accuracy or adding runtime overhead. Notably, REDS ensure the subnetworks' weights are stored in contiguous memory spaces, enhancing cache optimization and compatibility with hardware-specific enhancements like vector instructions. We formulate the design of REDS as an iterative knapsack problem and perform theoretical analysis of the solution space. Experimental evaluation on six benchmark architectures demonstrates the effectiveness of REDS on mobile and IoT devices and superior performance than the baselines.

Limitations. Our study has several limitations. First, we do not test the efficiency of REDS when combined with quantized models, which is left for future work. Secondly, the support of REDS for specific layers like skip connections is not explored and should be addressed in the future to support further advanced architectures.

Future work. Similarly to Fang et al. (2018), we plan to augment REDS with a task scheduler and deploy REDS in the context of a specific application with varying task priorities and dynamically changing resource constraints, such as on a flying drone.

Acknowledgements

The authors are grateful to Markus Gallacher for his support with the energy efficiency analysis of REDS. Christopher Hinterer and Julian Rudolf contributed to extending the TFL Micro framework to support REDS on Arduino Nano 33 BLE Sense. This research was funded in part by the Austrian Science Fund (FWF) within the DENISE doctoral school (grant number DFH 5). The results presented in this paper were computed using computational resources of HLR resources of the Zentralen Informatikdienstes of Graz University of Technology.

References

- Microchip introduces the industry’s first MCU with integrated 2D GPU and integrated DDR2 memory for groundbreaking graphics capabilities, 2017.
- NCNN: A high-performance neural network inference framework optimized for the mobile platform Topics. <https://github.com/Tencent/ncnn>, 2019.
- M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng. Tensorflow: a system for large-scale machine learning. 16 (2016):265–283, 2016.
- S. K. Ainsworth, J. Hayase, and S. Srinivasa. Git Re-Basin: Merging models modulo permutation symmetries. *arXiv preprint*, 2022. URL <https://arxiv.org/abs/2209.04836>.
- F. Bambusi, F. Cerizzi, Y. Lee, and L. Mottola. The case for approximate intermittent computing. *CoRR*, abs/2111.10726, 2021. URL <https://arxiv.org/abs/2111.10726>.
- M. Boehm, B. Reinwald, D. Hutchison, A. V. Evfimievski, and P. Sen. On optimizing operator fusion plans for large-scale machine learning in systemml. 2018. URL <https://arxiv.org/abs/1801.00829>.
- H. Cai, L. Zhu, and S. Han. Proxylessnas: Direct neural architecture search on target task and hardware. 2019. URL <https://arxiv.org/abs/1812.00332>.
- H. Cai, C. Gan, T. Wang, Z. Zhang, and S. Han. Once-for-all: Train one network and specialize it for efficient deployment. 2020. URL <https://arxiv.org/abs/1908.09791>.
- F. Corti, R. Entezari, S. Hooker, D. Bacciu, and O. Saukh. Studying the impact of magnitude pruning on contrastive learning methods. 2022. URL <http://arxiv.org/abs/2207.00200>.
- B. Dai, C. Zhu, and D. Wipf. Compressing neural networks using the variational information bottleneck. 2018. URL <https://arxiv.org/abs/1802.10399>.
- R. David, J. Duke, A. Jain, V. J. Reddi, N. Jeffries, J. Li, N. Kreeger, I. Nappier, M. Natraj, S. Regev, R. Rhodes, T. Wang, and P. Warden. TensorFlow Lite micro: Embedded machine learning for TinyML systems. *Proceedings of Machine Learning and Systems*, 3:800–811, 2021.
- F. Della Croce, U. Pferschy, and R. Scatamacchia. On approximating the incremental knapsack problem. *Discrete Applied Mathematics*, 264:26–42, 2019.
- R. Dirvin. Next-generation armv8.1-m architecture: Delivering enhanced machine learning and signal processing for the smallest embedded devices. <https://www.design-reuse.com/news/45581/arm-helium-armv8-1-m-architecture.html>, (last accessed: 2021-02-21), 2019.
- R. Entezari and O. Saukh. Class-dependent compression of deep neural networks. 2019. URL <http://arxiv.org/abs/1909.10364>.

- R. Entezari, H. Sedghi, O. Saukh, and B. Neyshabur. The role of permutation invariance in linear mode connectivity of neural networks. *arXiv preprint*, 2021. URL <https://arxiv.org/abs/2110.06296>.
- R. Entezari, M. Wortsman, O. Saukh, M. M. Shariatnia, H. Sedghi, and L. Schmidt. The role of pre-training data in transfer learning. 2023. URL <https://arxiv.org/abs/2302.13602>.
- Y. Faenza, D. Segev, and L. Zhang. Approximation algorithms for the generalized incremental knapsack problem. *Mathematical Programming*, 198:27–83, 2023.
- B. Fang, X. Zeng, and M. Zhang. NestDNN-device deep learning for continuous mobile vision. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, pages 115–127, 2018.
- E. Frantar and D. Alistarh. SPDY: accurate pruning with speedup guarantees. *CoRR*, abs/2201.13096, 2022. URL <https://arxiv.org/abs/2201.13096>.
- M.-P. Gherman, Y. Cheng, A. Gomez, and O. Saukh. Compensating Altered Sensitivity of Duty-Cycled MOX Gas Sensors with Machine Learning. In *IEEE SECON*, pages 368–379, 2021a.
- M.-P. Gherman, Y. Cheng, A. Gomez, and O. Saukh. Towards On-demand Gas Sensing. In *DCOSS*, pages 72–74, 2021b.
- A. Golubeva, B. Neyshabur, and G. Gur-Ari. Are wider nets better given the same number of parameters? 2020. URL <https://arxiv.org/abs/2010.14495>.
- K. Goto and R. A. v. d. Geijn. Anatomy of high-performance matrix multiplication. *ACM Transactions on Mathematical Software (TOMS)*, 34(3):1–25, 2008.
- Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2023. URL <https://www.gurobi.com>.
- S. Han, J. Pool, J. Tran, and W. Dally. Learning both weights and connections for efficient neural network. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, pages 1135–1143, 2015.
- S. Han, H. Mao, and W. J. Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. 2016. URL <http://arxiv.org/abs/1510.00149>.
- Y. Han, G. Huang, S. Song, and *et al.*. Dynamic neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44:7436–7456, 2022.
- B. Hassibi and D. G. Stork. Second order derivatives for network pruning: Optimal brain surgeon. *Proceedings of the 5th International Conference on Neural Information Processing Systems*, pages 164–171, 1993.
- T. Hoefler, D. Alistarh, and *et al.*. Sparsity in deep learning: Pruning and growth for efficient inference and training in neural networks, 2021. URL <https://arxiv.org/abs/2102.00554>.
- A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient cnns for mobile vision applications. 2017. URL <http://arxiv.org/abs/1704.04861>.
- H. Hu, R. Peng, Y.-W. Tai, and C.-K. Tang. Network trimming: A data-driven neuron pruning approach towards efficient deep architectures. 2016. URL <http://arxiv.org/abs/1607.03250>.
- S. Hymel, C. Banbury, D. Situnayake, A. Elium, C. Ward, M. Kelcey, M. Baaijens, M. Majchrzycki, J. Plunkett, D. Tischler, et al. Edge impulse: An mlops platform for tiny machine learning. 2022. URL <http://arxiv.org/abs/2212.03332>.
- S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, pages 448–456. pmlr, 2015.

- M. Jaderberg, A. Vedaldi, and A. Zisserman. Speeding up convolutional neural networks with low rank expansions. 2014. URL <http://arxiv.org/abs/1405.3866>.
- K. Jordan, H. Sedghi, O. Saukh, R. Entezari, and B. Neyshabur. Repair: Renormalizing permuted activations for interpolation repair. 2022. URL <http://arxiv.org/abs/2211.08403>.
- N. P. Jouppi, C. Young, and *et al.*. In-datacenter performance analysis of a tensor processing unit. In *Symposium on computer architecture*, pages 1–12, 2017.
- T. Kannan and H. Hoffmann. Budget RNNs: Multi-capacity nns to improve in-sensor inference under energy budgets. In *RTAS*, pages 143–156. IEEE, 2021.
- H. Kellerer, U. Pferschy, and D. Pisinger. *Knapsack Problems*. Springer, 2004.
- A. Krizhevsky, V. Nair, and G. Hinton. Cifar-100 and cifar-10 (canadian institute for advanced research), 2009. URL <http://www.cs.toronto.edu/~kriz/cifar.html>. MIT License.
- A. Kumar, S. Goyal, and M. Varma. Resource-efficient machine learning in 2 KB RAM for the internet of things. *Proceedings of the 34th International Conference on Machine Learning*, page 1935–1944, 2017.
- A. Kusupati, M. Singh, K. Bhatia, A. Kumar, P. Jain, and M. Varma. FastGRNN: A fast, accurate, stable and tiny kilobyte sized gated recurrent neural network. 2019. URL <https://arxiv.org/abs/1901.02358>.
- L. Lai, N. Suda, and V. Chandra. CMSIS-NN: Efficient neural network kernels for Arm Cortex-M CPUs. 2018. URL <https://arxiv.org/abs/1801.06601>.
- Y. LeCun and *et al.*. Optimal brain damage. *Advances in Neural Information Processing Systems*, pages 598–605, 1990.
- H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. P. Graf. Pruning filters for efficient convnets, 2017. URL <https://arxiv.org/abs/1608.08710>.
- E. Liberis, Ł. Dudziak, and N. D. Lane. μ NAS: Constrained neural architecture search for microcontrollers. *Proceedings of the 1st Workshop on Machine Learning and Systems*, pages 70–79, 2021.
- B. Maag, Z. Zhou, O. Saukh, and L. Thiele. BARTON: low power tongue movement sensing with in-ear barometers. In *ICPADS*, pages 9–16, 2017. doi: 10.1109/ICPADS.2017.00013. URL <https://doi.org/10.1109/ICPADS.2017.00013>.
- P. Molchanov, A. Mallya, S. Tyree, I. Frosio, and J. Kautz. Importance estimation for neural network pruning. *Computer Vision and Pattern Recognition*, 2019.
- M. C. Mozer and P. Smolensky. Skeletonization: A technique for trimming the fat from a network via relevance assessment. *NeurIPS*, 1:107–115, 1988.
- L. Perron and V. Furnon. Or-tools. URL <https://developers.google.com/optimization/>.
- Z. Qu, S. S. Sarwar, X. Dong, Y. Li, E. Sumbul, and B. De Salvo. DRESS: Dynamic real-time sparse subnets. 2022. URL <https://arxiv.org/abs/2207.00670>.
- D. Ravi, C. Wong, B. Lo, and G.-Z. Yang. Deep learning for human activity recognition. In *BSN*, pages 71–76. IEEE, 2016.
- RP2040. RP2040 Datasheet - A microcontroller by Raspberry Pi. <https://datasheets.raspberrypi.com/rp2040/rp2040-datasheet.pdf>, 2023.
- D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986.

- F. Sabry, T. Eltaras, and *et al.*. Machine learning for healthcare wearable devices: the big picture. *Journal of Healthcare Engineering*, 2022, 2022.
- M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- O. Saukh, D. Wang, X. He, and L. Thiele. Representing input transformations by low-dimensional parameter subspaces. 2023. URL <https://arxiv.org/abs/2305.13536>.
- M. Shen, H. Yin, P. Molchanov, L. Mao, J. Liu, and J. M. Alvarez. Structural pruning via latency-saliency knapsack. 2022. URL <https://arxiv.org/abs/2210.06659>.
- M. Tan and Q. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, pages 6105–6114. PMLR, 2019.
- M. Tan, B. Chen, R. Pang, V. Vasudevan, M. Sandler, A. Howard, and Q. V. Le. Mnasnet: Platform-aware neural architecture search for mobile, 2019. URL <https://arxiv.org/abs/1807.11626>.
- L. Timpl, R. Entezari, H. Sedghi, B. Neyshabur, and O. Saukh. Understanding the effect of sparsity on neural networks robustness. 2022. URL <https://arxiv.org/abs/2206.10915>.
- H. Vanholder. Efficient inference with tensorrt. In *GPU Technology Conference*, volume 1, page 2, 2016.
- P. Warden. Speech commands: A dataset for limited-vocabulary speech recognition. *CoRR*, 2018. URL <http://arxiv.org/abs/1804.03209>.
- P. Warden, M. Stewart, B. Plancher, C. Banbury, S. Prakash, E. Chen, Z. Asgar, S. Katti, and V. J. Reddi. Machine learning sensors. 2022. URL <https://arxiv.org/abs/2206.03266>.
- J. Wu, C. Leng, and *et. al.* Quantized convolutional neural networks for mobile devices. *CVPR*, pages 4820–4828, 2016.
- H. Xiao, K. Rasul, and R. Vollgraf. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. 2017. URL <http://arxiv.org/abs/1708.07747>.
- R. Xu, G. Li, J. Yang, and L. Lin. Unsupervised domain adaptation: An adaptive feature norm approach. *CoRR*, 2018. URL <http://arxiv.org/abs/1811.07456>.
- J. Yu, L. Yang, N. Xu, J. Yang, and T. Huang. Slimmable neural networks. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=H1gMCsAqY7>.
- Y. Zhang, N. Suda, L. Lai, and V. Chandra. Hello edge: Keyword spotting on microcontrollers. 2017. URL <http://arxiv.org/abs/1711.07128>.
- M. Zhu and S. Gupta. To prune, or not to prune: exploring the efficacy of pruning for model compression. 2017. URL <http://arxiv.org/abs/1710.01878>.

A Iterative Knapsack Theory

A.1 Knapsack for depth-wise convolutions

In this section we define a generalized knapsack framework for depth-wise convolutions. Let the input be a 2D matrix A with dimensions $m \times n$. The first layer L_0 is a standard convolutional layer (*cf.* Fig. 2) applied to A that consists of N_0 filters and hence produces N_0 output channels. It is followed by depth-wise convolutional blocks B_1, \dots, B_d , where each block i is built by a depth-wise layer B_i^D and a point-wise layer B_i^P for $i = 1, \dots, d$. A depth-wise layer applies one filter on each of the N_{i-1} input channels. The point-wise layer B_i^P then applies N_i filters, where the kernel number for each filter has to be equal to the number of filters of the previous layer. From a knapsack perspective, if we want to reduce the size of the network, we can decide on how many filters we use at L_0 and at each B_i^P . For example, if we choose k filters at L_0 the layer B_1^D will have exactly k filters and one filter of B_1^P will have k as kernels number. We give an integer programming formulation for this network structure. The resulting optimization method can be used for structural pruning of a neural network, *i.e.*, for choosing an optimal number ($\leq N_0$) of filters of L_0 and (at most N_i) filters of B_i^P such that we maximize performance of the network obeying a constraint on time or space complexity, which in REDS is the number of MACs.

Formally, we introduce an integer decision variable x_0 to decide on the number of filters we use at L_0 , and x_i on the number of filters to use at each B_i^P . Since every computational unit consumes computational resources and contributes to the overall accuracy, we also introduce decision variables that control whether a unit is used by a subnetwork or not. For L_0 we introduce N_0 binary variables y_1, \dots, y_{N_0} , and for every block B_i , $i = 1, \dots, d$, we introduce binary variables f_k^i and g_{kt}^i , where $k \in \{1, \dots, N_i\}$ and $t \in \{1, \dots, N_{i-1}\}$, f_k^i to indicate whether a filter k is used in B_i^P , and g_{kt}^i if filter k with kernel t of B_i^P is used. For B_i^D we introduce the decision variables d_t^i to decide if a depth-wise filter $t \in \{1, \dots, N_{i-1}\}$ is used.

In the model P_i^1 is the importance score of a standard convolution filter i in the first layer and W_1 is its number of MACs. For each depth-wise point-wise block i , P_t^i is the importance score for the depth-wise filter t , P_{kt}^i is the importance score of the corresponding kernel k of the point-wise filter t in the subsequent layer N_i . Due to permutation invariance (Sec. 3.1) the importance scores P^i and P_t^i are stored in descendant order. The problem is formulated as:

$$\begin{aligned}
 \max \quad & \sum_{i=1}^{N_0} y_i \cdot P_i^1 + \sum_{i=1}^d \sum_{t=1}^{N_{i-1}} \left(d_t^i \cdot P_t^i + \sum_{k=1}^{N_i} g_{kt}^i \cdot P_{kt}^i \right) \quad (1) \\
 s.t. \quad & \sum_{i=1}^{N_0} y_i \cdot W_1 + \sum_{i=1}^d \sum_{t=1}^{N_{i-1}} \left(d_t^i \cdot W_2 + \sum_{k=1}^{N_i} g_{kt}^i \cdot W_3 \right) \leq C \quad (2) \\
 & \sum_{i=1}^{N_0} y_i = x_0 \quad (3) \quad \text{and} \quad \sum_{t=1}^{N_{i-1}} d_t^i = x_{i-1} \quad \forall i \quad (4) \\
 & \sum_{k=1}^{N_i} f_k^i = x_i \quad \forall i \quad (5) \quad \text{and} \quad f_k^i \geq f_{k+1}^i \quad \forall i \quad (6) \\
 & g_{kt}^i \leq f_k^i \quad \forall i, k, t \quad (7) \quad \text{and} \quad f_k^i \leq \sum_{t=1}^{N_{i-1}} g_{kt}^i \quad \forall i, k \quad (8) \\
 & \sum_{t=1}^{N_{i-1}} g_{kt}^i \leq x_{i-1} \quad \forall i, k \quad (9) \\
 & \sum_{t=1}^{N_{i-1}} g_{kt}^i \geq x_{i-1} - (1 - f_k^i) \cdot N_{i-1} \quad \forall i, k \quad (10)
 \end{aligned}$$

The knapsack for depth-wise convolutions is described as:

- (1) The objective function maximizes the total importance score of the chosen architecture.
- (2) Solution MACs must comply with the constraint C.
- (3) The number of filters chosen in the first convolution layer.
- (4) The number of filters chosen in the depth-wise layer i must match the number of filters picked in the previous layer $i - 1$.
- (5) The number of the point-wise filters chosen in layer i .
- (6) Point-wise filters are chosen in ascending order to impose a contiguous solution.
- (7) If kernel t of filter k is chosen then the whole filter k is chosen.
- (8) A point-wise filter k is chosen only if one of its kernels is chosen.
- (9) The number of kernels in the filter k of point-wise layer i must \leq the number of filters taken in the previous depth-wise layer.
- (10) If filter k in layer i is chosen then the constraints (9) and (10) together ensure that the number of kernels t of filter k in layer i equals the number of point-wise filters in the previous block (*i.e.*, the number of filters in the depth-wise layer i). If filter k in layer i is not chosen, constraints (9) and (10) together imply that all kernels t of filter k at layer i are zero (the right-hand side of (9) is ≤ 0 , since N_{i-1} is an upper bound on x_{i-1}).

A.2 Iterative knapsack problem

In this section we prove that the order matters if we want to pack a knapsack iteratively. We are looking at the 2 stage iterative knapsack problem, where the items of a solution for capacity $c/2$ have to be a subset of the items of a solution for capacity c . We give our analysis and theoretical findings under the natural assumption that all items of the knapsack have weight $\leq c/2$. We first consider the bottom-up iterative knapsack heuristic and show that the quality of a worst-case solution is bounded by $\frac{2}{3} \cdot Opt$ and the bound is tight. We then analyze the top-down iterative knapsack heuristic and show that in this case a worst-case solution has a tight lower worst-case bound of $\frac{1}{2} \cdot Opt$. Approximation results for the related incremental knapsack problem were given in [Della Croce et al. 2019](#) and [Faenza et al. 2023](#). For a set of items I , let $P(I)$ ($W(I)$) denote the total profit (weight) of all items in I . For a binary knapsack problem (see [Kellerer et al. 2004](#) for a general overview) we say that a *split item* I_s according to some ordering O of the items and capacity c exists, if there is an item I_s with the following property: all the items I_b^O that appear in O before I_s fulfill $W(I_b^O) < c$ and $W(I_b^O \cup I_s) > c$.

Bottom-up knapsack. Let $I(Opt_c)$ denote the optimal solution set for a knapsack of capacity c . Consider the following iterative heuristic A_c : we first find an optimal solution of the knapsack with capacity $c/2$, then we fix the selected items $I(Opt_{c/2})$ and solve the knapsack defined on the remaining items and capacity $c - W(I(Opt_{c/2}))$. We denote this second set of items as $I(A_{c/2})$. Hence, the overall item set of this heuristic is given by $I(A_c) = I(Opt_{c/2}) \cup I(A_{c/2})$. Note that there may be $P(I(A_{c/2})) > P(I(Opt_{c/2}))$, since the corresponding knapsack capacity in the second step can be larger than $c/2$.

Theorem A.1. A_c yields a worst case ratio of $\frac{2}{3}$, *i.e.*, $P(I(A_c)) \geq \frac{2}{3} \cdot P(I(Opt_c))$, if all items of the knapsack have a weight $\leq c/2$. This bound is tight.

Proof. Let us arrange the items of $I(Opt_c)$ in the following order O : we first take all the items from $I(Opt_c) \cap I(Opt_{c/2})$ in an arbitrary order. Note that these are the items that are in both optimal solutions, *i.e.*, for both capacity c and $c/2$. Then we take all the items that are not included in $I(A_{c/2})$ followed by the items of $I(Opt_c) \cap I(A_{c/2})$ (again in arbitrary order). Now we have two cases:

Case 1: There does not exist a split item in $I(\text{Opt}_c)$ with respect to O and capacity $c/2$. Hence $W(I(\text{Opt}_{c/2})) = c/2$. It is easy to see that in this case $A_c = \text{Opt}_c$.

Case 2: Let I_s be the split item in $I(\text{Opt}_c)$ with respect to O and capacity $c/2$. In this case we get that the weight of all the items I_b^O before I_s as well as the weight of all the items I_f^O that follow I_s is smaller than $c/2$. It follows that $P(I(\text{Opt}_{c/2})) \geq P(I_b^O)$ and that $P(I(A_{c/2})) \geq P(I_f^O)$. Since all items have a weight $\leq c/2$ and by the fact that I_s is not contained in $I(\text{Opt}_{c/2})$ we know that its profit is less or equal than the minimum of $P(I(A_{c/2}))$ and $P(I(\text{Opt}_{c/2}))$. Therefore, it holds that $P(I_s) \leq \frac{1}{2}P(I(A_c))$. Hence we get:

$$\begin{aligned} P(I(\text{Opt}_c)) &= P(I_b^O) + P(I_s) + P(I_f^O) \leq P(I(A_c)) + \frac{1}{2}P(I(A_c)) \\ &= \frac{3}{2}P(I(A_c)) \end{aligned}$$

It remains to show the bound is tight. We introduce the following knapsack instance with four items and a large positive constant P .

item:	1	2	3	4
weight:	$c/3 + \epsilon$	$c/3$	$c/3$	$c/3$
profit:	$P + \epsilon$	P	P	P

Here $I(\text{Opt}_{c/2}) = \{1\}$ which only leaves space for one additional item for the larger capacity. Hence we get that $P(I(A_c)) = 2P + \epsilon$, whereas $P(I(\text{Opt}_c)) = 3P$. \square

Top-down knapsack. We now consider a heuristic $D_{c/2}$ consisting of an iterative top-down knapsack packing. We first solve the knapsack with capacity c to optimality, and then solve the knapsack problem defined only on the items $I(\text{Opt}_c)$ with capacity $c/2$ to optimality. $I(D_{c/2})$ corresponds to the items in this second and smaller knapsack.

Theorem A.2. $D_{c/2}$ yields a worst case ratio of $\frac{1}{2}$, i.e., $P(I(D_{c/2})) \geq \frac{1}{2} \cdot P(I(\text{Opt}_{c/2}))$ if all items of the knapsack have a weight $\leq c/2$. This bound is tight.

Proof. Consider a knapsack of size c with optimal solution set $I(\text{Opt}_c)$ and the knapsack problem with capacity $c/2$ defined on the restricted item set $I(\text{Opt}_c)$ with solution set $I(D_{c/2})$. We will show that: $P(I(D_{c/2})) \geq \frac{1}{2} \cdot P(I(\text{Opt}_{c/2}))$.

We first arrange the items of $I(\text{Opt}_c)$ in an ordering O' such that they start with those items contained also in $I(\text{Opt}_{c/2})$. Then we identify the split item I_s according to O' for capacity $c/2$ and partition $I(\text{Opt}_c)$ into three parts. $D_1 = I_b^{O'}$, $D_2 = I_s$ and D_3 contains all the remaining items. If no split item exists, we simply set $I_s = \emptyset$. We now show that:

$$\max(P(D_1), P(D_2), P(D_3)) \geq \frac{P(I(\text{Opt}_{c/2}))}{2} \tag{3}$$

Assuming that this is not the case, we would get that:

$$\max(P(D_1), P(D_2), P(D_3)) < \frac{P(I(\text{Opt}_{c/2}))}{2}$$

This would imply

$$P(I(\text{Opt}_c)) = P(D_1) + P(D_2) + P(D_3) < P(I(\text{Opt}_{c/2})) + P(D_3).$$

However, since $I(\text{Opt}_{c/2}) \cap D_3 = \emptyset$ and $W(I(\text{Opt}_{c/2})) \leq c/2$, $I(\text{Opt}_{c/2}) \cup D_3$ would constitute a feasible solution better than $I(\text{Opt}_c)$, which is a contradiction. Thus, we have shown (3).

For $i = 1, \dots, 3$, there is $W(D_i) \leq c/2$ and all items in D_i are available for $D_{c/2}$. Therefore, (3) implies

$$P(I(D_{c/2})) \geq \max(P(D_1), P(D_2), P(D_3)) \geq \frac{1}{2}P(I(Opt_{c/2})).$$

It remains to show the bound is tight. We introduce the following knapsack instance with four items and a large positive constant P .

Item:	1	2	3	4
Weight:	$c/3$	$c/3$	$c/3$	$c/2$
Profit:	$P + \epsilon$	$P + \epsilon$	$P + \epsilon$	$2P$

Here $I(Opt_c) = \{1, 2, 3\}$. $D_{c/2}$ then selects one of these items and no more items fit into the knapsack. $Opt_{c/2}$ selects item 4, which shows that the ratio of $\frac{1}{2}$ is tight. \square

Note that in case that we have instances, where the weight of certain items is greater than $c/2$, it is easy to construct instances with arbitrary bad ratios for both cases.

B Training Details

Table 3 summarizes the hyper-parameters used to train different networks. We refer to Zhang et al. 2017 regarding the description of the network architectures adopted in this paper (referred to as DNN, CNN and DS-CNN, sizes S and L). We used TensorFlow (Abadi et al., 2016) version 2.11. All models were trained on a workstation with 16 NVIDIA Tesla K80 GPUs and 32 Intel Xeon CPUs. To conduct all our experiments and compute the baselines we trained and optimized over 100 models. This translated into >1000h of compute. We always use 80:10:10 split for training, validation, and testing. All results constitute averages over three runs.

Hyper-parameter	DNN (S/L)	CNN (S/L)	DS-CNN (S/L)
Batch size		100	
Training epochs	75	75	250
Loss function		Cross-entropy	
Optimizer		Adam	
LR scheduler		PiecewiseConstantDecay	
LR	$10^{-3}, 10^{-4}$	$5 \cdot 10^{-4}, 10^{-4}, 2 \cdot 10^{-5}$	$10^{-3}, 10^{-4}$
LR steps	$1.5 \cdot 10^4,$ $3 \cdot 10^3$	$2 \cdot 10^4, 10^4, 10^4$	$10^4, 10^4, 10^4$

Table 3: Training hyper-parameters.

C Further REDS Performance Evaluation

C.1 REDS performance on DNN and CNN architectures

In addition to the results on DS-CNN reported in the main paper, we show in Table 4 and Table 5 REDS performance on DNN and CNN architectures (with full fine-tuning) and compare to training model of each capacity from scratch and training REDS from scratch. Despite full fine-tuning, the results for S architecture show superior performance of the BU heuristic over TD.

Fig. 8 shows the impact of the architecture on REDS structure found by the knapsack BU solver. We present the results for DNN S, L and CNN S, L from left to right, respectively.

MACs	Small (S) - Accuracy 83.82				Large (L) - Accuracy 86.87			
	Scratch	Knapsack BU	Knapsack TD	REDS training	Scratch	Knapsack BU	Knapsack TD	REDS training
100%	84.30 ±0.11	83.52 ±0.07	82.80 ±0.16	82.13 ±0.20	86.54 ±0.24	86.46 ±0.34	86.25±0.19	85.06 ±0.19
75%	83.77 ±0.23	82.29 ±0.35	81.88 ±0.30	81.23 ±0.21	85.96 ±0.13	86.38 ±0.65	86.09±0.03	84.93 ±0.20
50%	80.91 ±0.11	78.36 ±1.40	78.59 ±0.24	77.05 ±0.34	85.24 ±0.35	85.62 ±0.24	85.58±0.35	84.08 ±0.22
25%	69.77 ±0.67	64.42 ±1.99	61.43 ±3.35	63.69 ±3.26	84.22 ±0.13	82.61 ±0.61	83.00 ±0.62	82.03 ±0.59

Table 4: Test set accuracy [%] of training S and L fully-connected (DNN) architectures taken from Zhang et al. 2017: training a network of each size from scratch ("Scratch"), conversion from a pre-trained network using two knapsack versions ("Knapsack BU" and "Knapsack TD"), and training REDS structure from scratch ("REDS training"). Reported results from three independent runs. The accuracy of each 100% network reported in Zhang et al. 2017 is listed in the header row.

MACs	Small (S) - Accuracy 92.24				Large (L) - Accuracy 93.24			
	Scratch	Knapsack BU	Knapsack TD	REDS training	Scratch	Knapsack BU	Knapsack TD	REDS training
100%	91.10±0.23	91.60±0.39	91.20 ±0.35	88.89 ±0.26	92.94±0.20	92.83±0.26	92.97±0.15	90.97 ±0.21
75%	90.40±0.27	90.63±0.19	90.20 ±0.13	87.64 ±0.46	92.74±0.12	92.74±0.23	92.54 ±0.07	90.67 ±0.09
50%	89.07±0.24	88.39±0.31	88.39 ±0.52	85.99 ±0.19	92.44±0.30	91.95±0.22	91.93 ±0.28	90.36 ±0.30
25%	82.57±0.41	79.52±0.67	79.28 ±0.54	79.25 ±0.40	90.98±0.60	90.24±0.35	90.31 ±0.03	88.25 ±0.66

Table 5: The same as in Table 4 for S and L convolutional architectures (CNN) from Zhang et al. 2017.




Figure 8: Analysis of the subnetwork architecture obtained by the knapsack BU heuristics. From left to right: DNN S, DNN L, CNN S and CNN L on GOOGLE SPEECH COMMANDS. The patterns as to which computational units constitute a child subnetwork are architecture-specific.

C.2 REDS performance with 10 nested subnetworks

Table 6 and Fig. 9 show the performance of DNN, CNN and DS-CNN on GOOGLE SPEECH COMMANDS, when REDS structure comprises 10 subnetworks, compared to 4 subnetworks in the main paper. A larger number of subnetworks does not degrade model accuracy.

MACs	DNN		CNN		DS-CNN	
	Small	Large	Small	Large	Small	Large
100%	83.07 \pm 0.35	86.19 \pm 0.26	91.16 \pm 0.45	93.1 \pm 0.1	93.5 \pm 0.15	94.34 \pm 0.07
90%	82.93 \pm 0.4	86.17 \pm 0.36	90.4 \pm 0.47	92.84 \pm 0.27	93.33 \pm 0.11	94.32 \pm 0.1
80%	82.67 \pm 0.53	86.1 \pm 0.08	90.1 \pm 0.07	92.77 \pm 0.06	92.84 \pm 0.15	94.31 \pm 0.1
70%	81.67 \pm 0.53	85.78 \pm 0.11	89.43 \pm 0.41	92.25 \pm 0.47	92.64 \pm 0.24	94.21 \pm 0.14
60%	80.37 \pm 0.57	85.66 \pm 0.25	88.84 \pm 0.25	92.28 \pm 0.11	92.27 \pm 0.31	94.08 \pm 0.03
50%	78.26 \pm 0.53	85.33 \pm 0.09	88.4 \pm 0.2	92.03 \pm 0.23	91.04 \pm 0.51	93.97 \pm 0.04
40%	75.37 \pm 1.26	84.8 \pm 0.45	85.76 \pm 0.18	91.78 \pm 0.04	89.42 \pm 0.66	93.83 \pm 0.22
30%	67.76 \pm 1.59	82.66 \pm 0.18	81.92 \pm 0.49	90.52 \pm 0.54	87.63 \pm 1.03	93.59 \pm 0.14
20%	52.36 \pm 6.99	80.19 \pm 0.39	73.74 \pm 0.04	88.61 \pm 0.51	84.14 \pm 2.57	93.35 \pm 0.15
10%	23.93 \pm 5.88	50.7 \pm 7.5	58.87 \pm 5.27	81.15 \pm 1.39	58.46 \pm 3.35	90.38 \pm 0.17

Table 6: Test set accuracy [%] from Small (S) and Large (L) pretrained fully-connected (DNN), convolutional (CNN), and depth-wise separable convolutional (DS-CNN) networks taken from [Zhang et al. 2017](#). For each pre-trained architecture, REDS can support ten subnetworks obtained from the Knapsack BU formulation. The accuracies of the DS-CNN and CNN subnetworks do not degrade drastically until the lowest percentage of MACs considered. In contrast, the accuracies in the DNN subnetworks show a more pronounced drop from 30% MACs.

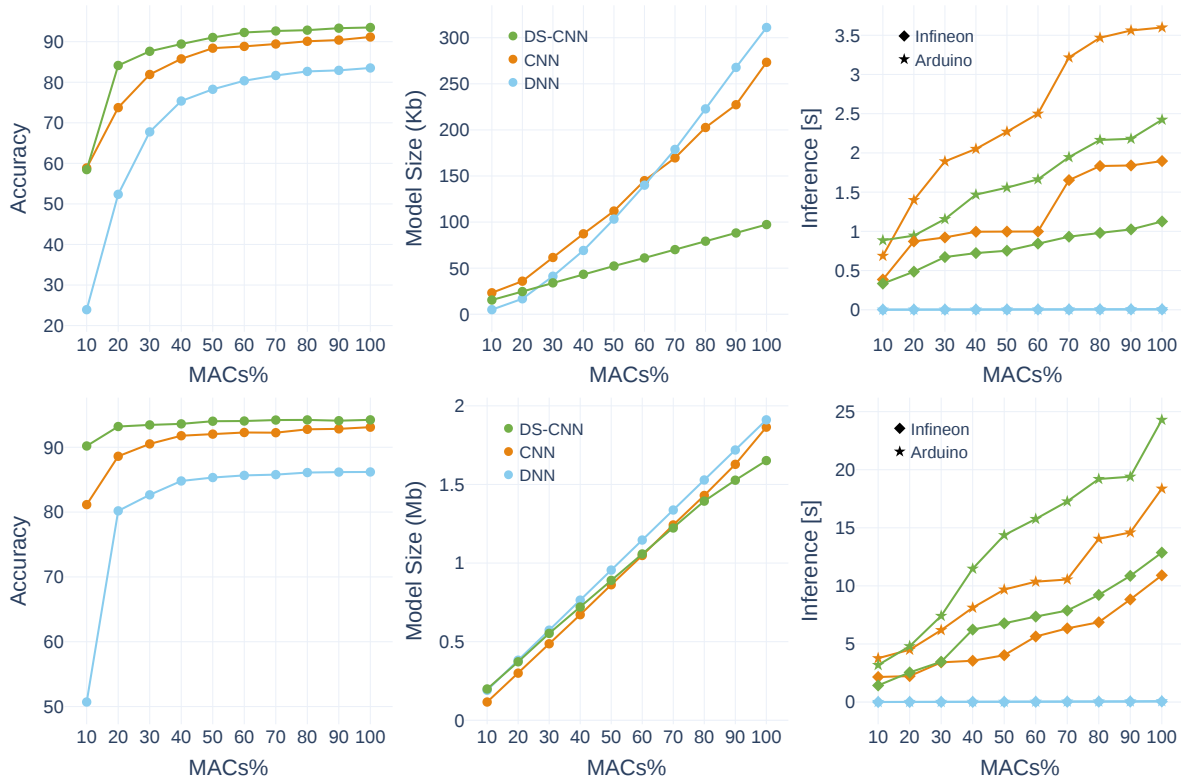


Figure 9: REDS size S (top row) and L (bottom row) architectures analysis finetuned on Google Speech Commands (Warden 2018) with ten subnetworks. The plots from left to right show the subnetworks size, the subnetworks accuracy and the subnetworks inference time as a function of MAC percentage.

C.3 REDS on FMNIST and CIFAR10

The results in Table 7 and Table 8 show REDS performance using DS-CNN architecture of size S on FMNIST and CIFAR10. BU heuristic was used to obtain the results. REDS supports a different data domain without degrading the accuracy of the pre-trained model, reported in the header row. Compared to the state-of-the-art such as μ NAS (Liberis et al. 2021), REDS demonstrates a faster architecture search time for both FMNIST and CIFAR10. In the former, REDS takes 19 minutes as opposed to 3 days; in the latter, REDS takes 90 minutes as opposed to 39 days while requiring less memory for model storage for both datasets. After finding and freezing the 25% MACs subnetwork architecture, the BU heuristic takes only a few seconds to find the other 50% and 75% MACs subnetworks architectures.

MACs	Acc (%) - Pre-trained 90.28	Model Size (Kb)	Time Taken (m)
100%	91.06 \pm 0.18	128.54	–
75%	90.56 \pm 0.1	107.73	1.58 [s]
50%	90 \pm 0.1	87.4	5.83 [s]
25%	87.86 \pm 0.39	66.63	19 [m]

Table 7: Analysis of the BU knapsack subnetworks obtained from a depth-wise separable convolutional (DS-CNN) S network, pre-trained on FMNIST (Xiao et al. 2017). REDS supports a different data domain without degrading the accuracy of the pre-trained model, reported in the header row.

MACs	Acc (%) - Pre-trained 78.28	Model Size (Kb)	Time Taken
100%	80.72 \pm 0.42	128.54	–
75%	78.8 \pm 0.75	109.41	2.89 [s]
50%	75.4 \pm 3.09	88.01	10.59 [s]
25%	62.34 \pm 1.22	69.63	90 [m]

Table 8: The same evaluation as in Table 7 for CIFAR10 (Krizhevsky et al. 2009).

C.4 REDS energy efficiency

Table 9 reports the energy consumption of inference time by different network architectures measured by Power Profiler Kit (PPK2) on Nordic nRF52840 (Arduino Nano 33 BLE Sense). In comparison, the model adaptation time takes less than 0.01 mJ.

MACs	DNN [mJ]	CNN [mJ]	DS-CNN [mJ]
100%	0.18	90.61	61
75%	0.15	86.01	44.03
50%	0.07	50.3	33.81
25%	0.05	44.81	20.44

Table 9: Energy consumption of REDS size S architectures measured by the Power Profiler Kit (PPK2) on Nordic nRF52840. The results are obtained by performing an inference pass for each subnetwork model and recording the inference current.