

Joslim: Joint Widths and Weights Optimization for Slimmable Neural Networks

Ting-Wu Chin¹✉, Ari S. Morcos², and Diana Marculescu^{1,3}

¹ Department of ECE, Carnegie Mellon University, Pittsburgh PA, USA

² Facebook AI Research, Menlo Park CA, USA

³ Department of ECE, The University of Texas at Austin, Austin TX, USA

✉ tingwuc@alumni.cmu.edu

Abstract. Slimmable neural networks provide a flexible trade-off front between prediction error and computational requirement (such as the number of floating-point operations or FLOPs) with the same storage requirement as a single model. They are useful for reducing maintenance overhead for deploying models to devices with different memory constraints and are useful for optimizing the efficiency of a system with many CNNs. However, existing slimmable network approaches either do not optimize layer-wise widths or optimize the shared-weights and layer-wise widths independently, thereby leaving significant room for improvement by joint width and weight optimization. In this work, we propose a general framework to enable joint optimization for both width configurations and weights of slimmable networks. Our framework subsumes *conventional and NAS-based slimmable methods* as special cases and provides flexibility to improve over existing methods. From a practical standpoint, we propose Joslim, an algorithm that jointly optimizes both the widths and weights for slimmable nets, which outperforms existing methods for optimizing slimmable networks across various networks, datasets, and objectives. Quantitatively, improvements up to 1.7% and 8% in top-1 accuracy on the ImageNet dataset can be attained for MobileNetV2 considering FLOPs and memory footprint, respectively. Our results highlight the potential of optimizing the channel counts for different layers *jointly* with the weights for slimmable networks. Code available at <https://github.com/cmu-enyac/Joslim>.

Keywords: Model Compression · Slimmable Neural Networks · Channel Optimization · Efficient Deep Learning

1 Introduction

Slimmable neural networks have been proposed with the promise of enabling multiple neural networks with different trade-offs between prediction error and the number of floating-point operations (FLOPs), *all at the storage requirement of only a single neural network* [42]. This is in stark contrast to channel pruning methods [4,39,14] that aim for a small standalone model. Slimmable neural

networks are useful for applications running on mobile and other resource-constrained devices. As an example, the ability to deploy multiple versions of the same neural network alleviates the maintenance overhead for applications which support a number of different mobile devices with different memory and storage constraints, as only one model needs to be maintained. On the other hand, slimmable networks can be critical for designing an efficient system that runs multiple CNNs. Specifically, an autonomous robot may execute multiple CNNs for various tasks at the same time. When optimizing the robot’s efficiency (overall performance *vs.* computational costs), it is unclear which CNNs should be trimmed by how much to achieve an overall best efficiency. As a result, methods based on trial-and-error are necessary for optimizing such a system. However, if trimming the computational requirement of any CNN requires re-training or fine-tuning, this entire process will be impractically expensive. In this particular case, if we replace each of the CNNs with their respective slimmable versions, optimizing a system of CNNs becomes practically feasible as slimmable networks can be slimmed without the need of re-training or fine-tuning.

A slimmable neural network is trained by simultaneously considering networks with different widths (or filter counts) using a single set of shared weights. The width of a child network is specified by a real number between 0 and 1, which is known as the “width-multiplier” [17]. Such a parameter specifies how many filters per layer to use proportional to the full network. For example, a width-multiplier of $0.35\times$ represents a network with 35% of the channel counts of the full network for all the layers. While specifying child networks using a single width-multiplier for all the layers has shown empirical success [40,42], such a specification neglects that different layers affect the network’s output differently [43] and have different FLOPs and memory footprint requirements [13], which may lead to sub-optimal results. As an alternative, neural architecture search (NAS) methods such as BigNAS [41] optimizes the layer-wise widths for slimmable networks, however, a sequential greedy procedure is adopted to optimize the widths and weights. As a result, the optimization of weights is not adapted to the optimization of widths, thereby leaving rooms for improvement by joint width and weight optimization.

In this work, we propose a framework for optimizing slimmable nets by formalizing it as minimizing the area under the trade-off curve between prediction error and some metric of interest, *e.g.*, memory footprint or FLOPs, with alternating minimization. Our framework subsumes both the universally slimmable networks [40] and BigNAS [41] as special cases. The framework is general and provides us with insights to improve upon existing alternatives and justifies our new algorithm Joslim, the first approach that jointly optimizes both shared-weights and widths for slimmable nets. To this end, we demonstrate empirically the superiority of the proposed algorithm over existing methods using various datasets, networks, and objectives. We visualize the algorithmic differences between the proposed method and existing alternatives in Fig. 1.

The contributions of this work are as follows:

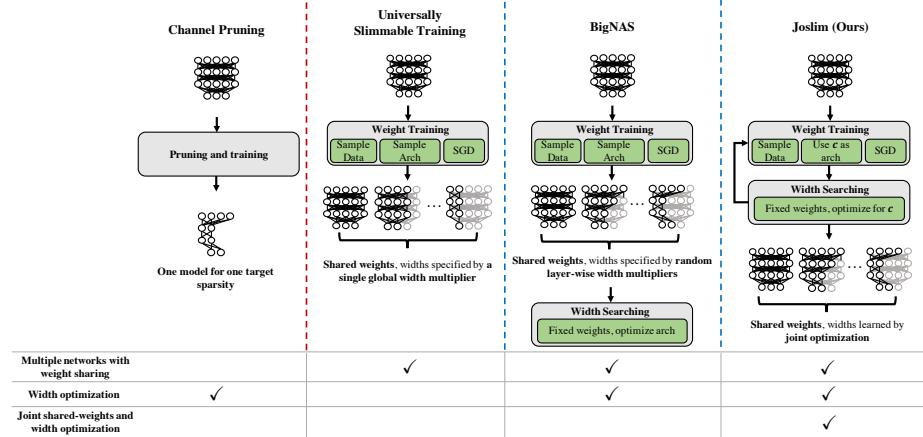


Fig. 1: Schematic overview comparing our proposed method with existing alternatives and channel pruning. Channel pruning has a fundamentally different goal compared to ours, *i.e.*, training slimmable nets. Joslim jointly optimizes both the widths and the shared weights.

- We propose a general framework that enables the joint optimization of the widths and their corresponding shared weights of a *slimmable net*. The framework is general and subsumes existing algorithms as special cases.
- We propose Joslim, an algorithm that jointly optimizes the widths and weights of slimmable nets. We show empirically that Joslim outperforms existing methods on various networks, datasets, and objectives. Quantitatively, improvements up to 1.7% and 8% in top-1 accuracy on ImageNet are attained for MobileNetV2 considering FLOPs and memory footprint, respectively.

2 Related work

2.1 Slimmable neural networks

Slimmable neural networks [42] enable multiple sub-networks with different compression ratios to be generated from a single network with one set of weights. This allows the network FLOPs to be dynamically configurable at run-time without increasing the storage requirement of the model weights. Based on this concept, better training methodologies have been proposed to enhance the performance of slimmable networks [40]. One can view a slimmable network as a dynamic computation graph where the graph can be constructed dynamically with different accuracy and FLOPs profiles. With this perspective, one can go beyond changing just the width of the network. For example, one can alter the network’s sub-graphs [31], network’s depth [5,18,21,19], and network’s kernel sizes and input resolutions [6,41,35,36]. Complementing prior work primarily focusing on generalizing slimmable networks to additional architectural paradigms, our

work provides the first principled formulation for jointly optimizing the weights and widths of slimmable networks. While our analysis focuses on the network widths, our proposed framework can be easily extended to other architectural parameters.

2.2 Neural architecture search

A slimmable neural network can be viewed as an instantiation of weight-sharing. In the literature for neural architecture search (NAS), weight-sharing is commonly adopted to reduce the search overhead [22,33,15,2,4,39]. Specifically, NAS methods use weight-sharing as a proxy for evaluating the performance of the sub-networks to reduce the computational requirement of iterative training and evaluation. However, the goal of NAS is the resulting architecture as opposed to both shared-weights and architecture. Exceptions are BigNAS [41] and Once-for-all (OFA) networks [6]; however, in neither case the architecture and shared-weights are jointly optimized. Specifically, both BigNAS and OFA employ a two-stage paradigm where the shared-weights are optimized before the architectures are optimized. This makes the trained weights oblivious to the optimized architectures.

While slimmable networks are inherently multi-objective, multi-objective optimization has also been adopted in NAS literature [11,7,26,12,37]. However, a crucial difference of the present work compared to these papers is that we are interested in learning a single set of weights from which multiple FLOP configurations can be used (as in slimmable networks) rather than finding architectures independently for each FLOP configuration that can be trained from scratch freely. Put another way, in our setting, both shared-weights and the searched architecture are optimized jointly, whereas in prior work, only searched architectures were optimized.

When it comes to joint neural architecture search and weight training, ENAS [29] and TuNAS [3] can both be seen as joint optimization. However, in stark contrast to our work, their search is dedicated to a single network of a single computational requirement (e.g., FLOPs) while our method is designed to obtain the weights that work for various architectures across a wide range of computational requirements.

2.3 Channel pruning

Reducing the channel or filter counts for a pre-trained model is also known as channel pruning. In channel pruning, the goal is to find a single small model that maximizes the accuracy while satisfying some resource constraints by optimizing the layer-wise channel counts [20,16,24,38,24,25,39,4,27,8,23,9]. While channel pruning also optimizes for non-uniform widths, the goal of channel pruning is crucially different from ours. The key difference is that channel pruning is concerned with a single pruned model while slimmable neural networks require a set of models to be trained using weight sharing. Nonetheless, we compare our work with pruning methods that conduct greedy channel pruning since

they naturally produce in a sequence of models that have different FLOPs. In particular, we compare our work with AutoSlim [39] in Appendix E.1 and demonstrate the effectiveness of our proposed Joslim.

3 Methodology

In this work, we are interested in jointly optimizing the network widths and network weights. Ultimately, when evaluating the performance of a slimmable neural network, we care about the trade-off curve between multiple objectives, *e.g.*, theoretical speedup and accuracy. This trade-off curve is formed by evaluating the two objectives at multiple width configurations using the same shared-weights. Viewed from this perspective, both the widths and shared-weights should be optimized in such a way that the resulting networks have a better trade-off curve (*i.e.*, larger area under curve). This section formalizes this idea and provides an algorithm to solve it in an approximate fashion.

3.1 Problem formulation

Our goal is to find both the weights and the width configurations that optimize the area under the trade-off curve between two competing objectives, *e.g.*, accuracy and inference speed. Without loss of generality, we use cross entropy loss as the accuracy objective and FLOPs as the inference speed objective throughout the text for clearer context. Note that FLOPs can also be replaced by other metrics of interest such as memory footprint. Since in this case both objectives are better when lower, the objective for the optimizing slimmable nets becomes to *minimize* the area under curve. To quantify the area under curve, one can use a Riemann integral. Let $w(c)$ be a width configuration of c FLOPs, one can quantify the Riemann integral by evaluating the cross entropy loss L_S on the training set \mathcal{S} using the shared weights $\boldsymbol{\theta}$ for the architectures that spread uniformly on the FLOPs-axis between a lower bound l and an upper bound u of FLOPs: $\{\mathbf{a}|\mathbf{a} = w(c), c \in [l, u]\}$. More formally, the area under curve \mathbb{A} is characterized as

$$\mathbb{A}(\boldsymbol{\theta}, w) \stackrel{\text{def}}{=} \int_l^u L_S(\boldsymbol{\theta}, w(c)) dc \quad (1)$$

$$\approx \sum_{i=0}^N L_S(\boldsymbol{\theta}, w(c_i)) \delta, \quad (2)$$

where equation 2 approximates the Riemann integral with the Riemann sum using N architectures that are spread uniformly on the FLOPs-axis with a step size δ . With a quantifiable area under curve, our goal for optimizing slimmable neural networks becomes finding both the shared-weights $\boldsymbol{\theta}$ and the architecture

function w to minimize their induced area under curve:

$$\arg \min_{\boldsymbol{\theta}, w} \mathbb{A}(\boldsymbol{\theta}, w) \approx \arg \min_{\boldsymbol{\theta}, w} \sum_{i=0}^N L_{\mathcal{S}}(\boldsymbol{\theta}, w(c_i)) \delta \quad (3)$$

$$= \arg \min_{\boldsymbol{\theta}, w} \frac{1}{N} \sum_{i=0}^N L_{\mathcal{S}}(\boldsymbol{\theta}, w(c_i)) \quad (4)$$

$$\approx \arg \min_{\boldsymbol{\theta}, w} \mathbb{E}_{c \sim U(l, u)} L_{\mathcal{S}}(\boldsymbol{\theta}, w(c)), \quad (5)$$

where $U(l, u)$ denotes a uniform distribution over a lower bound l and an upper bound u . Note that the solution to equation 5 is the shared-weight vector and a set of architectures, which is drastically different from the solution to the formulation used in the NAS literature [22,34], which is an architecture.

3.2 Proposed approach: Joslim

Since both the shared-weights $\boldsymbol{\theta}$ and the architecture function w are optimization variables of two natural groups, we start by using alternating minimization:

$$w^{(t+1)} = \arg \min_w \mathbb{E}_{c \sim U(l, u)} L_{\mathcal{S}}(\boldsymbol{\theta}^{(t)}, w(c)) \quad (6)$$

$$\boldsymbol{\theta}^{(t+1)} = \arg \min_{\boldsymbol{\theta}} \mathbb{E}_{c \sim U(l, u)} L_{\mathcal{S}}(\boldsymbol{\theta}, w^{(t+1)}(c)). \quad (7)$$

In equation 6, we maintain the shared-weights $\boldsymbol{\theta}$ fixed and for each FLOPs between l and u , we search for a corresponding architecture that minimizes the cross entropy loss. This step can be seen as a multi-objective neural architecture search given a fixed set of pre-trained weights, and can be approximated using smart algorithms such as multi-objective Bayesian optimization [28] or evolutionary algorithms [10]. However, even with smart algorithms, such a procedure can be impractical for every iteration of the alternating minimization.

In equation 7, one can use stochastic gradient descent by sampling from a set of architectures that spread uniformly across FLOPs obtained from solving equation 6. However, training such a weight-sharing network is practically $4 \times$ the training time of the largest standalone subnetwork [40] (it takes 6.5 GPU-days to train a slimmable ResNet18), which prevents it from being adopted in the alternating minimization framework.

To cope with these challenges, we propose targeted sampling, local approximation, and temporal sharing to approximate both equations.

Targeted sampling We propose to sample a set of FLOPs to approximate the expectation in equations 6 and 7 with empirical estimates. Moreover, the sampled FLOPs are shared across both steps in the alternating minimization so that one does not have to solve for the architecture function w (needed for the second step), but only solve for a set of architectures that have the corresponding

FLOPs. Specifically, we approximate the expectation in both equations 6 and 7 with the sample mean:

$$c_i^{(t)} \sim U(l, u) \quad \forall i = 1, \dots, M \quad (8)$$

$$w^{(t+1)} \approx \arg \min_w \frac{1}{M} \sum_{i=1}^M L_S \left(\boldsymbol{\theta}^{(t)}, w(c_i^{(t)}) \right) \quad (9)$$

$$\boldsymbol{\theta}^{(t+1)} \approx \arg \min_{\boldsymbol{\theta}} \frac{1}{M} \sum_{i=1}^M L_S \left(\boldsymbol{\theta}, w^{(t+1)}(c_i^{(t)}) \right). \quad (10)$$

From equation 9 and 10, we can observe that at any timestamp t , we only query the architecture function $w^{(t)}$ and $w^{(t+1)}$ at a fixed set of locations $c_i \forall i = 1, \dots, M$. As a result, instead of solving for the architecture function w , we solve for a fixed set of architectures $\mathcal{W}^{(t+1)}$ at each timestamp as follows:

$$\mathcal{W}^{(t+1)} := \{w^{(t+1)}(c_1), \dots, w^{(t+1)}(c_M)\} \quad (11)$$

where

$$\begin{aligned} w^{(t+1)}(c_i) &= \arg \min_{\mathbf{a}} L_S \left(\boldsymbol{\theta}^{(t)}, \mathbf{a} \right) \\ \text{s.t. FLOPs}(\mathbf{a}) &= c_i. \end{aligned} \quad (12)$$

With these approximations, for each iteration in the alternating minimization, we solve for M architectures with targeted FLOPs as opposed to solving for the entire approximate trade-off curve.

Local approximation To reduce the overhead for solving equation 10, we propose to approximate it with a few steps of gradient descent. Specifically, instead of training a slimmable neural network with sampled architectures until convergence in each iteration of alternating minimization (equation 10), we propose to perform K steps of gradient descent:

$$\begin{aligned} \mathbf{x}^0 &\stackrel{\text{def}}{=} \boldsymbol{\theta}^{(t)} \\ \mathbf{x}^{(k+1)} &\stackrel{\text{def}}{=} \mathbf{x}^{(k)} - \eta \frac{1}{M} \sum_{i=1}^M \nabla_{\mathbf{x}^{(k)}} L_S \left(\mathbf{x}^{(k)}, \mathcal{W}_i^{(t+1)} \right) \\ \boldsymbol{\theta}^{(t+1)} &\approx \mathbf{x}^{(K)}, \end{aligned} \quad (13)$$

where η is the learning rate. Larger K indicates better approximation with higher training overhead.

Temporal sharing Since we use local approximation, $\boldsymbol{\theta}^{(t+1)}$ and $\boldsymbol{\theta}^{(t)}$ would not be drastically different. As a result, instead of performing constrained neural architecture search *from scratch* (*i.e.*, solving for equation 12) in every iteration of

the alternating minimization, we propose to share information across the search procedures in different iterations of the alternation.

To this end, we propose to perform temporal sharing for multi-objective Bayesian optimization with random scalarization (MOBO-RS) [28] to solve equation 12. MOBO-RS itself is a sequential model-based optimization algorithm, where one takes a set of architectures \mathcal{H} , builds models (typically Gaussian Processes [30]) to learn a mapping from architectures to cross entropy loss g_{CE} and FLOPs g_{FLOPs} , scalarizes both models into a single objective with a random weighting λ (λ controls the preference for cross entropy and FLOPs), and finally optimizes the scalarized model to obtain a new architecture and stores the architecture back to the set \mathcal{H} . This entire procedure repeats for T iterations for one MOBO-RS.

To exploit temporal similarity, we propose MOBO-TS2, which stands for multi-objective Bayesian optimization with targeted scalarization and temporal sharing. Specifically, we propose to let $T = 1$ and share \mathcal{H} across alternating minimization. Additionally, we modify the random scalarization with targeted scalarization where we use binary search to search for the λ that results in the desired FLOPs. As such, \mathcal{H} grows linearly with the number of alternations. In such an approximation, for each MOBO in the alternating optimization, we reevaluate the cross-entropy loss for each $\mathbf{a} \in \mathcal{H}$ to build faithful GPs. We further provide theoretical analysis for approximation via temporal similarity for Bayesian optimization in Appendix D.

Joslim Based on this preamble, we present our algorithm, Joslim, in Algorithm 1. In short, Joslim has three steps: (1) build surrogate functions (*i.e.*, GPs) and acquisition functions (*i.e.*, UCBs) using historical data \mathcal{H} and their function responses, (2) sample M target FLOPs and solve for the corresponding widths (*i.e.*, \mathbf{a}) via binary search with the scalarized acquisition function and store them back to \mathcal{H} , and (3) perform K gradient descent steps using the solved widths. The first two steps solve equation 12 with targeted sampling and temporal sharing, and the final step solves equation 10 approximately with local approximation. In the end, to obtain the best widths, we use non-dominated sorting based on the training loss and FLOPs for $\mathbf{a} \in \mathcal{H}$.

3.3 Relation to existing approaches

For direct comparison with our work we consider the universally slimmable neural networks [40], which uses a single width multiplier to specify the widths of a slimmable network and NAS-based approaches such as OFA [6] and BigNAS [41], which have decoupled widths and weights optimization. To demonstrate the generality of the proposed framework, we show how these previously published works are special cases of our framework.

Slim Universally slimmable networks [40], or Slim for short, is a special case of our framework where the widths are not optimized but pre-specified by a single

Algorithm 1: Joslim

Input : Model parameters $\boldsymbol{\theta}$, lower bound for width-multipliers $w_0 \in [0, 1]$, number of full iterations F , number of gradient descent updates K , number of λ samples M

Output : Trained parameter $\boldsymbol{\theta}$, approximate Pareto front \mathcal{N}

```

1  $\mathcal{H} = \{\}$       (Historical minimizers  $\mathbf{a}$ )
2 for  $i = 1 \dots F$  do
3    $\mathbf{x}, y = \text{sample\_data}()$ 
4    $\mathbf{u}_{\text{CE}}, \mathbf{u}_{\text{FLOPs}} = L_{\text{CE}}(\mathcal{H}; \boldsymbol{\theta}, \mathbf{x}, y), \text{FLOPs}(\mathcal{H})$ 
5    $\mathbf{g}_{\text{CE}}, \mathbf{g}_{\text{FLOPs}} = \text{GP\_UCB}(\mathcal{H}, \mathbf{u}_{\text{CE}}, \mathbf{u}_{\text{FLOPs}})$ 
6   widths = []
7   for  $m = 1 \dots M$  do
8      $\mathbf{a} = \text{MOBO\_TS2}(\mathbf{g}_{\text{CE}}, \mathbf{g}_{\text{FLOPs}}, \mathcal{H})$       (Algorithm 2)
9     widths.append( $\mathbf{a}$ )
10  end
11   $\mathcal{H} = \mathcal{H} \cup \text{widths}$       (update historical data)
12  widths.append( $w_0$ )
13  for  $j = 1 \dots K$  do
14     $\text{SlimmableTraining}(\boldsymbol{\theta}, \text{widths})$ 
15    (line 3-16 of Algorithm 1 in [40])
16  end
17   $\mathcal{N} = \text{nonDominatedSort}(\mathcal{H}, \mathbf{u}_{\text{CE}}, \mathbf{u}_{\text{FLOPs}})$ 
18 end

```

Algorithm 2: MOBO-TS2

Input : Acquisition functions $\mathbf{g}_{\text{CE}}, \mathbf{g}_{\text{FLOPs}}$, historical data \mathcal{H} , search precision ϵ

Output : channel configurations \mathbf{a}

```

1  $c = \text{Uniform}(l, u)$       (Sample a target FLOPs)
2  $\lambda_{\text{FLOPs}}, \lambda_{\text{min}}, \lambda_{\text{max}} = 0.5, 0, 1$ 
3 while  $|\frac{\text{FLOPs}(\mathbf{a}) - c}{\text{FullModelFLOPs}}| > \epsilon$  do          // binary search
4    $c = \arg \min_c \text{Scalarize}(\lambda_{\text{FLOPs}}, \mathbf{g}_{\text{CE}}, \mathbf{g}_{\text{FLOPs}})$ 
5   if  $\text{FLOPs}(\mathbf{a}) > c$  then
6      $\lambda_{\text{min}} = \lambda_{\text{FLOPs}}$ 
7      $\lambda_{\text{FLOPs}} = (\lambda_{\text{FLOPs}} + \lambda_{\text{max}})/2$ 
8   else
9      $\lambda_{\text{max}} = \lambda_{\text{FLOPs}}$ 
10     $\lambda_{\text{FLOPs}} = (\lambda_{\text{FLOPs}} + \lambda_{\text{min}})/2$ 
11  end
12 end

```

global width multiplier. This corresponds to solving equation 5 with w given as a function that returns the width that satisfies some FLOPs by controlling a single global width multiplier. Our framework is more general as it introduces the freedom for optimizing the widths of slimmable nets.

OFA and BigNAS OFA and BigNAS use the same approach when it comes to the channel search space⁴. They are also a special case of our framework where the optimization of the widths and the shared-weights are carried out greedily. Specifically, BigNAS first trains the shared-weights by random layer-wise width multipliers. After convergence, BigNAS performs evolutionary search to optimize the layer-wise width multipliers considering both error and FLOPs. This greedy algorithm can be seen as performing one iteration of alternating minimization by solving equation 7 followed by solving equation 6. From this perspective, one can observe that the shared-weights θ are not jointly optimized with the widths. Our framework is more general and enables joint optimization for both widths and weights.

As we demonstrate in Section 4.2, our comprehensive empirical analysis reveals that Joslim is superior to either approach when compared across multiple networks, datasets, and objectives.

4 Experiments

4.1 Experimental setup

For all the Joslim experiments in this sub-section, we set K such that Joslim only visits 1000 width configurations throughout the entire training ($|\mathcal{H}| = 1000$). Also, we set M to be 2, which follows the conventional slimmable training method [40] that samples two width configurations in between the largest and the smallest widths. As for binary search, we conduct at most 10 binary searches with ϵ set to 0.02, which means that the binary search terminates if the FLOPs difference is within a two percent margin relative to the full model FLOPs. On average, the procedure terminates by using 3.4 binary searches for results on ImageNet. The dimension of \mathbf{a} is network-dependent and is specified in Appendix A and the training hyperparameters are detailed in Appendix B. To arrive at the final set of architectures for Joslim, we use non-dominated sort based on the training loss and FLOPs for $\mathbf{a} \in \mathcal{H}$.

4.2 Performance gains introduced by Joslim

We consider three datasets: CIFAR-10, CIFAR-100, and ImageNet. To provide informative comparisons, we verify our implementation for the conventional slimmable training with the reported numbers in [40] using MobileNetV2 on

⁴ Since we only search for channel counts, the progressive shrinking strategy proposed in OFA does not apply. As a result, both OFA and BigNAS have the same approach.

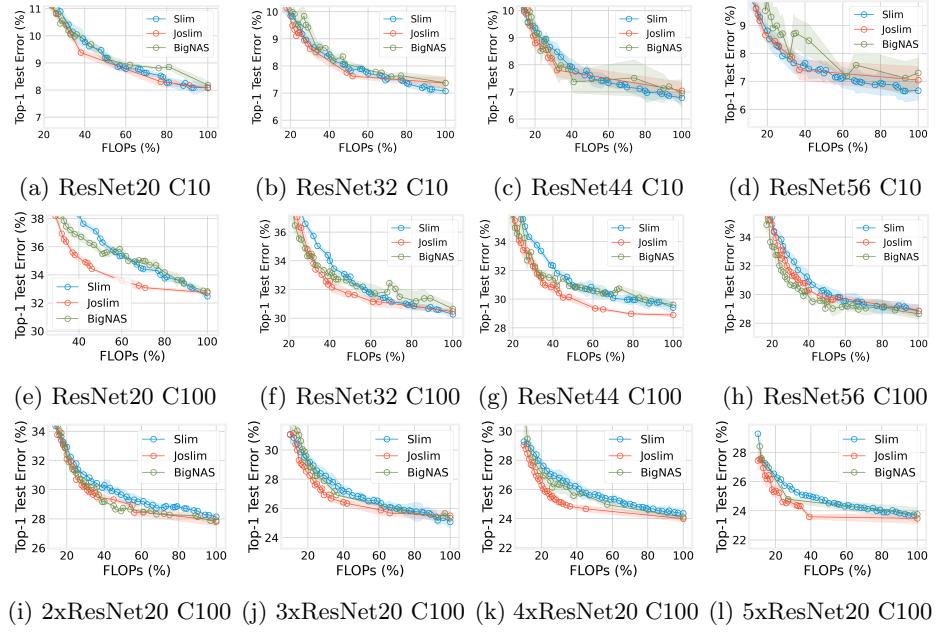


Fig. 2: Comparisons among Slim, BigNAS, and Joslim. C10 and C100 denote CIFAR-10/100. We perform three trials for each method and plot the mean and standard deviation. nx ResNet20 represents a n times wider ResNet20.

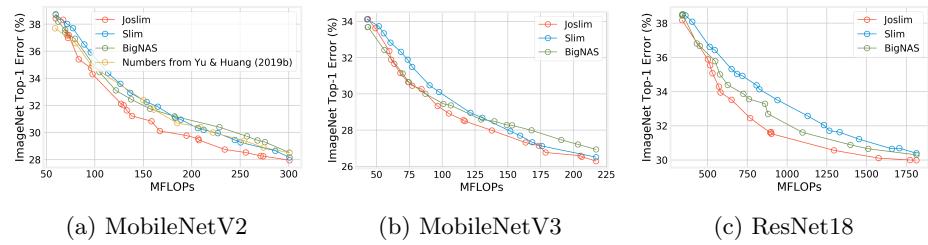


Fig. 3: Comparisons among Slim, BigNAS, and Joslim on ImageNet.

ImageNet. Our results follow closely to the reported numbers as shown in Fig. 3a, which makes our comparisons on other datasets convincing.

We compare to the following baselines:

- **Slim:** the conventional slimmable training method (the universally slimmable networks by [40]). We select 40 architectures uniformly distributed across FLOPs and run a non-dominated sort using training loss and FLOPs.
- **BigNAS:** disjoint optimization that first trains the shared-weights, then uses search methods to find architectures that work well given the trained weights (similar to OFA [6]). To compare fairly with Joslim, we use MOBO-RS for the

MobileNetV2				MobileNetV3				ResNet18			
MFLOPs	Slim	BigNAS	Joslim	MFLOPs	Slim	BigNAS	Joslim	MFLOPs	Slim	BigNAS	Joslim
59	61.4	61.3	61.5	43	65.8	66.3	65.9	339	61.5	61.5	61.8
84	63.0	63.1	64.6	74	68.1	68.1	68.8	513	63.4	64.2	64.5
102	64.7	65.5	65.5	85	69.1	70.0	70.0	650	64.7	65.6	66.5
136	67.1	67.5	68.2	118	71.0	71.4	71.4	718	65.1	66.1	67.5
149	67.6	68.2	69.1	135	71.5	71.5	72.1	939	66.5	67.3	68.5
169	68.2	68.8	69.9	169	72.7	72.0	72.8	1231	68.0	68.4	69.4
212	69.7	69.6	70.6	184	73.0	72.5	73.2	1659	69.3	69.3	69.9
300	71.8	71.5	72.1	217	73.5	73.1	73.7	1814	69.6	69.7	70.0

Table 1: Comparing the top-1 accuracy among Slim, BigNAS, and Joslim on ImageNet. Bold represents the highest accuracy of a given FLOPs.

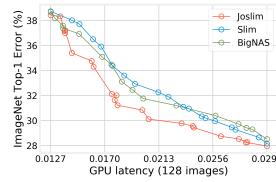


Fig. 4: A latency-*vs.*-error view of Fig. 3a.

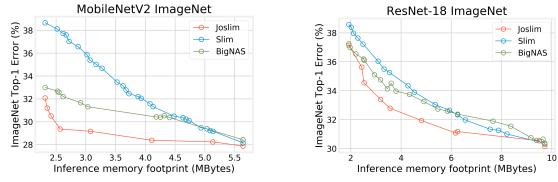


Fig. 5: Prediction error *vs.* inference memory footprint for MobileNetV2 and ResNet18 on ImageNet.

search. After optimization, we run a non-dominated sort for all the visited architectures \mathcal{H} using training loss and FLOPs.

The main results for the CIFAR dataset are summarized in Fig. 2 with results on ImageNet summarized in Figure 3 and Table 1. Compared to *Slim*, the proposed Joslim has demonstrated much better results across various networks and datasets. This suggests that channel optimization can indeed improve the efficiency of slimmable networks. Compared to *BigNAS*, Joslim is better or comparable across networks and datasets. This suggests that joint widths and weights optimization leads to better overall performance for slimmable nets. From the perspective of training overhead, Joslim introduced minor overhead compared to Slim due to the temporal similarity approximation. More specifically, on ImageNet, Joslim incurs approximately 20% extra overhead compared to Slim.

Note that the performance among these three methods are similar for the CIFAR-10 dataset. This is plausible since when a network is more over-parameterized, there are many solutions to the optimization problem and it is easier to find solutions with the constraints imposed by weight sharing. In contrast, when the network is relatively less over-parameterized, compromises have to be made due to the constraints imposed by weight sharing. In such scenarios, Joslim outperforms Slim significantly, as it can be seen in CIFAR-100 and ImageNet experiments. We conjecture that this is because Joslim introduces a new optimization variable (width-multipliers), which allows better compromises to be attained. Similarly, from the experiments with ResNets on CIFAR-100 (Fig. 2e to Fig. 2h), we

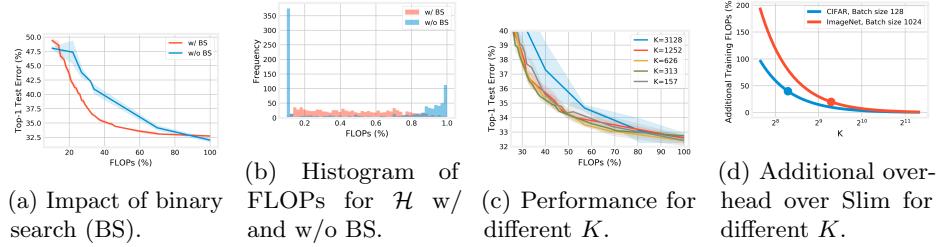


Fig. 6: Ablation study for the introduced binary search and the number of gradient descent updates per full iteration using ResNet20 and CIFAR-100. Experiments are conducted three times and we plot the mean and standard deviation.

find that shallower models tend to benefit more from joint channel and weight optimization than their deeper counterparts.

As FLOPs may not necessarily reflect latency improvements since FLOP does not capture memory accesses, we in addition plot latency-*vs.*-error for the data in Fig. 3a in Fig. 4. The latency is measured on a single V100 GPU using a batch size of 128. When visualized in latency, Joslim still performs favorably compared to Slim and BigNAS for MobileNetV2 on ImageNet.

Lastly, we consider another objective that is critical for on-device machine learning, *i.e.*, inference memory footprint [42]. Inference memory footprint decides whether a model is executable or not on memory-constrained devices. We detailed the memory footprint calculation in Appendix E. Since Joslim is general, we can replace the FLOPs calculation with memory footprint calculation to optimize for memory-*vs.*-error. As shown in Fig. 5, Joslim significantly outperform other alternatives. Notably, Joslim outperforms Slim by up to 8% top-1 accuracy for MobileNetV2. Such a drastic improvement comes from the fact that memory footprint depends mostly on the largest layers. As a result, slimming all the layers equally to arrive at networks with smaller memory footprint (as done in Slim) is less than ideal since only one layer contributes to the reduced memory. In addition, when comparing Joslim with BigNAS, we can observe significant improvements as well, *i.e.*, around 2% top-1 accuracy improvements for MobileNetV2, which demonstrates the effectiveness of joint width and weights optimization.

4.3 Ablation studies

In this subsection, we ablate the hyperparameters that are specific to Joslim to understand their impact. We use ResNet20 and CIFAR-100 for the ablation with the results summarized in Fig. 6.

Binary search Without binary search, one can also consider sampling the scalarization weighting λ uniformly from $[0, 1]$, which does not require any binary search and is easy to implement. However, the issue with this sampling strategy is that uniform sampling λ does not necessarily imply uniform sampling in the

objective space, *e.g.*, FLOPs. As shown in Fig. 6a and Fig. 6b, sampling directly in the λ space results in non-uniform FLOPs and worse performance compared to binary search.

Number of gradient descent steps In the approximation, the number of architectures ($|\mathcal{H}|$) is affected by the number of gradient descent updates K . In previous experiments for CIFAR, we have $K = 313$, which results in $|\mathcal{H}| = 1000$. Here, we ablate K to 156, 626, 1252, 3128 such that $|\mathcal{H}| = 2000, 500, 250, 100$, respectively. Given a fixed training epoch and batch size, Joslim produces a better approximation for equation 10 but a worse approximation for equation 9 with larger K . The former is because of the local approximation while the latter is because there are overall fewer iterations put into Bayesian optimization due to temporal sharing. As shown in Fig. 6c, we observe worse results with higher K . On the other hand, the improvement introduced by lower K saturates quickly. The overhead of Joslim as a function of K compared to Slim is shown in Fig. 6d where the dots are the employed K .

5 Conclusion

In this work, we are interested in optimizing both the architectural components and shared-weights of slimmable neural networks. To achieve this goal, we propose a general framework that optimizes slimmable nets by minimizing the area under the trade-off curve between cross entropy and FLOPs (or memory footprint) with alternating minimization. We further show that the proposed framework subsumes existing methods as special cases and provides flexibility for devising better algorithms. To this end, we propose Joslim, an algorithm that jointly optimizes the weights and widths of slimmable nets, which empirically outperforms existing alternatives that either neglect width optimization or conduct widths and weights optimization independently. We extensively verify the effectiveness of Joslim over existing techniques on three datasets (*i.e.*, CIFAR10, CIFAR100, and ImageNet) with two families of network architectures (*i.e.*, ResNets and MobileNets) using two types of objectives (*i.e.*, FLOPs and memory footprint). Our results highlight the importance and superiority in results of jointly optimizing the channel counts for different layers and the weights for slimmable networks.

Acknowledgement

This research was supported in part by NSF CCF Grant No. 1815899, NSF CSR Grant No. 1815780, and NSF ACI Grant No. 1445606 at the Pittsburgh Supercomputing Center (PSC).

References

1. Balandat, M., Karrer, B., Jiang, D.R., Daulton, S., Letham, B., Wilson, A.G., Bakshy, E.: BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In: NeurIPS (2020)

2. Bender, G., Kindermans, P.J., Zoph, B., Vasudevan, V., Le, Q.: Understanding and simplifying one-shot architecture search. In: ICML. pp. 550–559 (2018)
3. Bender, G., Liu, H., Chen, B., Chu, G., Cheng, S., Kindermans, P.J., Le, Q.V.: Can weight sharing outperform random architecture search? an investigation with tunas. In: CVPR. pp. 14323–14332 (2020)
4. Berman, M., Pishchulin, L., Xu, N., Medioni, G., et al.: Aows: Adaptive and optimal network width search with latency constraints. CVPR (2020)
5. Bolukbasi, T., Wang, J., Dekel, O., Saligrama, V.: Adaptive neural networks for efficient inference. In: ICML (2017)
6. Cai, H., Gan, C., Wang, T., Zhang, Z., Han, S.: Once-for-all: Train one network and specialize it for efficient deployment. In: ICLR (2020)
7. Cheng, A.C., Dong, J.D., Hsu, C.H., Chang, S.H., Sun, M., Chang, S.C., Pan, J.Y., Chen, Y.T., Wei, W., Juan, D.C.: Searching toward pareto-optimal device-aware neural architectures. In: ICCAD. pp. 1–7 (2018)
8. Chin, T.W., Ding, R., Zhang, C., Marculescu, D.: Towards efficient model compression via learned global ranking. In: CVPR (2020)
9. Chin, T.W., Marculescu, D., Morcos, A.S.: Width transfer: on the (in) variance of width optimization. In: CVPR Workshops. pp. 2990–2999 (2021)
10. Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.: A fast and elitist multiobjective genetic algorithm: Nsga-ii. IEEE transactions on evolutionary computation **6**(2), 182–197 (2002)
11. 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: ECCV (2018)
12. Elsken, T., Metzen, J.H., Hutter, F.: Efficient multi-objective neural architecture search via lamarckian evolution. arXiv preprint arXiv:1804.09081 (2018)
13. Gordon, A., Eban, E., Nachum, O., Chen, B., Wu, H., Yang, T.J., Choi, E.: Morphnet: Fast & simple resource-constrained structure learning of deep networks. In: CVPR. pp. 1586–1595 (2018)
14. Guo, S., Wang, Y., Li, Q., Yan, J.: Dmcp: Differentiable markov channel pruning for neural networks. In: CVPR. pp. 1539–1547 (2020)
15. Guo, Z., Zhang, X., Mu, H., Heng, W., Liu, Z., Wei, Y., Sun, J.: Single path one-shot neural architecture search with uniform sampling. ECCV (2020)
16. He, Y., Liu, P., Wang, Z., Hu, Z., Yang, Y.: Filter pruning via geometric median for deep convolutional neural networks acceleration. In: CVPR. pp. 4340–4349 (2019)
17. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861 (2017)
18. Huang, G., Chen, D., Li, T., Wu, F., van der Maaten, L., Weinberger, K.Q.: Multi-scale dense networks for resource efficient image classification. ICLR (2018)
19. Kaya, Y., Hong, S., Dumitras, T.: Shallow-Deep Networks: Understanding and mitigating network overthinking. In: ICML. Long Beach, CA (Jun 2019)
20. Li, H., Kadav, A., Durdanovic, I., Samet, H., Graf, H.P.: Pruning filters for efficient convnets. arXiv preprint arXiv:1608.08710 (2016)
21. Li, H., Zhang, H., Qi, X., Yang, R., Huang, G.: Improved techniques for training adaptive deep networks. In: ICCV. pp. 1891–1900 (2019)
22. Liu, H., Simonyan, K., Yang, Y.: Darts: Differentiable architecture search. arXiv preprint arXiv:1806.09055 (2018)
23. Liu, Z., Mu, H., Zhang, X., Guo, Z., Yang, X., Cheng, K.T., Sun, J.: Metapruning: Meta learning for automatic neural network channel pruning. In: ICCV (2019)
24. Liu, Z., Li, J., Shen, Z., Huang, G., Yan, S., Zhang, C.: Learning efficient convolutional networks through network slimming. In: ICCV. pp. 2736–2744 (2017)

25. Louizos, C., Welling, M., Kingma, D.P.: Learning sparse neural networks through l_0 regularization. arXiv preprint arXiv:1712.01312 (2017)
26. Lu, Z., Whalen, I., Boddeti, V., Dhebar, Y., Deb, K., Goodman, E., Banzhaf, W.: Nsga-net: neural architecture search using multi-objective genetic algorithm. In: GECCO. pp. 419–427 (2019)
27. Ma, X., Triki, A.R., Berman, M., Sagonas, C., Cali, J., Blaschko, M.B.: A bayesian optimization framework for neural network compression. In: ICCV (2019)
28. Paria, B., Kandasamy, K., Póczos, B.: A flexible framework for multi-objective bayesian optimization using random scalarizations. In: Globerson, A., Silva, R. (eds.) UAI (2019)
29. Pham, H., Guan, M., Zoph, B., Le, Q., Dean, J.: Efficient neural architecture search via parameters sharing. In: ICML. pp. 4095–4104. PMLR (2018)
30. Rasmussen, C.E.: Gaussian processes in machine learning. In: Summer School on Machine Learning. pp. 63–71. Springer (2003)
31. Ruiz, A., Verbeek, J.: Adaptative inference cost with convolutional neural mixture models. In: ICCV. pp. 1872–1881 (2019)
32. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: CVPR. pp. 4510–4520 (2018)
33. Stamoulis, D., Ding, R., Wang, D., Lymberopoulos, D., Priyantha, B., Liu, J., Marculescu, D.: Single-path nas: Designing hardware-efficient convnets in less than 4 hours. In: ECML-PKDD (2019)
34. Tan, M., Chen, B., Pang, R., Vasudevan, V., Sandler, M., Howard, A., Le, Q.V.: Mnasnet: Platform-aware neural architecture search for mobile. In: CVPR. pp. 2820–2828 (2019)
35. Wang, D., Gong, C., Li, M., Liu, Q., Chandra, V.: Alphanet: Improved training of supernet with alpha-divergence. In: ICML (2021)
36. Yang, T., Zhu, S., Chen, C., Yan, S., Zhang, M., Willis, A.: Mutualnet: Adaptive convnet via mutual learning from network width and resolution. In: ECCV. pp. 299–315. Springer (2020)
37. Yang, Z., Wang, Y., Chen, X., Shi, B., Xu, C., Xu, C., Tian, Q., Xu, C.: Cars: Continuous evolution for efficient neural architecture search. In: CVPR (June 2020)
38. Ye, J., Lu, X., Lin, Z., Wang, J.Z.: Rethinking the smaller-norm-less-informative assumption in channel pruning of convolution layers. In: ICLR (2018)
39. Yu, J., Huang, T.: Autoslim: Towards one-shot architecture search for channel numbers. arXiv preprint arXiv:1903.11728 **8** (2019)
40. Yu, J., Huang, T.S.: Universally slimmable networks and improved training techniques. In: ICCV. pp. 1803–1811 (2019)
41. Yu, J., Jin, P., Liu, H., Bender, G., Kindermans, P.J., Tan, M., Huang, T., Song, X., Pang, R., Le, Q.: Bignas: Scaling up neural architecture search with big single-stage models. In: ECCV (2020)
42. Yu, J., Yang, L., Xu, N., Yang, J., Huang, T.: Slimmable neural networks. In: ICLRa (2019)
43. Zhang, C., Bengio, S., Singer, Y.: Are all layers created equal? arXiv preprint arXiv:1902.01996 (2019)

A Width parameterization

For ResNets with CIFAR, \mathbf{a} has six dimensions and is denoted by $\mathbf{a}_{1:6} \in [0.316, 1]$, *i.e.*, one parameter for each stage and one for each residual connected layers in three stages. More specifically, the network is divided into three stages according to the output resolution, and as a result, there are three stages for all the ResNets designed for CIFAR. For example, in ResNet20, there are 7, 6, and 6 layers for each of the stages, respectively. Also, the layers that are added together via residual connection have to share the same width-multiplier, which results in one width-multiplier per stage for the layers that are connected via residual connections.

For MobileNetV2, $\mathbf{a}_{1:25} \in [0.42, 1]$, and therefore there is one dimension for each independent convolutional layer. Note that while there are in total 52 convolutional layers in MobileNetV2, not all of them can be altered independently. More specifically, for layers that are added together via residual connection, their widths should be identical. Similarly, the depth-wise convolutional layer should have the same width as its preceding point-wise convolutional layers. The same logic applies to MobileNetV3, which has 47 convolutional layers (excluding squeeze-and-excitation layers) and $\mathbf{a}_{1:22} \in [0.42, 1]$. In MobileNetV3, there are squeeze-and-excitation (SE) layers and we do not alter the width for the expansion layer in the SE layer. The output width of the SE layer is set to be the same as that of the convolutional layer where the SE layer is applied to. Note that there is no concept of expansion ratio for the inverted residual block in MobileNets in our width optimization. More specifically, the convolutional layer that acts upon expansion ratio is in itself just a convolutional layer with tunable width. Also, we do not quantize the width to be multiples of 8 as adopted in the previous work [32,40]. Due to these reasons, our $0.42 \times$ MobileNetV2 has 59 MFLOPs, which has the same FLOPs as the $0.35 \times$ MobileNetV2 in [40,32].

B Training hyperparameters

We use PyTorch as our deep learning framework and we build MOBO-TS on BoTorch [1], which works seamlessly with PyTorch. More specifically, for the covariance function of Gaussian Processes, we use the commonly adopted Matérn Kernel without changing the default hyperparameters provided in BoTorch. Similarly, we use $\text{beta} = 0.1$ for the Upper Confidence Bound acquisition function. To perform the optimization of line 6 in Algorithm 2, we make use of the API “*optimize_acqf*” provided in BoTorch. As a reference, with a single 1080Ti GPU, one can train a Joslim-ResNet20 on CIFAR-100 with around 3 hours. On the other hand, with 8 V100 GPUs on a single machine, one can train a Joslim-ResNet18 on ImageNet with 19 hours.

CIFAR The training hyperparameters for the independent models are 0.1 initial learning rate, 200 training epochs, 0.0005 weight decay, 128 batch size, SGD with nesterov momentum, and cosine learning rate decay. The accuracy on the

validation set is reported using the model at the final epoch. For slimmable training, we keep the same exact hyperparameters but train $2\times$ longer compared to independent models, *i.e.*, 400 epochs.

ImageNet Our training hyperparameters follow that of [40]. Specifically, we use initial learning rate of 0.5 with 5 epochs linear warmup (from 0 to 0.5), linear learning rate decay (from 0.5 to 0), 250 epochs, $4e^{-5}$ weight decay, 0.1 label smoothing, and we use SGD with 0.9 nesterov momentum. We use a batch size of 1024. For data augmentation, we use the “RandomResizedCrop” and “RandomHorizontalFlip” APIs in PyTorch. For MobileNetV2 we follow [40] and use random scale between 0.25 to 1. For MobileNetV3, we use the default scale parameters, *i.e.*, from 0.08 to 1. The input resolution we use is 224. Besides scaling and horizontal flip, we follow [40] and use color and lighting jitters data augmentation with parameter of 0.4 for brightness, contrast, and saturation; and 0.1 for lighting. These augmentations can be found in the official repository of [40]⁵. The hyperparameters for training ResNet18 is identical to MobileNetV2 except that we train it for 100 epochs only. The training for ImageNet is done using 8 NVIDIA V100 GPUs.

C Width differences

In Fig. 7, we visualize the widths learned by Joslim and contrast them with Slim for MobileNetV2 and MobileNetV3. Note that both Joslim and Slim are slimmable networks with shared weights and from the top row to the bottom row represent three points on the trade-off curve for Fig. 3a and Fig. 3b.

D Theoretical analysis for temporal approximation

The intuition behind the proposed approximation in Section 3.2 is the similarity for $\boldsymbol{\theta}$ across alternating minimization. In an extreme case, if we hold $\boldsymbol{\theta}$ constant throughout the training procedure, the approximation is equivalent to the original multi-objective BO. With that said, $\boldsymbol{\theta}$ changes gradually throughout training. To proceed with further theoretical understanding, we assume the loss $L_S(\boldsymbol{\theta})$ is L-Lipschitz. More formally,

$$L_S(\boldsymbol{\theta}^t) - L_S(\boldsymbol{\theta}^{t+1}) \leq L\|\boldsymbol{\theta}^t - \boldsymbol{\theta}^{t+1}\|_1, \forall \boldsymbol{\theta}^t, \boldsymbol{\theta}^{t+1}. \quad (14)$$

Now, consider using stochastic gradient descent to update the weights $\boldsymbol{\theta}$, *i.e.*, $\boldsymbol{\theta}^{t+1} = \boldsymbol{\theta}^t - \eta^t \mathbf{g}^t$ where \mathbf{g}^t is the gradient of loss with respect to the weights and η^t is the learning rate at iteration t . Since L_S is L-Lipschitz, we have $\|\mathbf{g}\|_1 \leq L$. Assuming using an exponential decaying learning rate with a factor $\gamma < 1$, we

⁵ https://github.com/JiahuiYu/slimmable_networks/blob/master/train.py#L43



Fig. 7: Comparing the width-multipliers between Joslim and Slim. The title for each plot denotes the relative differences (Joslim - Slim) and the numbers in the parenthesis are for Joslim.

can further upper bound the functional differences across K iterations of gradient descents as follows:

$$L_S(\boldsymbol{\theta}^t) - L_S(\boldsymbol{\theta}^{t+n}) \leq \sum_{i=t}^{t+K} \eta^i \|g^i\| \leq K\eta^t L. \quad (15)$$

Aligning with our intuition, the analysis reveals that larger K implies poorer approximation for Bayesian optimization to share history. In multi-objective Bayesian optimization [28], the hyperparameter is searched over stationary objectives. In our case, due to temporal approximation, our cross entropy changes over time and the change is upper-bounded by $K\eta^t L$. As a result, we can plug such an upper bound in the regret bound analysis of Bayesian optimization [28] to understand how K , η , and γ affect the optimality of Bayesian optimization. Specifically, we upper bound $L_S(\boldsymbol{\theta}^t)$ with $L_S(\boldsymbol{\theta}^{t+K}) + K\eta^t L$ and use it in Lemma 2 and Lemma 3 from [28] in Appendix B.1. With such a technique, a regret bound will have the following overhead in addition to the original regret bound

in equation (14) of [28]:

$$\frac{2\eta^0}{1-\gamma} K L \mathbb{E}[L_\lambda] K', \quad (16)$$

where we have utilized the geometric progression of the exponential learning rate decay and L , $\mathbb{E}[L_\lambda]$, and K' are the notations used by [28]. In other words, without a decaying learning rate, the overhead can be unbounded. This analysis reveals that larger initial learning rate η^0 and K results in a worse regret bound.

E Inference memory footprint calculation

To demonstrate the generality of proposed Joslim, we in addition consider optimizing for the trade-off curve between prediction error and inference memory footprint. The inference memory footprint is a critical factor when it comes to deploying deep CNNs onto resource-constrained devices such as mobile phones or micro-controllers as motivated in the original slimmable neural network paper [42]. We use a single image per batch to calculate the memory footprint. Specifically the inference memory footprint is characterized as follows:

$$\begin{aligned} FM_{in}^l &= W_{in}^l \times H_{in}^l \times C_{in}^l \\ FM_{out}^l &= W_{out}^l \times H_{out}^l \times C_{out}^l \\ \text{Weights}^l &= K_w^l \times K_h^l \times C_{in}^l \times C_{out}^l / G^l \\ \text{Skip}^l &= W_{out}^l \times H_{out}^l \times C_{skip}^l \\ MEM &= \max_l \left(FM_{in}^l + FM_{out}^l + \text{Weights}^l + \text{Skip}^l \right), \end{aligned} \quad (17)$$

where FM_{in}^l and FM_{out}^l denote the input and output feature map sizes of layer l , Weights^l denotes the size of the weights in layer l , and Skip^l denotes the memory requirement of storing the feature maps from skip connections. W and H represent the width and height of the feature map. K_w and K_h denote the kernel size. Lastly, C_{in} , C_{out} and G denote the input channel, output channel, and the number of groups for convolutional layer l .

E.1 Comparisons with AutoSlim

AutoSlim [39] is a NAS method proposed to do channel search for *standalone models*. While it also provides non-uniform widths for different layers, it is not a method derived to solve for equation 5. Specifically, AutoSlim first perform weight-sharing training (equation 7) for a short amount of period, i.e., 10% to 20% of the full training epochs. Then, AutoSlim conducts greedy pruning to greedily remove a fixed amount of channels from the layer that affects the loss the least. Such a greedy procedure naturally results in a sequence of models of different computational requirements. Crucially, this greedy pruning procedure is not solving equation 12 since the computational requirement (FLOPs or memory

footprint) does not affect the ranking among filters to be pruned. Nonetheless, we can adopt AutoSlim to obtain a sequence of models and train them via weight-sharing to form a slimmable network. We compared with AutoSlim using ResNet18 on ImageNet with both FLOPs and memory footprint. As shown in Fig. 8, Joslim performs similarly with AutoSlim when it comes to FLOPs, and this is because ResNet18 has balanced FLOPs for all the layers. On the other hand, when it comes to memory footprint, Joslim outperforms AutoSlim. In hindsight, this result is not surprising as Joslim is designed to solve equation 5 while AutoSlim is not.

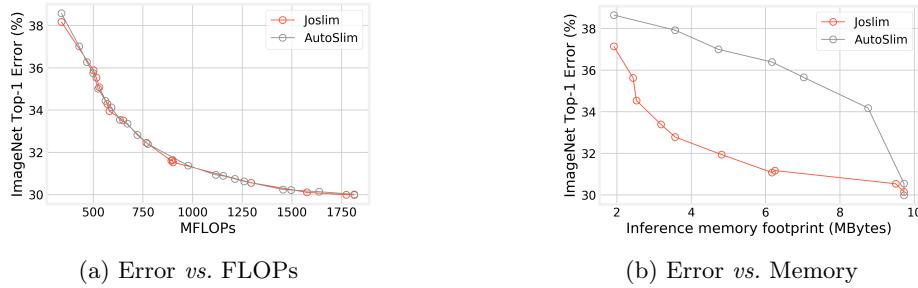


Fig. 8: Comparing Joslim and AutoSlim on ResNet18. Since ResNet18 has similar FLOPs across different layers, greedy pruning used by AutoSlim perform comparably to Joslim. However, Joslim outperforms AutoSlim when it comes to optimizing for memory consumption since the greedy pruning procedure adopted by AutoSlim is not multi-objective.