QuantNAS for super resolution: searching for efficient quantization-friendly architectures against quantization noise

Egor Shvetsov*   Dmitry Osin*   Alexey Zaytsev   Ivan Koryakovskiy   Valentin Buchnev

Ilya Trofimov   Evgeny Burnaev

Abstract

There is a constant need for high-performing and computationally efficient neural network models for image super-resolution (SR) often used on low-capacity devices. One way to obtain such models is to compress existing architectures, e.g. quantization. Another option is a neural architecture search (NAS) that discovers new efficient solutions. We propose a novel quantization-aware NAS procedure for a specifically designed SR search space. Our approach performs NAS to find quantization-friendly SR models. The search relies on adding quantization noise to parameters and activations instead of quantizing parameters directly. Our QuantNAS finds architectures with better PSNR/BitOps trade-off than uniform or mixed precision quantization of fixed architectures. Additionally, our search against noise procedure is up to 30% faster than directly quantizing weights.

1. Introduction

Neural networks (NNs) have become a default solution for many problems because of their high performance. Wider adoption of NNs requires high computational efficiency. Researchers either compress, search or jointly search and compress architectures aiming for more computationally effective solutions [26]. Computational efficiency is especially important for models utilized on mobile devices.

In this paper, we choose to combine techniques of NAS and quantization to search for efficient quantization-friendly models for SR. The NAS problem is hard because we should either define a differentiable NAS procedure or use discrete optimization in a high-dimensional space of architectures. The problem is even more challenging for quantization-aware NAS because quantization is an indif-ferentiable operation. Therefore, optimization of quantized models is more difficult than full precision models. An additional technical challenge arises for SR, as Batch Norm (BN) in SR models damages final performance [18] but training models without BN is much slower. Finally, we should define an appropriate search space given numerous recent advances in quantized SR architectures and take into account that the size of the discrete search space grows exponentially with the introduction of new parts of the architecture, making the optimization problem harder.

Our contributions:

- We propose the first end2end approach to NAS for mixed precision quantization of SR architectures.

- To search for robust, quantization-friendly architectures we approximate model degradation caused by quantization with Quantization Noise (QN) instead of directly quantizing model weights during the search phase. For sampling QN we follow procedure proposed in [7]. Such reparametrization allows differentiability crucial for a differentiable NAS and is up to 30% faster than quantizing weights directly.

- We design a specific search space for SR models. The proposed search space is simpler than the current SOTA - TrilevelNAS [24] and leads to about 5 times faster search procedure. We show that the search space design is equally important as search methods and argue that the community should pay more attention to the search space design.

- Quantization-aware NAS with Search Against Noise (SAN) yields results with a better trade-off between quality measured in PSNR and efficiency measured in BitOps compared to uniform and mixed precision quantization of fixed architectures. Thus, the joint quantization-aware NAS is a better choice than separate quantization and NAS.

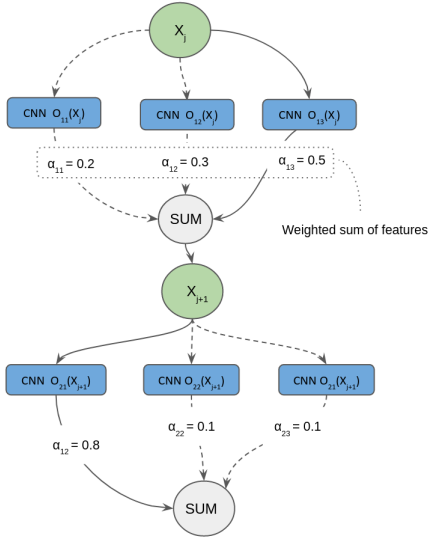---

*Indicates equal contribution.

2. Related works

Figure 1. The search space for NAS is equivalent to an over-parametrized supernet represented as a graph. In this graph, multiple possible operations connect nodes that are outputs of each layer. α values represent the importance of edges. Joint training of parameters of operations and their importance allow differentiable NAS. The final architecture is the result of the selection of edges with the highest importance for each consecutive pair of nodes. They have solid lines in the figure.
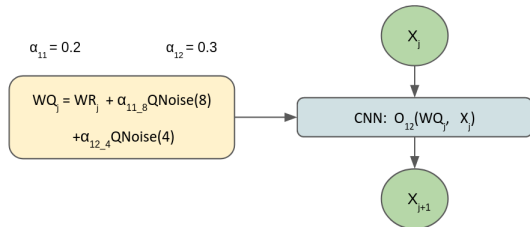

Figure 2. SAN approach, $QNoise(b)$ - is some function to generate quantiziation noise, $QNoise(b)$ does not depend on weights and does not require gradients approximation. For quantization aware search, each blue operation on Figure 1 becomes SAN operation with noisy weights. WR - real valued weights, WQ - pseudo quantized weights, α is a vector of trainable parameters. By adjusting α we can search for acceptable model degradation caused by quantization procedure.

Differentiable NAS (DNAS) [23, 17, 11, 24] is a differentiable method of selecting a directed acyclic sub-graph (DAG) from an over-parameterized supernet. An example of such selection is on Figure 1. Each node represents a feature map of intermediate layers of inputs and outputs. Edges

are operations between those nodes. During the search procedure, we aim to assign importance weights for each edge and consequently select a sub-graph using edges with the highest importance weights.

The weights assignment can be done in several ways. The main idea of DNAS is to update importance weights α with respect to a loss function parameterized on supernet weights W. Consequently, hardware constraints are easy to introduce as an extension of an initial loss function.

DNAS has been proven to be efficient to search for computationally-optimized models. FBnet [23] focuses on optimizing FLOPs and latency. Authors mainly focus on classification problems. AGD [11] and TrilevelNAS[24] further extend resource constrained NAS for super resolution problem (SR) for full precision models.

Quantization aware DNAS. DNAS can be employed to search for architectures with desired properties. In OQAT [19], authors perform quantization-aware NAS with uniform quantization. They show that architectures found with a quantization-aware search perform better when quantized compared to architectures found without accounting for quantization. However, uniform quantization is less flexible and leads to suboptimal solutions compared to mixed-precision quantization (MPQ) - where each operation and activation has its own quantization bits. This idea was explored in EdMIPS [5]. MPQ used in EdMIPS is a NAS procedure where all operations are fixed, and we search for different quantization levels. One on hand MPQ-NAS is a natural extension of EdMIPS or OQAT but on another hand joint optimization of quantization bits and operations has a high computational costs. Additionally, instability of optimization with quantized weights was highlighted in DiffQ [7], authors considers QN to perform MPQ on existing architectures.

Two problems above make joint optimization of an over-parameterized supernet with mixed-precision bits a challenging task. We propose a procedure that is simultaneously effective and relatively robust. Due to the usage of supernet, we turn our problem into a continuous one, and by the use of quantization noise, we can make the solution of this problem fast and stable.

Search space design is crucial for achieving good results. It should be both flexible and contain known best-performing solutions. Even a random search can be a reasonable method with a good search space design. In AGD [11] authors apply NAS for SR, they search for (1) a cell - a block which is repeated several times, and (2) kernel size along with other hyperparameters like the number of input and output channels. TrilevelNAS[24] extends the previous work by adding (3) network level that optimizes the position of the network upsampling layer. Both arti-

cles expand the search space, making it more flexible while more tricky to search in, possibly leading to local and sub-optimal solutions.

In our work, we choose to focus on a simpler search space consisting of: (1) a cell block - search is performed for operations within the block and (2) quantization levels different for each operation (for quantization-aware search). We show that this design leads to architectures with similar performance as TrilevelNAS[24] but search time is much faster.

Sparsification for differentiable architecture search was discussed in several works [3, 21, 6, 25, 24, 27]. The problem arises because operations in a supernet co-adapt. A selected from the supernet subgraph depends on all the left in the supernet operations. We can make the problem better suitable for NAS if we enforce the sparsification of a graph with most of the importances for a particular connection between nodes being close to zero.

The sparsification strategy depends on the graph structure of a final model. In our work, we use the Single-Path strategy - one possible edge between two nodes, more in appendix C. For the Single-Path strategy, the sum of node outputs is a weighted sum of features, it can be seen in Figure 1. The co-adaptation problem becomes obvious. Second layer convolutions are trained on a weighted sum of features, but after discretization (selecting a subgraph), only one source of features remains. Therefore, sparsification for the vector of α is necessary. In BATS [3], sparsification is achieved via scheduled temperature for softmax. Entropy regularization proposed in Discretization-Aware search [21]. In [6], authors proposed an ensemble of Gumbels to sample sparse architectures for the Mixed-Path strategy and in [25], Sparse Group Lasso (SGL) regularization is used. In ISTA-NAS [27], authors tackle sparsification as a sparse coding problem. Trilvel NAS [24] proposed sorted Sparsestmax. In our work, we used entropy regularization [21] to disentangle the final model from the supernet.

3. Methodology

We follow the basic definitions provided in the previous section with the description of our approach. The description has three parts. We start with (i) subsection 3.1 that describes the space of considered architectures. The procedure for searching of an architecture in this space is in (ii) subsection 3.2. It includes the description of the used loss function. Finally, in (iii) subsection 3.3 we provide details on reparametrization with quantization noise.

3.1. Search space design

We design our search space taking into account recent results in this area and, in particular, the SR quantization challenge [14]. The challenge was in manually designing of quantization-friendly architectures.

We combine most of these ideas in the search design depicted in Figure 3. The deterministic part of our search space includes the upsampling layer in tail block of the architecture, the number of channels in convolutions and specific for SR AdaDM [18] block. The AdaDM block is used only in quantization-aware search. The variable part is quantization bit values and operations within head, body, up-sample, and tail blocks. We perform all experiments with 3 body blocks, unless specified otherwise. Additional, parallel convolutional layer of the body block is used to increase representation power of quantized activations.

Batch Norm for SR and modification of AdaDM. Variation in a signal is crucial for identifying small details. In particular, the residual feature's standard deviation shrinks after the layers' normalisation, and SR models with BN perform worse [18]. On another hand training overparameterized supernet without BN can be challenging. The authors of [18] proposed to rescale the signal after BN based on its variation before BN layers. Empirically we found that AdaDM with removed second BN improves overall performance of quantized models. Original AdaDM block and our modification are depicted in Figure 4. All the residual blocks in our search design have the modified AdaDM part: the body block, all the repeated layers within the body block and the tail block.

3.2. The search procedure

We consider the selection of blocks and quantization bits during NAS. We assign separate α values to optimize during NAS for a Cartesian product of possible operations and the number of quantization bits. Search and training procedures are performed as two independent steps.

For search, we alternately update supernet's weights W and edges importances α. Two different subsets of training data are used to calculate the loss function and derivatives for updating W and α similar to [11]. Hardware constraints and entropy regularisation are applied as additional terms in the loss function for updating α. To calculate the output of l-th layer x_{l+1} we weight the output of separate edges according to importance values α_{ibl} of each of O^l operations and bits:

$$x_{l+1} = \sum_{i=1}^{|O^l|} \sum_{b=1}^{|B|} \alpha_{ibl} o_i^l(x_l), \qquad (1)$$

where $\sum_{i=1}^{|O^j|} \sum_{b=1}^{|B|} \alpha_{ibl} = 1$ and all $\alpha_{ibl} \geq 0$.

Then the final architecture is derived by choosing a single operator with the maximal α_j among the ones for this layer. Finally, we train the obtained architecture from scratch.

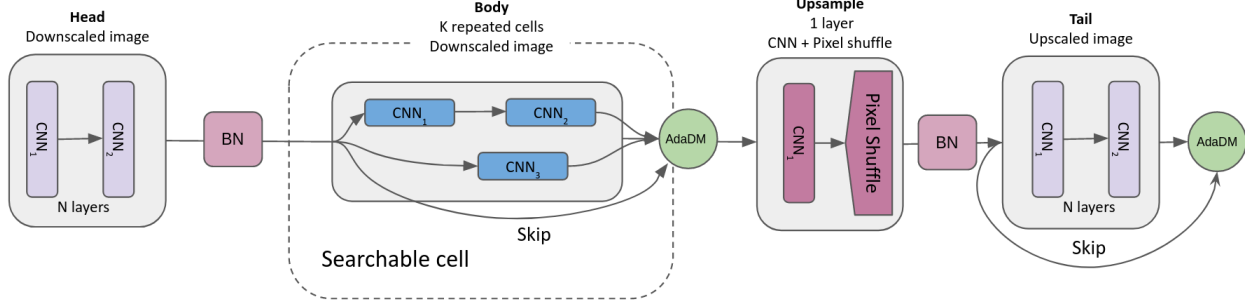To optimize α we compute the following loss:

Figure 3. The search space design. We separate the whole architecture into 4 parts: head, body, upsample, and tail. The head and the tail have $N = 2$ convolutional layers. The identical body part is repeated 3 times, unless specified otherwise. The number of channels for all the blocks equals 36, except for the head's first layer, upsample, and the tail's first layers. All the blocks with skip connections incorporate AdaDM with BN.
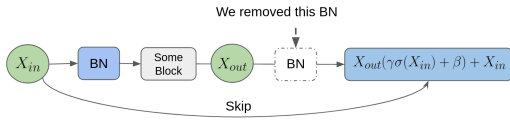

Figure 4. Our modification of AdaDM [18]. *Some Block* represents any residual block with several layers within, $\sigma(X_{in})$ is a variance of input signal, γ and β are learnable scalars. We remove the second BN after X_{out}.

$$L(\boldsymbol{\alpha}) = L_1(\boldsymbol{\alpha}) + \eta L_{cq}(\boldsymbol{\alpha}) + \mu(t)L_e(\boldsymbol{\alpha}),$$

where η and $\mu(t)$ are regularization constants. $\mu(t)$ depends on the iteration t. $L_1(\boldsymbol{\alpha})$ is l_1-distance between high resolution and restored images, and $L_{cq}(\boldsymbol{\alpha})$ is the hardware constraint and $L_e(\boldsymbol{\alpha})$ is the entropy loss that enforces sparsity of the vector $\boldsymbol{\alpha}$. The last two losses are defined in two subsections below.

3.2.1 Hardware constraint regularization

The hardware constraint is proportional to the number of floating point operations FLOPs for full precision models and the number of quantized operations BitOps for mixed-low precision models. $F_{fp}(o, x)$ is the function computing FLOPs value based on the input image size x and the properties of a convolutional layer o: kernel size, number of channels, stride, and the number of groups. We use the same number of bits for weights and activations in our setup. Therefore, BitOps can be computed as $F_q(o, x) = b^2 F_{fp}(o, x)$ where b is the number of bits. Then, the corresponding hardware part of the loss L_{cq} is:

$$L_{cq}(\boldsymbol{\alpha}) = \sum_{j=1}^{|S|} \sum_{i=1}^{|O^j|} \sum_{b=1}^{|B|} \alpha_{jib} b^2 F_{fp}(o_i^j, x_j),$$

$$\text{where} \quad \sum_{i=1}^{|O^j|} \sum_{b=1}^{|B|} \alpha_{jib} = 1 \quad \text{and} \quad \forall \quad \alpha_{jib} \geq 0, \quad (2)$$

where S is a supernet's block or layer consisting of several operations, the layer-wise structure is presented in Figure 1. We normalize $L_{cq}(\boldsymbol{\alpha})$ value by the value of this loss at initialization with the uniform assignment of α, as the scale of the unnormalized hardware constraint reaches 10^{12}.

3.2.2 Sparsity regularization

After the architecture search, the model keeps only one edge between two nodes. Let us denote as $\boldsymbol{\alpha}_l$ all alphas that correspond to edges that connect a particular pair of nodes. They include different operations and different bits. At the end of the search, we want $\boldsymbol{\alpha}_l$ to be a vector close to the vector with one value close to 1 and all remaining components to be 0. We found that the entropy loss works the best with our settings.

The sparsification loss $L_e(\boldsymbol{\alpha})$ for $\boldsymbol{\alpha}$ update step has the following form:

$$L_e(\boldsymbol{\alpha}) = \sum_{l=1}^{|S|} H(\boldsymbol{\alpha}_l), \qquad (3)$$

where H is the entropy function. The coefficient before this loss $\mu(t)$ depends on the training epoch t. The detailed new procedure for regularization scheduling is given in Appendix A.

3.3. Quantization

Our aim is to find quantization-friendly architectures that perform well after quantization. (i) To obtain a trained and quantized model, we perform Quantization Aware Training QAT [15]. During training, (a) we approximate quantization with QN; (b) compute gradients for quantized weights; (c) update full precision weights. (ii) Then, a model with found architecture is trained from scratch. Below we provide details for these steps.

3.3.1 Quantization Aware Training

Let's consider the following one-layer neural network (NN),

$$y = f(a(x)) = Wa(x), \tag{4}$$

where a is a non linear activation function and f is a function parametrized by a tensor W. In (4) f is a linear function, but it also can be a convolutional operation. To decrease the computational complexity of the network, we replace expensive float-point operations with low-bit width ones. Quantization occurs for both weights W and activation a.

The quantized output has the following form:

$$y_q = f_q(a_q(x)) = o(G(a(x), b), Q(W, b)). \tag{5}$$

$Q(W, b)$ is a quantization function for weights, we use Learned Step Quantization (LSQ) [10] with learnable step value. To quantize activations we use $G(W, b)$ a half wave Gaussian quantization function [4]. Where, quantization level is b and a convolution layer is o.

3.3.2 Quantization Aware Search with Shared Weights (SW)

To account for further quantization during the search phase, we perform Quantization Aware Search similar to QAT [15]. One way to do so is to quantize model weights and activations during the search phase the same way as during training.

To improve computational efficiency we can quantize weights of identical operations with different quantization bits instead of using different weights for each quantization bit, this idea was studied in [5]. Here, x_{j+1} is the output of j-th layer with input x_j and parameters W_j.

$$x_{j+1} = \sum_{i=1}^{|O^j|} o_{ij} \left(\sum_{b \in B} \alpha_{jib} G(a(x_j), b), \sum_{b \in B} \alpha_{jib} Q(W_{ji}, b) \right). \tag{6}$$

$$\sum_{i=1}^{|O^j|} \sum_{b \in B} \alpha_{jib} = 1, \forall i, j, b \; \alpha_{ijb} \geq 0.$$

The effectiveness of SW can be seen from (6): it requires fewer convolutional operations and less memory to store the weights.

3.3.3 Quantization Aware Search Against Noise

To further improve computational efficiency and performance of search phase we introduce SAN. Model degradation caused by weights quantization is equivalent to adding the quantization noise $QNoise_b(W) = Q(W, b) - W$.

Then, quantized weights is $Q(W, b) = W + QNoise_b(W)$ and (6) is:

$$x_{j+1} = \sum_{i=1}^{|O^j|} o_{ij} \Big(\sum_{b \in B} \alpha_{jib} QNoise_b(a(x_j)) + a(x_j), \\ \sum_{b \in B} \alpha_{jib} QNoise_b(W_{ji}) + W_{ji} \Big). \tag{7}$$

This procedure (i) does not require weights quantization and (ii) differentiable unlike quantization. $QNoise_b$ is a function of W because it depends on its shape and magnitude of values. Given the quantization noise, we can more efficiently run forward and backward passes for our network, similar to the reparametrization trick.

Adding quantization noise is similar to adding independent uniform variables from $[-\Delta/2, \Delta/2]$ with $\Delta = \frac{1}{2^b - 1}$ [22]. However, for the noise sampling, we use the following procedure [7]:

$$QNoise(b) = \frac{\Delta}{2} z, z \sim \mathcal{N}(0, 1), \tag{8}$$

as it performs slightly better than the uniform distribution [7].
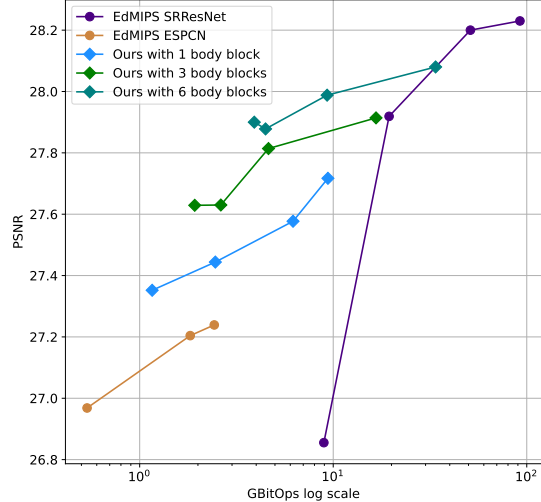

Figure 5. Our quantization aware NAS approach and fixed quantized architectures. Ours - QuantNAS with AdaDM and SAN with following hardware penalties: 0, 1e-4, 1e-3, 5e-5 and a different number of body blocks. Mixed precision quantization by EdMIPS [5] for SRResNet [16] and ESPCN [20] with following hardware penalties: 0, 1e-3, 1e-2, 1e-1. PSNR was computed on Set14 and BitOPs for image size 32x32. Our search procedure found a significantly more efficient architecture compared to manually designed and then quantized ESPCN within the same BitOPs range.
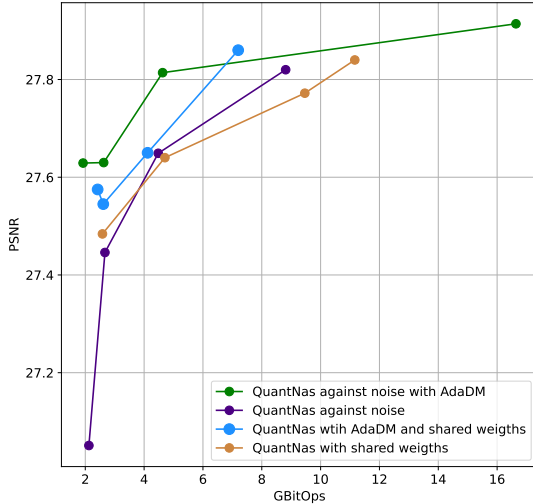
Figure 6. QuantNAS, models found with different hardware penalty values: 0, 1e-4, 1e-3, 5e-5. QuantNAS with AdaDM in blue, SAN with AdaDM in green, SAN without AdaDM and without BN in purple, shared weights procedure without AdaDM and without BN in brown. Results are presented for Set14.

4. Results

We provide the code for our experiments here.

4.1. Evaluation protocol

For all experiments we consider the following setup if not stated otherwise. A number of body blocks is set to 3. As the training dataset we use DIV2K dataset [1]. As the test datasets we use Set14 [28], Set5 [2], Urban1100 [13], Manga109 [12], with scale 4. In the main body of the paper we present results on Set14. The results for other datasets are presented in Appendix. For training, we use RGB images. For PSNR score calculation, we use only the Y channel similarly to [11, 24]. Evaluation of FLOPs and BitOPs is done for fixed image sizes 256×256 and 32×32, respectively.

Search space. For full precision search, we use ten different possible operations as candidates for a connection between two nodes. For quantization-aware search, we limit the number of operations to 4 to obtain a search space of a reasonable size. For quantization, we consider two options as possible quantization bits: 4 or 8 bits for activations and weights.

4.2. Different number of body blocks

A straightforward way to improve model performance is to increase the number of layers. We study how our method

Model	GBitOPs	PSNR	Method
SRResNet	23.3	27.88	LSQ(4-U)
SRResNet	23.3	27.42	HWGQ(4-U)
ESPCN	2.3	27.26	LSQ(8-U)
ESPCN	2.3	27.00	HWGQ(8-U)
SRResNet	19.4	27.919	EdMIPS(4,8)
Our(**body3**)	4.6	27.814	QuantNAS(4,8)
Our(**body6**)	9.3	27.988	QuantNAS(4,8)

Table 1. Quantitative results for different quantization methods for different models. For EdMIPS and QuantNAS, we present models found with different hardware penalties 1e-3 and 1e-4, respectively. "U" - stands for uniform quantization - all bits are the same for all layers. BitOPs were computed for 32x32 image size. Our (body 3) quantized architecture with body block repeated 3 times is depicted in Figure 15.

scales by performing search with a different number of body blocks: 1, 3, and 6. Three constellations are presented in Figure 5. We observe that increasing the number of blocks improves final performance and increases the number of BitOPs for architectures found with the highest hardware regularization - each constellation is slightly shifted to the right.

4.3. QuantNAS and quantization of fixed architectures

We compare QuantNAS with (1) uniform quantization and (2) mixed precision quantization for two existing architectures, ESPCN [20] and SRResNet [16]. **For uniform quantization**, we use LSQ [10] and HWGQ [4]. **For mixed precision quantization**, we use EdMIPS [5]. Our setup for EdMIPS is matching the original setup and search is performed for different quantization bits for weights and activations. Unlike in QuantNAS, quantization bits for activations and weights are the same.

In Figure 5, we compare our procedure with two architectures quantized by EdMIPS. We can see that QuantNAS finds architectures with better PSNR/BitOps trade-offs within a range where BitOps values overlap. Performance gain is especially notable between quantized ESPCN and our approach.

We note that due to computational limits, our search is bounded in terms of the number of layers. Therefore, we can't extend our results beyond SRResNet in terms of BitOps to provide a more detailed comparison.

From Table 1, we can see that QuantNAS finds architectures with a better PSNR/BitOps trade-off than uniform quantization techniques. We compare within the same range of BitOPs values, 8 bits for ESPCN and 4 bits for SRResNet.
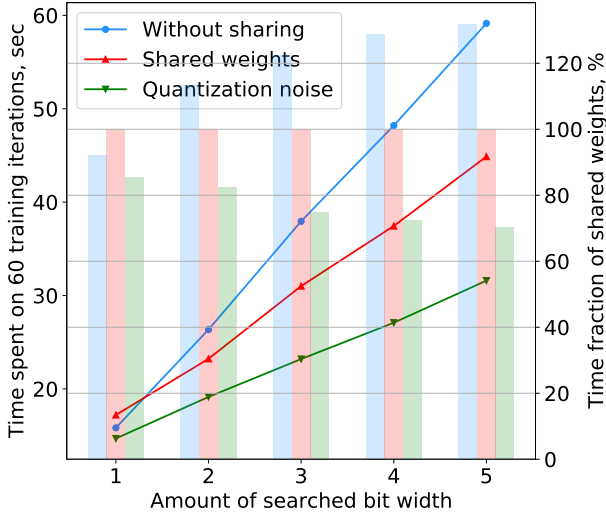
Figure 7. Time comparison of quantization noise and weights sharing strategy during the search phase of quantization-aware NAS. Y-axis (on the left) shows time spent on 60 training iterations (line plot). The secondary Y-axis (on the right) presents the time fraction of SW strategy (bar plot). SAN provides speed-up.

4.4. Time efficiency of QuantNAS

We measure the average training time for three considered quantisation approaches: without sharing weights, with sharing (SW), and with quantization noise (ours QuantNAS or briefly SAN).

We run the same experiment for different amounts of searched quantization bits. Figure 7 shows the advantage of our approach in training time. As the number of searched bits grows, so does the advantage. On average, we get up to 30% speedup.

4.5. Ablation studies

4.5.1 Adaptive Deviation Modulation

We start with comparing the effect of AdaDM [18] and Batch Normalization on two architectures randomly sampled from our search space. In Table 2, we can see that both original AdaDM and Batch Normalization hurt the final performance, while AdaDM with our modification improves PSNR scores. In Figure 6, we observe that architectures found with AdaDM are better in terms of both PSNR and BitOPs. Interestingly, we did not notice any improvement with AdaDM for full-precision models. Our best full precision model in Table 2 was obtained without AdaDM.

4.5.2 Entropy regularization

We consider three settings to compare QuantNas with and without Entropy regularization: (A) reduced search space, SGD optimizer; (B) full search space, Adam [8] optimizer;

Model	Model M1	Model M2
Without Batch Norm	27.55	28
With Batch Norm	27	27.16
Original AdaDM	27.33	27.84
Our AdaDM	**27**.**68**	**28**.**046**

Table 2. PSNR of SR models with scaling factor 4 for Set14 dataset. M1 and M2 are two arbitrary mixed precision models randomly sampled from our search space.

Training settings	w/o Entropy	w Entropy
A	27.99 / 111	**28**.**10** / 206
B	28.00 / 30	**28**.**12** / 19
C	27.92 / 61	**28**.**11** / 321

Table 3. PSNR/GFFLOPs values of search procedure with and without Entropy regularisation. Models were searched in different settings A, B, and C.

(C) reduced search space, Adam [8] optimizer. All the experiments were performed for full precision search. For full and reduced search spaces, we refer to Appendix C.1. We perform the search without hardware penalty to analyze the effect of the entropy penalty.

Quantitative results for Entropy regularisation are in Table 3. Entropy regularisation improves performance in terms of PSNR for all the experiments.

Figure 8 demonstrates dynamics of operations importance for joint NAS with quantization for 4 and 8 bits. 4 bits edges are depicted in dashed lines. Only two layers are depicted: the first layer for the head (HEAD) block and the skip (SKIP) layer for the body block. With entropy regularization, the most important block is evident from its $alpha$ value. Without entropy regularization, we have no clear most important block. So, our search procedure has two properties: (a) the input to the following layer is mostly produced as the output of a single operation from the previous layer; (b) an architecture at final search epochs is very close to the architecture obtained after selecting only one operation per layer with the highest importance value.

4.5.3 Comparison with existing full precision NAS for SR approaches

Here we examine the quality of our procedure for full precision NAS. We did not use the AdaDM block for the full precision search.

Method	GFLOPs	PSNR	Search cost
SRResNet	166.0	28.49	Manual
AGD*	140.0	28.40	1.8 GPU days
Trilevel NAS*	33.3	28.26	6 GPU days
Our FP best	**29**.**3**	28.22	1.2 GPU days

Table 4. Quantitative results of PSNR-oriented models with SR scaling factor 4 for Set14 dataset. * results are from paper [24]
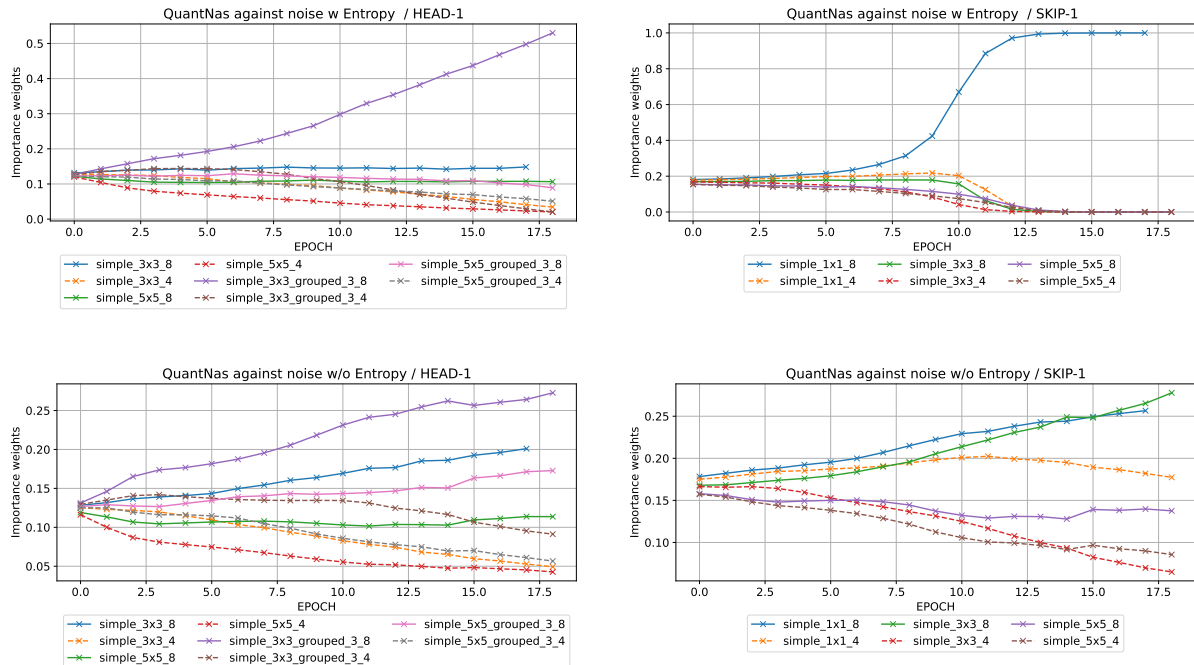
Figure 8. Importance weights for different operations through epochs for QuantNas search, for 8 and 4 bits, solid and dashed lines respectively. Supernet sparsification with entropy regularisation on the top for two layers HEAD-1 and a parallel conv layer in BODY (SKIP-1), regularisation value is set to 1e-3 (on the top). Training the supernet without sparsification with entropy regularization at the bottom.

The results are in Table 4. Our search procedure achieves comparable results with TrilevelNAS[24] with a relatively simpler search design and about 5 times faster training. The best performing full precision architecture was found with a hardware penalty of value $1e − 3$. This architecture is depicted in Appendix Figure 14.

Additionally, we compare results with a popular SR architecture SRResNet [16]. Visual examples of the obtained super-resolution pictures are presented in Appendix Figure 16 for Set14 [28], Set5 [2], Urban1100 [13], Manga109 [12] with scale factor 4.

5. Limitations

For our NAS procedure, the overall NAS limitation applies: the computational demand for joint optimization of many architectures is high. The search procedure takes about 24 hours to finish for a single GPU TITAN RTX. Moreover, obtaining the full Pareto frontier requires running the same experiment multiple times.

On Figure 6 all most right points (within one experiment/color) have 0 hardware penalty. It clearly shows that limited search space creates an upper bound for the top model performance. Therefore, our results do not fall within the same BitOps range as SRResNet.

We found that our search design is sensitive to hyper-parameters. In particular, optimal coefficients for hardware penalty and entropy regularization can vary across different search settings. Moreover, we expect that there is a connection between optimal coefficients for the hardware penalty, entropy regularization and search space size. Different strategies or search settings require different values of hardware penalties. Applying the same set of values for different settings might not be the best option, but it is not straightforward how to determine them beforehand.

6. Conclusion

To the best of our knowledge, we are the first to deeply explore NAS with mixed-precision search for Super-Resolution (SR) tasks.

We proposed the method QuantNAS for obtaining computationally efficient and accurate architectures for SR using jointly NAS and mixed-precision quantization.

Our method is better than others due to (1) specifically tailored search space design; (2) differentiable SAN procedure; (3) adaptation of AdaDM; (4) the entropy regularization to avoid co-adaptation in super-nets during differentiable search.

Experiments on standard SR tasks demonstrate the high quality of our search. Our method leads to better solutions compared to mixed-precision quantization of popular SR architectures with [5]. Moreover, our search is up to 30% faster than a share-weights approach.

References

[1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 126–135, 2017.

[2] Marco Bevilacqua, Aline Roumy, Christine Guillemot, and Marie Line Alberi-Morel. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. 2012.

[3] Adrian Bulat, Brais Martinez, and Georgios Tzimiropoulos. Bats: Binary architecture search. *ECCV2020*, 2020.

[4] Zhaowei Cai, Xiaodong He, Jian Sun, and Nuno Vasconcelos. Deep learning with low precision by half-wave gaussian quantization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5918–5926, 2017.

[5] Zhaowei Cai and Nuno Vasconcelos. Rethinking differentiable search for mixed-precision neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2349–2358, 2020.

[6] Jianlong Chang, Yiwen Guo, Gaofeng Meng, Shiming Xiang, Chunhong Pan, et al. Data: Differentiable architecture approximation. *Conference on Neural Information Processing Systems*, 2019.

[7] Alexandre Défossez, Yossi Adi, and Gabriel Synnaeve. Differentiable model compression via pseudo quantization noise. *arXiv:2104.09987v2*, 2021.

[8] Jimmy Ba Diederik P. Kingma. Adam: A method for stochastic optimization. In *In Proceedings of the International Conference on Learning Representations*, pages 126–135, 2015.

[9] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. *ECCV*, 2014.

[10] Steven K Esser, Jeffrey L McKinstry, Deepika Bablani, Rathinakumar Appuswamy, and Dharmendra S Modha. Learned step size quantization. *arXiv preprint arXiv:1902.08153*, 2019.

[11] Yonggan Fu, Wuyang Chen, Haotao Wang, Haoran Li, Yingyan Lin, and Zhangyang Wang. Autogandistiller: Searching to compress generative adversarial networks. *ICML*, 2020.

[12] Azuma Fujimoto, Toru Ogawa, Kazuyoshi Yamamoto, Yusuke Matsui, Toshihiko Yamasaki, and Kiyoharu Aizawa. Manga109 dataset and creation of metadata. In *Proceedings of the 1st international workshop on comics analysis, processing and understanding*, pages 1–5, 2016.

[13] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5197–5206, 2015.

[14] Andrey Ignatov, Radu Timofte, Maurizio Denna, and Abdel Younes. Real-time quantized image super-resolution on mobile npus, mobile ai 2021 challenge: Report. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2525–2534, 2021.

[15] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. *In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2704–2713*, 2018.

[16] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image superresolution using a generative adversarial network. *In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 4681–4690*, 2017.

[17] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *ICLR*, 2019.

[18] Jie Liu, Jie Tang, and Gangshan Wu. Adadm: Enabling normalization for image super-resolution. *arXiv preprint arXiv:2111.13905*, 2021.

[19] Mingzhu Shen, Feng Liang, Chuming Li, Chen Lin, Ming Sun, Junjie Yan, and Wanli Ouyang. Once quantized for all: Progressively searching for quantized efficient models. *arXiv preprint arXiv:2010.04354*, 2020.

[20] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. *In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 1874–1883*, 2016.

[21] Yunjie Tian, Chang Liu, Lingxi Xie, Qixiang Ye, et al. Discretization-aware architecture search. *Pattern Recognition*, 2021.

[22] Bernard Widrow, Istvan Kollar, and Ming-Chang Liu. Statistical theory of quantization. *IEEE Transactions on instrumentation and measurement, 45(2): 353–361.*, 1996.

[23] Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Pe-

ter Vajda, Yangqing Jia, and Kurt Keutzer. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.

[24] Yan Wu, Zhiwu Huang, Suryansh Kumar, Rhea Sanjay Sukthanker, Radu Timofte, and Luc Van Gool. Trilevel neural architecture search for efficient single image super-resolution. *Computer Vision and Pattern Recognition*, 2021.

[25] Yan Wu, Aoming Liu, Zhiwu Huang, Siwei Zhang, and Luc Van Gool. Neural architecture search as sparse supernet. *arXiv:2007.16112*, 2020.

[26] Wenming Yang, Xuechen Zhang, Yapeng Tian, Wei Wang, Jing-Hao Xue, and Qingmin Liao. Deep learning for single image super-resolution: A brief review. *IEEE Transactions on Multimedia*, 21(12):3106–3121, 2019.

[27] Yibo Yang, Hongyang Li, Shan You, Fei Wang, Chen Qian, and Zhouchen Lin. Ista-nas: Efficient and consistent neural architecture search by sparse coding. *Conference on Neural Information Processing Systems*, 2020.

[28] Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In *International conference on curves and surfaces*, pages 711–730. Springer, 2010.

Appendix

A. Technical details

During the search phase, we consider architectures with a fixed number of 36 channels for each operation unless channel size is changed due to operations properties. The search is performed for 20 epochs. For updating supernet weights we use the following hyper-parameters: batch size 16, initial learning rate (lr) 1e-3, cosine learning rate scheduler, SGD with momentum 0.9, weight decay 3e-7. For updating alphas we use fixed lr 3e-4 and 0 weight decay.

During the training phase, an obtained architecture is optimized for 30 epochs with the following setting: batch size 16, initial lr 1e-3, lr scheduler with the weight decay 3e-7.

A.1. Entropy schedule

For entropy regularization, we gradually increase the regularization value α according to Figure 9, for the first two epochs regularization is zero. Entropy regularization is multiplied by an initial coefficient and coefficient factor. Initial coefficients are 1e-3 and 1e-4 for experiments with full precision and the quantization-aware search.
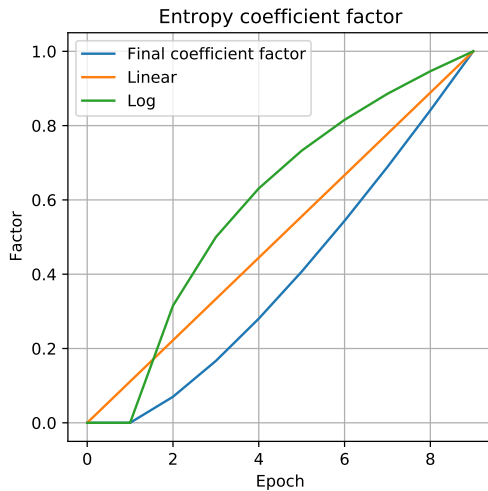

Figure 9. Entropy coefficient regularization is a product of log and linear functions

B. Scaling SR models with initial up-sampling

To maintain good computational efficiency, it is common for recent SR models to operate on down-sampled images and then up-sample them with some up-sampling layers. This idea was introduced first in ESPCN [20]. Since then, there were not many works in the literature exploring SR models on initially up-scaled images.

Therefore, we were interested in how this approach scales in terms of quality and computational efficiency
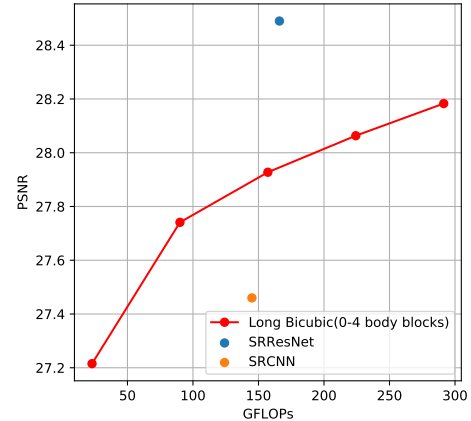

Figure 10. Orange point is the original SRCNN [9], blue point is SRResNet [16]. For Long Bicubic we initially upscale an image with bicubic interpolation and then add an efficient block found in our experiments. The block consists of 3 convolutions layers with 32 filters and it is added 1, 2,3,4 times. PSNR is reported on Set14.

given arbitrary many layers. Results are presented in Figure 10. We start with one fixed block, similar to our body block in Figure 3 and then increase it by one each time. We compare our results with SRResNet [16] and SRCNN [9]. As we can see SRResNet [16] operates on down-scaled images and yields better results given the same computational complexity.

C. Single-path search space

There are several ways to select directed acyclic subgraph from a supernet. DARTS [17] uses Multi-Path strategy - one node can have several input edges. Such a strategy makes a search space significantly larger. Single-Path strategy - each searchable layer in the network can choose only one operation from the layer-wise search space Figure 1. It has been shown in FBNet [23] that simpler Single-Path approach yields comparable with Multi-Path approach results for classification problems. Also, since it aligns more with SR search design in our work, we use Single-Path approach.

C.1. Different search spaces

For detailed operations description we refer to our code.

D. Search space

We have a fixed number of channels equal to 36 for all the layers unless specified.

- **Head** : two layers, first layer computes images with three channels;
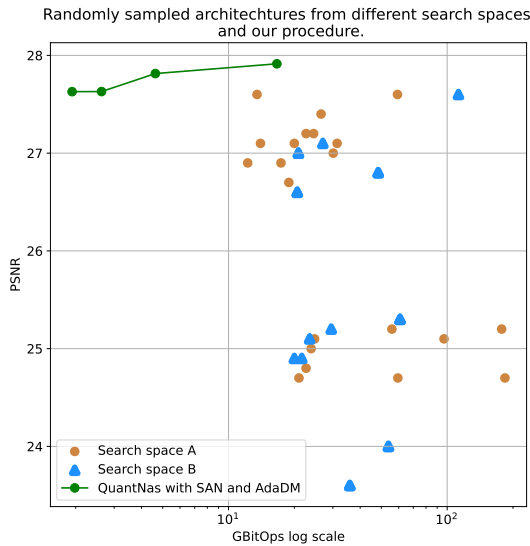
Figure 11. Randomly sampled architectures from two search spaces. PSNR was computed on Set14 and BitOPs for image size 256x256. Search space A is described in section D.2. We observe that two search spaces provide with slightly different results with random sampling. Results in green for architecture search were obtained with Search space A.
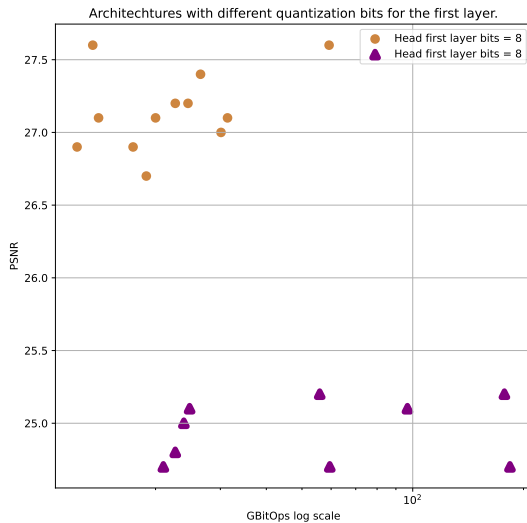

Figure 12. Influence of the first layer quantization on final performance. First layer low bit quantization significantly hurts model performance. Two clusters above and below 25.5 PSNR value attribute to different quantization levels of the first layer.

- **Body**: a repeatable cell with a skip connection and three layers;

- **Skip**: a single convolutional layer within body cell connecting input and output (do not confuse with conventional skip operation);

- **Upsample**: one layer before pixel shuffle, this layer increases a number of channels before pixel shuffle and pixel shuffle operation outputs 3 channels image;

- **Tail** : two layers, final layer outputs images with 3 channels.

D.1. Search space for full precision experiments

Possible operations block-wise:

- **Head** 8 operations: simple 3x3, simple 5x5, growth2 5x5, growth2 3x3, simple 3x3 grouped 3, simple 5x5 grouped 3, simple 1x1 grouped 3, simple 1x1;

- **Body** 7 operations: simple 3x3, simple 5x5, simple 3x3 grouped 3, simple 5x5 grouped 3, decenc 3x3 2, decenc 5x5 2, simple 1x1 grouped 3;

- **Skip** 4 operations: decenc 3x3 2, decenc 5x5 2, simple 3x3, simple 5x5;

- **Upsample** 12 operations: conv 5x1 1x5, conv 3x1 1x3, simple 3x3, simple 5x5, growth2 5x5, growth2 3x3, decenc 3x3 2, decenc 5x5 2, simple 3x3 grouped 3, simple 5x5 grouped 3, simple 1x1 grouped 3, simple 1x1;

- **Tail** 8 operations: simple 3x3, simple 5x5, growth2 5x5, growth2 3x3, simple 3x3 grouped 3, simple 5x5 grouped 3, simple 1x1 grouped 3, simple 1x1;

D.2. Reduced search space for Quantization experiments

Possible operations block-wise:

- **Head** 5 operations: simple 3x3, simple 5x5, simple 3x3 grouped 3, simple 5x5 grouped 3;

- **Body** 4 operations: conv 5x1 1x5, conv 3x1 1x3, simple 3x3, simple 5x5;

- **Skip** 3 operations: simple 1x1, simple 3x3, simple 5x5;

- **Upsample** 4 operations: conv 5x1 1x5, conv 3x1 1x3, simple 3x3, simple 5x5;

- **Tail** 3 operations: simple 1x1, simple 3x3, simple 5x5;

Conv 5x1 1x5 and conv 3x1 1x3 are depth-wise separable convolution convolutions. For operations description we refer to our code.

E. Results on other datasets

In Figure 13 we provide quantative results obtained on differetn test datasets: Set14 [28], Set5 [2], Urban1100 [13], Manga109 [12] with scale 4.

In Figure 16 we provide with visual results for quantized and full precision models.

F. QuantNAS algorithm

Algorithm 1 QuantNAS steps

1: Initialize parameters W and edge values α
2: **for** $iteration = 1, 2, \ldots, N$ **do**
3:      Add QN to W as in 8 and 7
4:      Compute the loss function $L(\alpha)$ as in 3.2
5:      Run backpropagation to get derivatives for α
6:      Update α
7:      Add QN to W as in 8 and 7
8:      Compute the loss function $L_1(W)$
9:      Run backpropagation to get derivatives for W
10:      Update W
11: **end for**
12: Select edges with the highest α
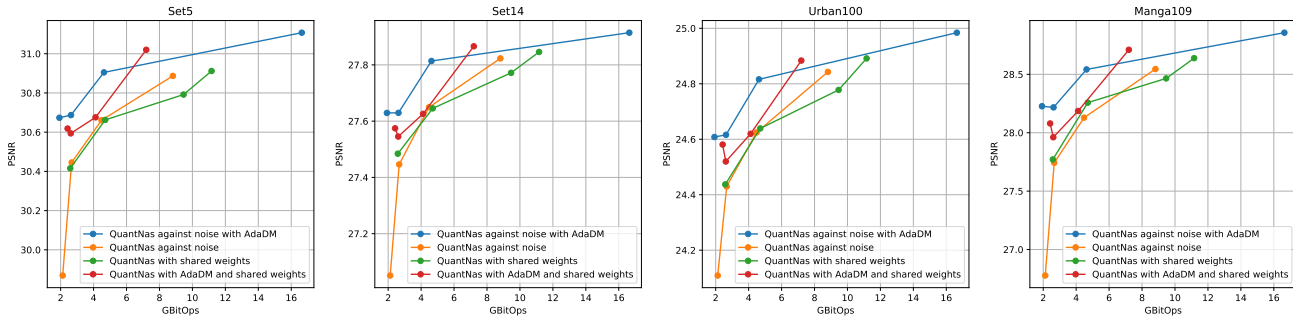13: Train the final architecture from scratch

---

Figure 13. The same as Figure 6 but for different datasets. QuantNAS, models found with different hardware penalty values: 0, 1e-4, 1e-3, 5e-5. QuantNAS with AdaDM in blue, SAN with AdaDM in green and original procedure with shared weights in brown. Better view in zoom.
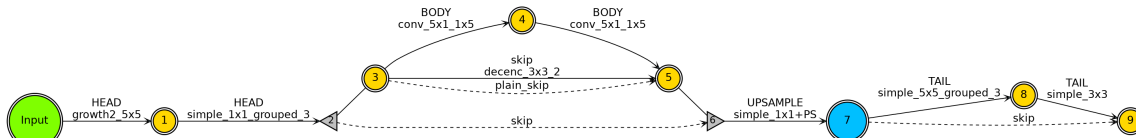

Figure 14. Our best FP (full precision) architecture, 29.3 GFLOPs (image size 265x265), PSNR: 28.22 dB. PSNR was computed on Set14. Body block is repeated three times for the both architectures.
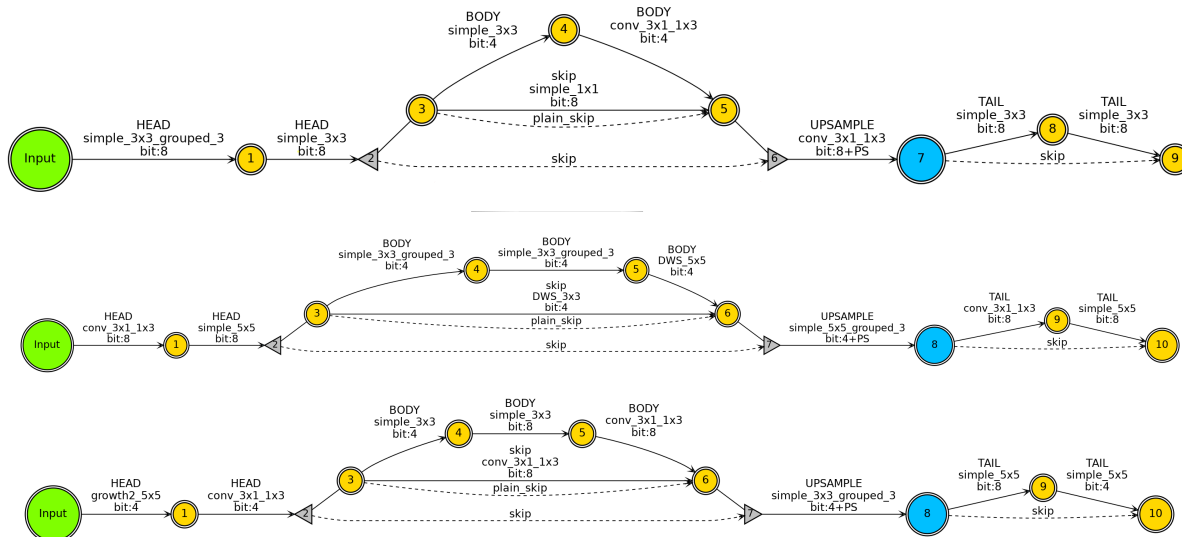

Figure 15. Examples of quantized architechtures. PSNR from top to bottom: 27.814 dB, 27.2 db, 24.8 db. On the top our (body 3) quantized architecture, more details are given in Table 1. PSNR was computed on Set14 with scale 4. Body block is repeated three times for all the architectures. Architecture on the bottom were sampled randomly.
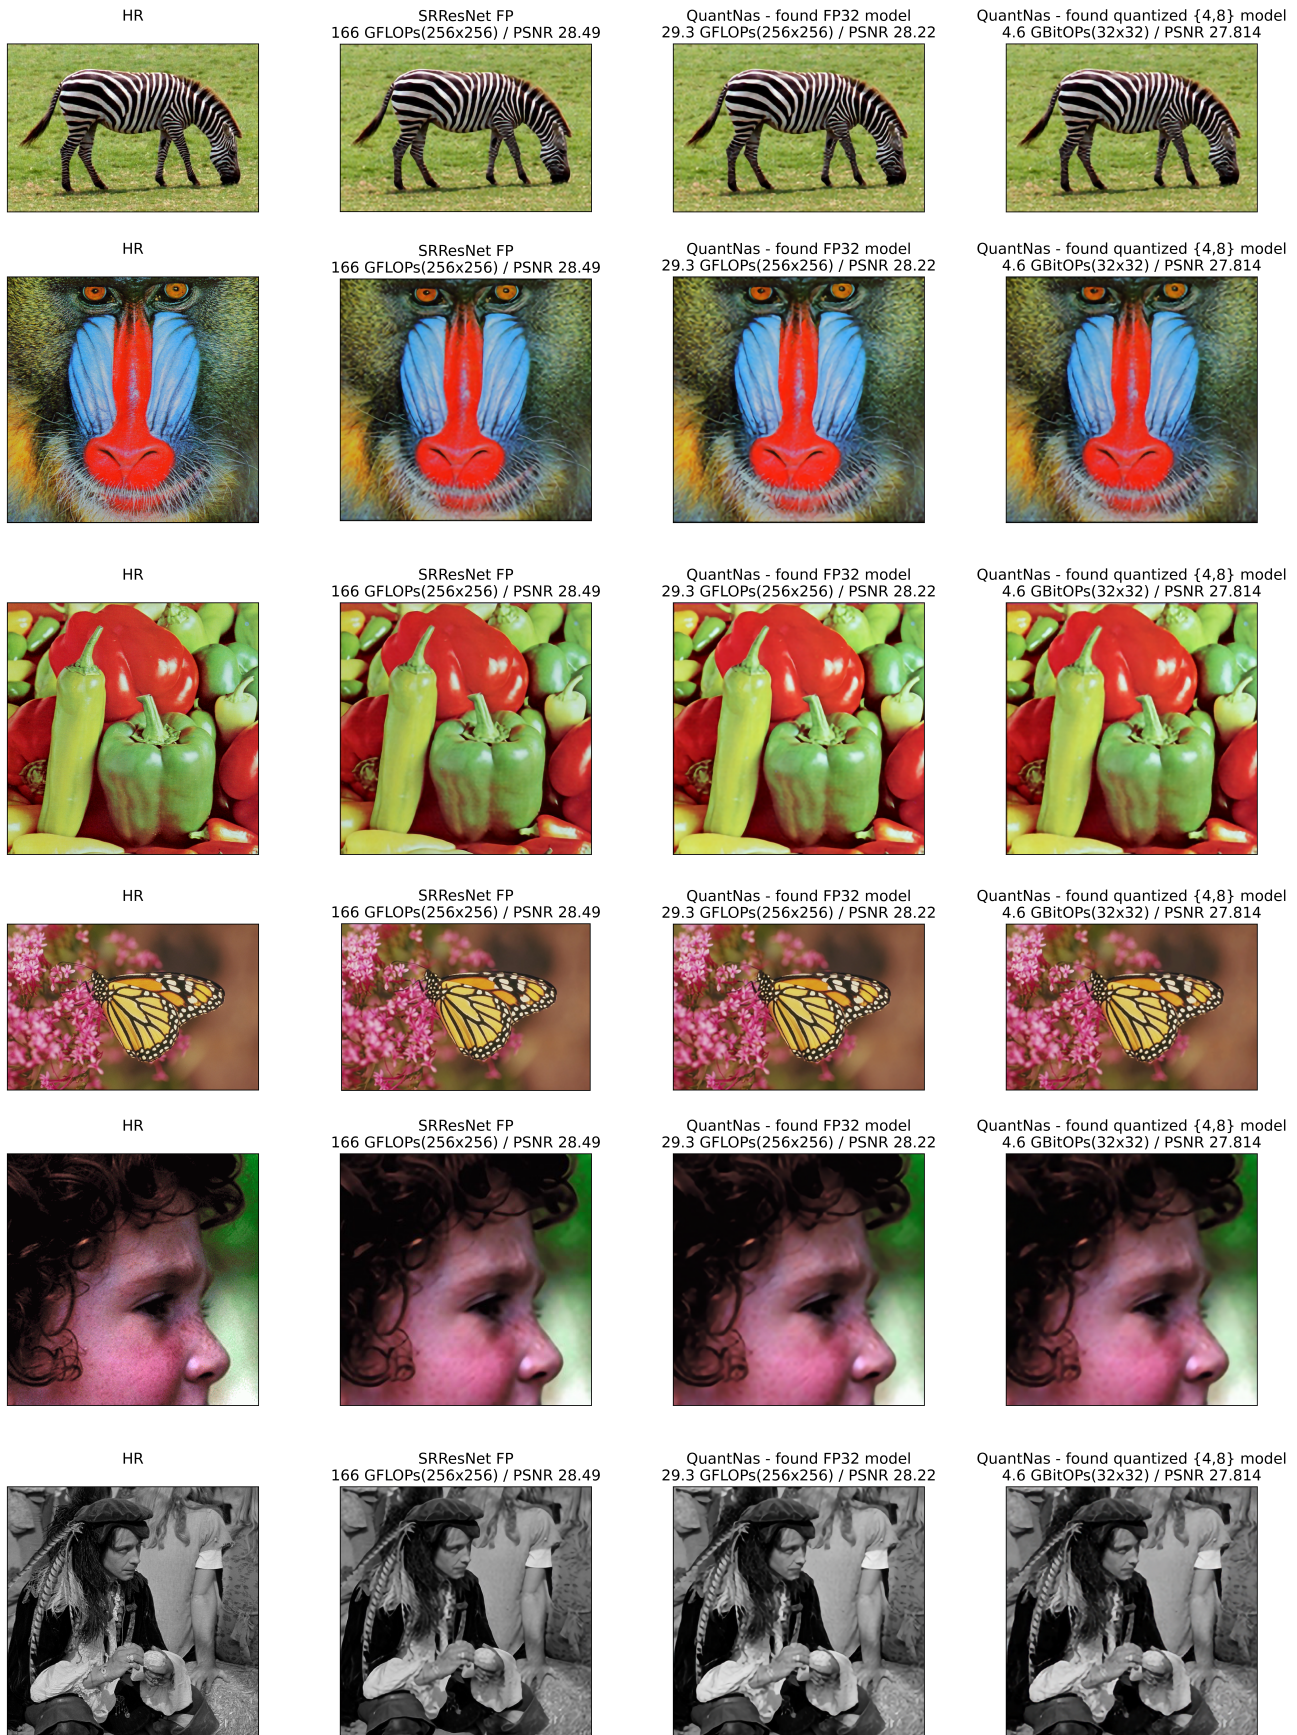
Figure 16. Visual comparison of results for Set14. Better view in zoom. Note, we present results for quantized models with the body block repeated 3 times. Model with with the body block repeated 6 times has better PSNR values, see in Table 1