

# Pushing Joint Image Denoising and Classification to the Edge

Thomas C. Markhorst<sup>\*</sup>, Jan C. Van Gemert, and Osman S. Kayhan

CV Lab, Delft University of Technology  
{t.c.markhorst}@tudelft.nl

**Abstract.** In this paper, we jointly combine image classification and image denoising, aiming to enhance human perception of noisy images captured by edge devices, like low-light security cameras. In such settings, it is important to retain the ability of humans to verify the automatic classification decision and thus jointly denoise the image to enhance human perception. Since edge devices have little computational power, we explicitly optimize for efficiency by proposing a novel architecture that integrates the two tasks. Additionally, we alter a Neural Architecture Search (NAS) method, which searches for classifiers to search for the integrated model while optimizing for a target latency, classification accuracy, and denoising performance. The NAS architectures outperform our manually designed alternatives in both denoising and classification, offering a significant improvement to human perception. Our approach empowers users to construct architectures tailored to domains like medical imaging, surveillance systems, and industrial inspections.

**Keywords:** Image Denoising · Neural Architecture Search · Image Classification

## 1 Introduction

The intersection of edge devices, such as security cameras, and deep learning has sparked an interest in optimizing neural networks for inference time, further referred to as latency. Common tasks to optimize for such efficiency are object classification and object detection, which unlock automatic recognition. However, in noisy settings, the recognition accuracy might not be perfect and it is important to allow the ability to validate the automatic recognition by human inspection. Thus, in addition to automatic recognition, the perceptual quality of the processed image is equally significant. In particular, this is relevant for images containing noise, which can arise from various sources such as low-light conditions, sensor noise, or other recording conditions. We focus on using an efficient model that can be used on the edge with the aim of enhancing human perception for validating the recognition output of noisy images.

Domains relying on human image perception but challenged by noisy images, like medical imaging [25], surveillance systems [34], and industrial in-



**Fig. 1:** We take a noisy image (a), which can be interpreted as an animal (b) or human (c). We denoise and classify the image (a), aiming to improve human perception resulting in (d). Note, in a real application (b) and (c) would not be available, which increases the difficulty of interpreting the noisy image. Artist: DALL-E-2 [30]

spectations [10], can benefit from recently proposed denoising Convolutional Neural Networks (CNNs) [14, 46]. As CNNs denoise better than traditional methods [7, 13]. Fast CNN denoisers [12, 47] are required to accommodate the real-time requirement of the affected domains. However, denoisers are not able to remove all noise, which is not always enough for human image perception.

We further improve human understanding of the image by combining denoising with machine perception, like image classification. From the Human-Computer Cooperation strategies in [40], we use the Classifier as Aid to Human. Where the image classifier can be used as a direct advisor or an independent agent to the security guard, the latter being analogous to a second opinion in medical diagnosis. In different use cases, fusing skills of humans and computers has been shown to improve performance beyond using only one type [2, 23, 40]. Therefore, we investigate models that can leverage the benefits of both denoising and classification to enhance human understanding in real time.

A model combining both denoising and classification is studied in [20], focusing on denoising performance. In addition, we optimize for efficiency, which is required for edge devices, and classification. Our efficiency definition is based on two elements: (i) latency reduction while (ii) retaining denoising performance and classification accuracy. These elements could be optimized using independent classification and denoising models. However, we propose an architecture combining the tasks more efficiently.

First, we employ established model design approaches to enhance independent denoising and classification models, such as model scaling [24, 36] and efficient operators [32]. Although the models are optimized, they still operate separately, resulting in unnecessary overhead. Hence we propose and compare two methods that join both tasks, yielding a novel and efficient architecture.

Adjusting this architecture for each device and desired latency can be laborious and requires expert knowledge. These issues have recently garnered interest, leading to the emergence of new automated architecture search techniques, which have achieved competitive results in image classification [37, 41]. Moreover, recent Neural Architecture Search (NAS) approaches incorporate latency in their

loss function, enabling the design of architectures tailored to specific latency requirements [15, 18, 41]. Combining NAS with the proposed architecture provides a seamless and efficient approach to designing denoising and classification models for diverse use cases.

We find that our proposed efficiency-focused architecture consistently outperforms our more straightforward one. This is observed for both the manually and NAS designed models. In addition, our NAS models significantly outperform the manually designed ones in denoising and classification performance.

We have the following contributions. (i) We introduce a novel architecture to efficiently combine denoising and classification. The novelty lies in sharing an encoder between the denoiser and the classifier. (ii) We propose modifications to an existing NAS method for classification [18] to stabilize its search, improving the performance of the found architectures. (iii) We extend an existing NAS method to search for a model that combines denoising and classification, optimized for a target latency, classification accuracy, and denoising performance.

Since no prior work proposes a joint efficient model for denoising and classification, we study the tasks both separately and joint in Sec. 3. The findings are used as expert knowledge to construct the NAS method in Sec. 4.<sup>1</sup>

## 2 Related work

**Denoising.** Image denoising aims to reconstruct a clean image  $x$  from its observed noisy variant  $y$ . This relation can be formulated as  $y = x + n$ , where we assume  $n$  to be additive white Gaussian noise (AWGN). Neural network-based denoisers offer faster inference and good performance compared to traditional denoising methods like BM3D [7] and WNNM [13]. The interest in deep learning for denoising started with DnCNN [46], a simple Convolutional Neural Network (CNN). Encoder-decoder architectures became popular due to their efficient hierarchical feature extraction. Specifically, UNet [31] whose skip-connections between the encoder and decoder enhance the denoising process as shown in follow-up methods [14, 22, 28]. The interest in the UNet structure continues with transformer architectures [11, 38]. In this paper, our denoisers are based on UNet, ensuring our findings can translate to most related work.

**Efficient classification.** Optimization for efficiency is generally achieved by either compressing pre-trained networks [27] or designing small networks directly [32, 37]. We focus on efficient design, for which handcrafted models and neural architecture search (NAS) play essential roles. Studies proposing handcrafted models often introduce efficient operators [16, 32, 48] or scaling methods [36]. These efficient operators are used in NAS methods [37, 41] aiming for the automated design of efficient neural networks. Such an operator is the inverted residual with a linear bottleneck (MBConv), as introduced in MobileNetV2 [32]. In our models, we study scaling methods and MBConv’s efficiency characteristic.

**Neural Architecture Search.** The use of reinforcement learning (RL) for neural architecture search introduced efficient architectures with competitive

<sup>1</sup> Project site: <https://thomas-markhorst.github.io>

classification performance [15, 29, 35, 37]. However, their discrete search space is computationally expensive. Differentiable NAS (DNAS) methods [3, 21, 41] significantly reduce this cost by relaxing the search space to be continuous using learnable vectors  $\alpha$  for selecting candidate operations, which allows for gradient-based optimization. The popularity of DNAS started with DARTS [21], which searches a cell structure. Due to the complex design and repetitiveness throughout the network of the cell structure, follow-up works [18, 41] search operators for every layer instead of constructing repeating cells.

Pitfalls of DNAS are the collapse of search into some fixed operations and a performance drop when converting from the continuous search network to the discretized inference network [5, 43, 44]. TF-NAS [18] addresses these issues with an adaptation in the search algorithm, which lets the search model mimic the discrete behavior of the inference model. In addition, TF-NAS searches an architecture with a target latency by adding a latency loss to the search optimization. Because of these properties, we use TF-NAS as a baseline for our NAS study.


Existing NAS methods for denoising are either not reproducible [26], have a cell-based search space [45], or do not have an encoder-decoder [4] architecture. Instead, we use a layer-based search space and encoder-decoder structure.

**Joint classification and denoising.** In [42], the positive influence of denoising methods on classification performance is discussed. Moreover, [20] proposed a joint model where a VGG classifier [33] is attached to a denoiser similar to UNet. This method [20] reports a qualitative improvement of the denoised images when adding the classification loss to the denoiser’s optimization, whereas [42] reports a quantitative improvement. Although these models denoise and classify well, they are not optimized for efficiency. In this paper, we design a joint image denoising and classification method for edge devices.

### 3 Exploiting Expert Knowledge

We start in a controlled setting with separate baseline models for classification and denoising. Additionally, methods to increase their respective efficiency are studied, resulting in a reduced version of the baseline denoiser and classifier. Both the construction and efficiency improvement of the models are described in Suppl. A, where a UNet (Fig. 2) and simple 2-block CNN (Fig. 2.i and 2.ii) are used as baseline denoiser and classifier respectively. This section describes how the different sizes of the classifiers and denoisers are used to study joining methods and their efficiency.

**Dataset & settings.** For the experiments in this section, we generate a controlled synthetic data set to study the behavior of the classifier and denoiser when applying model scaling, replacing the convolutional operations, and combining both models. The dataset consists of 30k images, each with a random constant background in a gray tint [0.1 - 0.3] with two randomly placed non-overlapping MNIST [9] digits. We use two digits to increase the complexity of the denoising task. For experiments including classification, the two digits are extracted from the image using ground truth locations. These extracted digits



**Fig. 2:** A UNet, with hyperparameters base feature map width ( $b$ ), depth ( $d$ ), channel multiplier ( $m$ ) and convolutions per layer ( $c$ ). For the joint model either attach the classifier to form (i) the Sequential model or (ii) the Integrated model.

are separately used as input for the classifier. In the experiments where noise is required, for either denoising or noisy classification, synthetic Gaussian noise is added. This noise is zero mean, and the intensity of the noise is controlled using the standard deviation ( $\sigma$ ) of the distribution. Fig. 3a shows a sample, and Fig. 3b its noisy variant. To test the model behavior on an extensive noise range, every model is trained and tested on eleven  $\sigma$  values evenly spaced on the interval  $[0, 1]$  (Tabs. 1, 7 and 8). The models are trained using Adam optimizer with  $1E-3$  learning rate (LR), plateau LR scheduler, and 100 epochs.


Since the experiments with the controlled data set are not targeted at a specific device, the metric defining efficiency should not depend on a device. Such a metric is computational power, most commonly defined as Floating Point Operations (FLOPs), which we use as the primary metric. Despite being device dependent, we assess latency as a secondary metric. The latency is measured with a batch size of 32, 100 warm-up inference passes and averaged over 1000 inference passes. Classification performance is quantified using accuracy, while for denoising performance the Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM) metrics [39] are used. Higher is better for all our metrics.

### 3.1 Joint model: DC-Net

**Experimental setup.** We construct a baseline and reduced joint model, Denoising-Classifying Network (DC-Net). Both the baseline and reduced DC-Net use the same classifier (MB2.5-M Suppl. A). Whereas UNet-S and UNet are used for the reduced and baseline DC-Net respectively.

For joining the denoiser and classifier, we propose two models: (i) a Sequential model where the classifier is attached after the denoiser (Fig. 2.i), and (ii) an Integrated model, the classifier is attached to the UNet encoder (Fig. 2.ii). For the Integrated model, classification and denoising branches share the encoder.

The benefits of the Integrated model could come in threefold. First, using a shared encoder removes the need for a second large classifier, as in the Sequential method. Second, the decoder and classifier branches could run in parallel compared to running sequentially, which can result in lower latency. Thirdly, the



**Fig. 3:** Ground-truth sample (a), which is the target for the denoiser when given noisy image (b). S stands for the reduced model and L for the baseline. (c-f) are the cropped denoised outputs for input (b) and the red squares indicate the zoomed-in regions. For higher noise levels, the denoising performance of the Sequential model is worse than the Integrated model.

decoder is only optimized for denoising, since the optimization of the classifier does not influence it anymore. It can result in better image quality.

The models are trained using a weighted combination of the Cross-Entropy and Charbonnier loss [1, 19] (Eq. 1). We report the metrics averaged over all 11 noise levels,  $\sigma$  in  $[0, 1]$ .


$$\mathcal{L} = 0.1 \cdot \mathcal{L}_{CE} + 0.9 \cdot \mathcal{L}_{Char} \quad (1)$$

**Exp 1. Integrated vs. Sequential.** Which joining method performs better for the baseline, and does the same conclusion hold when reducing its size? We compare the *Sequential* and *Integrated* models. In Tab. 1, we see that for both the baseline and reduced DC-Net models, the Integrated version performs significantly better at denoising, while the Sequential version performs better at classification. The difference in denoising performance is visualized in Fig. 3. We see that both the reduced (3c) and baseline (3e) Integrated models reconstruct the digit clearly. Whereas both sizes of the Sequential model (3d and f) fail to reconstruct the digit.


**Conclusion.** The integrated model has a slightly lower classification accuracy compared to the Sequential model, yet it has superior performance in terms of image quality. When aiming for improved human perception, it is still required for the human to see the content of the image. Therefore, the Integrated model is more suitable for joint denoising and classification and is called DC-Net.

**Table 1:** Comparison of the reduced and baseline joint models. Both the Integrated and Sequential methods trained on the synthetic noise dataset. The integrated model performs significantly better in denoising and slightly worse in classification. The integrated model also scales down better.


DC-Net	Type	FLOPs (M) ↓	Lat. (ms) ↓	PSNR ↑	SSIM ↑	Acc. (%) ↑
Baseline	<b>Integrated</b>	1301.8	<b>7.14</b>	<b>32.8</b>	<b>0.97</b>	88.1
	Sequential	1302.1	7.55	27.1	0.95	<b>89.6</b>
Reduced	<b>Integrated</b>	51.2	<b>2.41</b>	<b>29.9</b>	<b>0.97</b>	86.2
	Sequential	51.5	2.83	25.2	0.92	<b>87.6</b>



**Fig. 4:** Stage-5:block-4’s  $\alpha$  values for Removed and Learned  $\beta$ . Search is more stable when  $\beta$  is removed.



**Fig. 5:** Acc for different search spaces, showed for different target latencies. Using fewer operations is the most robust.



**Fig. 6:** Comparing classifiers with similar latency on Imagenet100. Our model outperforms other methods.

## 4 Neural Architecture Search

We follow similar experimentation strategies as in the previous section. TF-NAS [18] is used to construct a classifier, which we use as a basis for our denoiser and joint model. All our proposed models in this Section contain searchable blocks, the models and which parts are searchable are defined in Figure 7.

**Dataset & settings.** The following experiments are conducted on Imagenet [8], randomly cropped to 224x224 pixels. To reduce search and training time, 100 classes (Imagenet 100) from the original 1000 classes were chosen, as in [18]. In the experiments requiring noise, Gaussian noise is sampled uniformly with a continuous range of  $\sigma$  in  $[0, 1]$  (Tabs. 2, 3, 4, 5 and 6).

The models are searched using SGD with momentum, 2E-2 LR with 90 epochs. Afterward, the found architecture is trained from scratch with 2E-1 LR for 250 epochs. All other settings are similar to [18]. The loss function depends on the task of the experiment, Cross-Entropy with label smoothing for classification ( $\mathcal{L}_{CE}$ ), combined Charbonnier and SSIM losses for denoising ( $\mathcal{L}_{Den}$ ), and a weighted combination for the joint model ( $\mathcal{L}_{Both}$ ), see Eq. 2, 3.

$$\mathcal{L}_{Den} = 0.8 \cdot \mathcal{L}_{Char} + 0.2 \cdot \mathcal{L}_{SSIM} \quad (2)$$

$$\mathcal{L}_{Both} = 0.1 \cdot \mathcal{L}_{CE} + 0.9 \cdot \mathcal{L}_{Den} \quad (3)$$

Since our NAS method uses a latency look-up table constructed for our device, these experiments target a specific device, GeForce RTX 3090 GPU. Therefore latency is suitable for defining efficiency in the NAS experiments.

### 4.1 Classification: C-NAS

**Experimental Setup.** Since TF-NAS [18] learns  $\beta$ ’s to control the number of convolutional operators per stage,  $\beta$ ’s can reduce the model size. However, in

the models proposed by [18], only 2 out of 24 stages are reduced by  $\beta$ . So the  $\beta$ 's have little effect on the found architectures, yet they make the search space more complex. Therefore, we propose a version of TF-NAS where the  $\beta$ 's are removed so that all convolutional blocks are used.

The candidate operations in the search space of TF-NAS are MBConvs with 8 different configurations, see Suppl. B. The configurations differ in kernel size, expansion rate, and in- or excluding a squeeze- and excitation layer (SE) [17].

The classification experiments are performed using data without noise, as the aim is to examine the NAS method, which is designed for clean images. We investigate key components of TF-NAS and try to improve its stability and classification performance.

**Exp 1. Learned vs. Removed  $\beta$ .** We conduct an experiment to study the effect of removing  $\beta$  on the search quality. The SE-layer is excluded from the candidate blocks, halving the search space to ensure the number of candidate operations does not cause search instability. We set a low target latency of 6 ms, as learning  $\beta$  should have a positive effect on small networks. For both the learned and removed settings, we run two searches, search 1 and 2.

Fig. 4 shows that when  $\beta$  is learned, the  $\alpha$ 's selecting a candidate operation oscillate and therefore do not decide on an architecture. Whereas with Removed  $\beta$ , the search is stable. This stability is reflected in the performance, as the average accuracy of the Removed  $\beta$  models is 86.3%, compared to 84.2% for Learned  $\beta$ . The separate results for each model are shown in Suppl. C.

**Exp 2. Number of operators in search space.** Does reducing the number of operators during search positively influence the performance of the found models? We test this by comparing the performance of architectures searched with three different search space sizes, {4, 6, or 8} operations, defined in Suppl. B. For each of these search spaces, three different latency targets are used: {6, 8, and 12} ms.

In Fig. 5, we see that for lower target latencies, 6 and 8 ms, using fewer operations in the search space does not alter performance significantly. When targeting 12 ms latency, reducing the number of operations in the search space does show a significant improvement. Additionally, we find that when using the larger search spaces, the operators from the small search space are still preferred for lower latencies.

**Exp 3. Compare with original TF-NAS.** How do architectures found using our proposed changes to TF-NAS perform compared to models with similar latency? We compare our model, C-NAS M, with TF-NAS C, MobileNetV2, and ResNet-18. MobileNetV2 our model have similar latency, architecture, and operator types. ResNet only differs in that it uses the Conv operator. We include these standard baseline architectures to indicate where C-NAS, see Fig. 7.i, stands on Imagenet100.

Fig. 6 shows that the model found using our method has lower latency yet higher accuracy than TF-NAS C as proposed in [18]. The model is searched with target latency 8.0. We observe that our search method is able to find a model that matches its target latency. Although ResNet-18 and MobileNetV2



run faster than our model, our classification accuracy is superior, especially when compared to ResNet-18, which only uses Convs.

**Conclusion.** By Removing  $\beta$  and reducing the number of operators used in the search, the search stability increases, and we find architectures that have better accuracy. An architecture found using our changes classifies better than a TF-NAS architecture with similar latency.

The comparison between our model and ResNet-18 shows that our search space is able to compete with widely accepted Conv-based classifiers. Moreover, our model performs on par with MobileNetV2, a manually designed classifier using MBConvs.

## 4.2 Denoising: D-NAS

**Experimental setup.** To construct a denoiser, D-NAS (Fig. 7.ii), we use the first six stages of a found C-NAS classifier, which has four levels of resolution. Afterwards, we attach a UNet style decoder by using both a transposed convolution and two normal convolutions for each decoder level. Like UNet, we also add skip connections between the encoder and decoder layers. The decoder is not searched.

**Exp 1. D-NAS vs UNet denoiser.** Does our denoiser D-NAS perform similarly to the UNet denoisers? For this experiment, we use UNet-S (Sec. A.2)  $\{d = 4, b = 8, c = 2, m = 1.5\}$ , with a latency of 9.2 ms and the larger UNet-M,  $\{d = 4, b = 16, c = 2, m = 2\}$  with a latency of 16.9 ms. We compare them with our D-NAS M, with similar latency.

**Table 2:** Comparison of D-NAS and UNet variants for denoising. D-NAS outperforms slightly faster UNet-S, but UNet-M denoises best at the cost of 45% higher latency.

Model	UNet params:		Lat. (ms) ↓	PSNR ↑	SSIM ↑
	d	b	m		
UNet-S	4	8	1.5	9.2	0.69
UNet-M	4	16	2	16.9	0.72
D-NAS M	-	-	-	11.6	0.71

Tab. 2 shows that D-NAS M outperforms UNet-S by 0.6 dB PSNR and 2% SSIM, at the cost of 2.4 ms latency. However, the 7.7 ms slower UNet variant, UNet-M, denoises better than our proposed model, by 0.3 dB and 1% SSIM.

**Conclusion.** D-NAS performs similarly to our baseline UNets. Therefore D-NAS is a suitable denoising architecture and it can form the backbone of our Integrated model.

## 4.3 Joint Model: DC-NAS

**Experimental setup.** To construct the joint model, we use the Integrated setup. The Integrated model, DC-NAS, is constructed similarly to D-NAS. We

**Table 3:** Comparison of DC-NAS models searched for three different latencies, with their corresponding C-NAS model, classifier baseline, and denoiser baseline. Our Integrated models perform similar or better than their corresponding baselines, with the advantage of having a joint denoising and classification network. \*See Suppl. A.2.


Model	Type	Lat. (ms) ↓	Classification	Denoising	
			Acc. (%) ↑	PSNR ↑	SSIM ↑
MobileNetV3 [15]	Classifier	4.9	70.4	-	-
C-NAS S ( <i>ours</i> )	Classifier	5.9	73.5	-	-
UNet-S* [31]	Denoiser	9.2	-	25.0	0.69
DC-Net S ( <i>ours</i> )	Integrated	10.0	61.9	24.5	0.68
DC-NAS S ( <i>ours</i> )	Integrated	10.3	<b>74.3</b>	<b>25.4</b>	<b>0.70</b>
EfficientNetV2-b0 [37]	Classifier	9.0	75.4	-	-
C-NAS M ( <i>ours</i> )	Classifier	7.9	75.5	-	-
LPIENet 0.25x [6]	Denoiser	12.7	-	24.1	0.65
DC-NAS M ( <i>ours</i> )	Integrated	13.7	<b>76.0</b>	<b>25.4</b>	<b>0.70</b>
EfficientNetV2-b1 [37]	Classifier	11.8	<b>76.7</b>	-	-
C-NAS L ( <i>ours</i> )	Classifier	12.0	76.0	-	-
UNet-M* [31]	Denoiser	16.9	-	<b>25.9</b>	<b>0.72</b>
LPIENet 0.5x [6]	Denoiser	19.8	-	24.7	0.68
DC-NAS L ( <i>ours</i> )	Integrated	17.9	76.4	25.2	0.70

connect the decoder after the first six stages of C-NAS (Fig. 7, but still use the remaining C-NAS stages(Fig. 7.iii)) as a classification branch. The design choices for DC-NAS are discussed in the ablations study (Sec. 4.4).

Using our search method, we search for DC-NAS models of three different sizes {S, M, L}. Apart from our manually designed Integrated model, we compare our searched models with separate state-of-the-art classifiers and denoisers, as there are no existing models that jointly optimize denoising, classification, and efficiency. For each DC-NAS model, we also separately train the classifier (C-NAS) to evaluate the influence of joint denoising. The classifier and denoiser baselines are chosen to have similar latency as the C-NAS or D-NAS model on which the corresponding DC-NAS is based.

**Results.** We discuss the results in Tab. 3 in three separate sections for the different target latencies. Our smallest Integrated model, DC-NAS S, outperforms both of its classifier baselines MobileNetV3 and C-NAS S. Note, that the latter shows that the classifier, C-NAS S, performs better when integrated into DC-NAS S. Moreover, our Integrated model denoises better than its baseline UNet-S (Suppl. A). DC-NAS S also significantly outperforms our manually designed DC-Net S (Reduced), which is the only Integrated baseline. We display the denoising results of DC-Net, UNet-S, and DC-NAS S in Fig. 8. We observe better denoising on smooth areas, sharper edges, and more realistic color reconstruction for DC-NAS S.

The results of DC-NAS M follow a similar pattern, where our DC-NAS outperforms its baseline denoiser and classifier, using C-NAS M in the Integrated model boosts accuracy by 0.5%. When comparing DC-NAS M to DC-NAS S,



**Fig. 7:** C-NAS and D-NAS architecture. Connecting block (iii) for DC-NAS and block (iv) for DC-NAS<sub>seq</sub>. During search models with various latencies can be obtained. Only the orange stages are searchable in the encoder and classifier.

the classification performance improves by 1.7%, yet the denoising performance plateaus. LPIENet denoises the worst. Comparing DC-NAS M’s denoising performance to 3.2 ms slower UNet-M we observe slightly worse denoising performance. However, our Integrated model denoises and classifies with lower latency.

For DC-NAS L, we observe that both the classification and denoising baseline slightly outperform our Integrated model. EfficientNetV2-b1 has 0.3% higher classification accuracy than DC-NAS L, and UNet-M improves with 0.7 dB PSNR and 2% SSIM. However, our Integrated model performs both denoising and classification at a similar latency as UNet-M, which only denoises. When comparing DC-NAS L with DC-NAS M, we again note an improvement in classification performance. However, the PSNR score drops by 0.2 dB while the more important SSIM score remains at 0.70.

**Conclusion.** Our results demonstrate that the proposed DC-NAS models perform similar or better than their denoising and classification baselines for their target latency. In addition, the searched model performs better than our manually designed joint denoiser and classifier.

#### 4.4 Ablation Study

**Exp 1. Encoder search.** C-NAS forms the encoder of DC-NAS and contains the searchable operations within DC-NAS. We test multiple search approaches:

(i) using clean images, and (ii) using noisy images. For both approaches, we search the encoder using only classification loss  $\mathcal{L}_{\text{Cls}}$ . In addition, we also search the DC-NAS encoder on noisy images using the combined denoising and classification loss  $\mathcal{L}_{\text{Both}}$ . Therefore it searches for the optimal encoder for both tasks within the Integrated model DC-NAS. Regardless of the search method, the found models are trained using noisy images and the combined loss.

Tab. 4, shows that when using  $\mathcal{L}_{\text{Cls}}$  with noisy images during search improves classification accuracy by 0.3%. Surprisingly, the denoising performance is the same. Using both the denoising and classification objectives during the search reduces the classification accuracy. Caused by the denoising loss complicating the search, without improving denoising performance. Therefore, we search our DC-NAS models by only using  $\mathcal{L}_{\text{Cls}}$  loss.

**Exp 2. Compare Integrated vs. Sequential.** We compare DC-NAS and DC-NAS<sub>seq</sub> models with similar latency. Where DC-NAS<sub>seq</sub> is our Sequential model, which is constructed by attaching C-NAS to the output of D-NAS, see Fig. 7.iv. Since the searched classifier is used twice in DC-NAS<sub>seq</sub>, the Sequential model has a higher latency than the Integrated variant. To counter this, a smaller C-NAS model is used in both the encoder and classifier of DC-NAS<sub>seq</sub>. The classifier, C-NAS, used to construct DC-NAS<sub>seq</sub> L has a latency of 6.7 ms. Whereas in DC-NAS L, the classifier has a latency of 12 ms. Note, that these models were searched by using clean instead of noisy images, as this holds for both models it is still a fair comparison.

We see that both models have similar latency and the same classification accuracy, however, DC-NAS L improves denoising performance with 0.5 dB PSNR and 1% SSIM (Tab. 5). This improvement is caused by DC-NAS L’s Integrated design as this allows for a bigger encoder without increasing latency.


**Table 4:** Different search strategies for DC-NAS, using (i) clean or noisy images and (ii)  $\mathcal{L}_{\text{Cls}}$  or  $\mathcal{L}_{\text{Cls}} + \mathcal{L}_{\text{Den}}$ . Searching on **Noisy** images with only  $\mathcal{L}_{\text{Cls}}$  performs best.

Search	Lat. (ms)	Acc. (%)	↑ SSIM	↑ PSNR
Clean $\mathcal{L}_{\text{Cls}}$	13.9	75.7	25.4	0.70
Noisy $\mathcal{L}_{\text{Cls}}$	<b>13.7</b>	<b>76.0</b>	25.4	0.70
Noisy $\mathcal{L}_{\text{Both}}$	13.8	75.5	25.4	0.70

**Exp 3. Decoder tuning.** The DC-NAS models found in Tab. 3 and 5, have similar denoising performance. These models differ only in the type of MBConvs that are selected during search in the encoder. We test the hypothesis if the denoising performance is influenced by adjusting the operators in the decoder while retaining the latency. DC-NAS M is used as a basis in this experiment. We construct three alternatives. First, the convolutional operators in the decoder are replaced with MBConvs (MB-k3-e3) [32], which significantly increases the latency of the model. To account for this, we scale down the decoder by (i)

**Table 5:** Comparing Sequential and Integrated DC-NAS, classification performance is similar, yet the Integrated model is faster and denoises better.

Model	Lat. (ms)	Acc. (%)	PSNR $\uparrow$	SSIM $\uparrow$
DC-NAS <sub>seq</sub> L	18.3	76.0	25.0	0.69
DC-NAS L	<b>17.9</b>	76.0	<b>25.5</b>	<b>0.70</b>

**Fig. 8:** Denoising performance of DC-NAS S and its baselines. Left to right: noisy image, the denoiser outputs of size S, and the clean image. Comparing (d) and (b,c), we see better performance in smooth areas and more correct colors in (d). With (i) and (g), we observe a better color reconstruction for (i). Moreover, (i) has less artifacts than (h). Hence, DC-NAS S denoises better than the other denoisers of similar latency.

using 1 instead of 2 convolutional operations (MBCConv) per layer or (ii) using 3 instead of 4 decoder layers.

In Tab. 6, we see that using the MBCConvs compared to Convs improves the denoising performance. However, at the cost of a 14 ms latency increase, only caused by the MBCConv decoder. When reducing the complexity of the MBCConv decoder with *1 operator* and *3 layers*, the denoising performance reduces to the original level again, but the latency is still higher than for DC-NAS M which has only standard convolutional layers in the decoder block.

**Conclusion.** We have seen that the Integrated combining method outperforms its Sequential counterpart in denoising. To construct the integrated model (DC-NAS), we find that searching for a classifier on noisy data, without taking the denoising objective into account results in the best classification performance.

**Table 6:** The influence of altering the Conv operators in the DC-NAS M decoder to MBConv and scaling down the MBConv alternative by reducing the number of operators or decoder layers. Using the standard convolutional layers is more efficient.

Decoder		Lat. (ms)	Acc. (%)	PSNR $\uparrow$	SSIM $\uparrow$
Operator	Scaling				
Conv	-	<b>13.7</b>	76	25.4	0.70
MBConv	-	27.7	75.5	25.8	0.71
MBConv	1 operator	16.4	75.4	25.3	0.70
MBConv	3 layers	22.1	75.1	25.4	0.70

Surprisingly, the search method does not influence the denoising performance. Furthermore, manually altering the decoder does not benefit denoising efficiency either. However, the NAS denoising experiments demonstrate that our denoising setup is competitive. Since tuning the decoder operators does not improve performance, our method is focused on searching for only the encoder of the integrated model. The models found by this approach, outperform our manually designed models with similar latency.

## 5 Limitations & Conclusion

One limitation of our NAS method is its inability to alter the decoder. It is designed this way as manually altering the decoder does not improve efficiency. However, when targeting a significantly different latency, a change in denoising architecture could be required. Therefore, designing model scaling rules for the searched models is of interest, similar to the EfficientNets [36, 37].

Another limitation is the fixation of  $\beta$  in our NAS method. Although this improves the stability of search and network performance, learning  $\beta$  while retaining a stable search would be preferred. This would introduce more possibilities in the search space for optimizing efficiency.

In addition, the latency of Integrated models can be reduced further by running the denoising and classification branches in parallel.

To conclude, we show that using efficient operators and scaling methods proposed in previous work [24, 32, 36] are relevant for denoising and noisy classification. In addition, we present the Integrated model DC-Net to join the two tasks efficiently and show that the Integrated design is more suitable across various latencies than the Sequential variant. To simplify the design process of the joint model when targeting a latency, we present a NAS method. We alter an existing NAS method to improve the stability and performance of the search. This method searches a classifier. Using the searched classifier as a basis, we build the Integrated DC-NAS model. We demonstrate that the proposed model outperforms the manually constructed model. We believe that our study can be a precursor of efficient joint low-level and high-level computer vision tasks.

## References

1. Barron, J.T.: A more general robust loss function. CoRR **abs/1701.03077** (2017), <http://arxiv.org/abs/1701.03077>
2. Bosch, N., D’Mello, S.K.: Can computers outperform humans in detecting user zone-outs? implications for intelligent interfaces. *ACM Trans. Comput.-Hum. Interact.* **29**(2) (jan 2022). <https://doi.org/10.1145/3481889>, <https://doi.org/10.1145/3481889>
3. Cai, H., Zhu, L., Han, S.: Proxylessnas: Direct neural architecture search on target task and hardware. CoRR **abs/1812.00332** (2018), <http://arxiv.org/abs/1812.00332>
4. Cheng, A., Wang, J., Zhang, X.S., Chen, Q., Wang, P., Cheng, J.: DPNAS: neural architecture search for deep learning with differential privacy. CoRR **abs/2110.08557** (2021), <https://arxiv.org/abs/2110.08557>
5. Chu, X., Zhou, T., Zhang, B., Li, J.: Fair DARTS: eliminating unfair advantages in differentiable architecture search. CoRR **abs/1911.12126** (2019), <http://arxiv.org/abs/1911.12126>
6. Conde, M.V., Vasluianu, F., Vazquez-Corral, J., Timofte, R.: Perceptual image enhancement for smartphone real-time applications. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. pp. 1848–1858 (January 2023)
7. Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K.: Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on Image Processing* **16**(8), 2080–2095 (2007). <https://doi.org/10.1109/TIP.2007.901238>
8. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *2009 IEEE conference on computer vision and pattern recognition*. pp. 248–255. Ieee (2009)
9. Deng, L.: The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine* **29**(6), 141–142 (2012)
10. Dey, B., Halder, S., Khalil, K., Lorusso, G., Severi, J., Leray, P., Bayoumi, M.A.: SEM image denoising with unsupervised machine learning for better defect inspection and metrology. In: Adan, O., Robinson, J.C. (eds.) *Metrology, Inspection, and Process Control for Semiconductor Manufacturing XXXV*. vol. 11611, p. 1161115. International Society for Optics and Photonics, SPIE (2021). <https://doi.org/10.1117/12.2584803>, <https://doi.org/10.1117/12.2584803>
11. Fan, C.M., Liu, T.J., Liu, K.H.: Sunet: Swin transformer unet for image denoising. In: *2022 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE (May 2022). <https://doi.org/10.1109/iscas48785.2022.9937486>, <http://dx.doi.org/10.1109/ISCAS48785.2022.9937486>
12. Gu, S., Li, Y., Gool, L.V., Timofte, R.: Self-guided network for fast image denoising. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (October 2019)
13. Gu, S., Zhang, L., Zuo, W., Feng, X.: Weighted nuclear norm minimization with application to image denoising. In: *2014 IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2862–2869 (2014). <https://doi.org/10.1109/CVPR.2014.366>
14. Gurrola-Ramos, J., Dalmau, O., Alarcón, T.E.: A residual dense u-net neural network for image denoising. *IEEE Access* **9**, 31742–31754 (2021). <https://doi.org/10.1109/ACCESS.2021.3061062>

15. Howard, A., Sandler, M., Chu, G., Chen, L., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., Le, Q.V., Adam, H.: Searching for mobilenetv3. CoRR **abs/1905.02244** (2019), <http://arxiv.org/abs/1905.02244>
16. 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. CoRR **abs/1704.04861** (2017), <http://arxiv.org/abs/1704.04861>
17. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. CoRR **abs/1709.01507** (2017), <http://arxiv.org/abs/1709.01507>
18. Hu, Y., Wu, X., He, R.: TF-NAS: rethinking three search freedoms of latency-constrained differentiable neural architecture search. CoRR **abs/2008.05314** (2020), <https://arxiv.org/abs/2008.05314>
19. Ignatov, A., Timofte, R., Liu, S., Feng, C., Bai, F., Wang, X., Lei, L., Yi, Z., Xiang, Y., Liu, Z., Li, S., Shi, K., Kong, D., Xu, K., Kwon, M., Wu, Y., Zheng, J., Fan, Z., Wu, X., Zhang, F., No, A., Cho, M., Chen, Z., Zhang, X., Li, R., Wang, J., Wang, Z., Conde, M.V., Choi, U.J., Perevozchikov, G., Ershov, E., Hui, Z., Dong, M., Lou, X., Zhou, W., Pang, C., Qin, H., Cai, M.: Learned smartphone isp on mobile gpus with deep learning, mobile ai & aim 2022 challenge: Report (2022), <https://arxiv.org/abs/2211.03885>
20. Liu, D., Wen, B., Jiao, J., Liu, X., Wang, Z., Huang, T.S.: Connecting image denoising and high-level vision tasks via deep learning. CoRR **abs/1809.01826** (2018), <http://arxiv.org/abs/1809.01826>
21. Liu, H., Simonyan, K., Yang, Y.: DARTS: differentiable architecture search. CoRR **abs/1806.09055** (2018), <http://arxiv.org/abs/1806.09055>
22. Liu, P., Zhang, H., Zhang, K., Lin, L., Zuo, W.: Multi-level wavelet-cnn for image restoration. CoRR **abs/1805.07071** (2018), <http://arxiv.org/abs/1805.07071>
23. Lopera Tellez, O.: Underwater threat recognition: Are automatic target classification algorithms going to replace expert human operators in the near future? In: OCEANS 2019 - Marseille. pp. 1–4 (2019). <https://doi.org/10.1109/OCEANSE.2019.8867168>
24. Matuszewski, D.J., Sintorn, I.M.: Reducing the u-net size for practical scenarios: Virus recognition in electron microscopy images. *Computer Methods and Programs in Biomedicine* **178**, 31–39 (2019). <https://doi.org/https://doi.org/10.1016/j.cmpb.2019.05.026>, <https://www.sciencedirect.com/science/article/pii/S0169260719300859>
25. Mohd Sagheer, S.V., George, S.N.: A review on medical image denoising algorithms. *Biomedical Signal Processing and Control* **61**, 102036 (2020). <https://doi.org/https://doi.org/10.1016/j.bspc.2020.102036>, <https://www.sciencedirect.com/science/article/pii/S1746809420301920>
26. Możejko, M., Latkowski, T., Łukasz Treszczotko, Szafraniuk, M., Trojanowski, K.: Superkernel neural architecture search for image denoising (2020)
27. O’Neill, J.: An overview of neural network compression. CoRR **abs/2006.03669** (2020), <https://arxiv.org/abs/2006.03669>
28. Park, B., Yu, S., Jeong, J.: Densely connected hierarchical network for image denoising. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 2104–2113 (2019). <https://doi.org/10.1109/CVPRW.2019.00263>
29. Pham, H., Guan, M.Y., Zoph, B., Le, Q.V., Dean, J.: Efficient neural architecture search via parameter sharing. CoRR **abs/1802.03268** (2018), <http://arxiv.org/abs/1802.03268>



30. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents (2022)
31. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. CoRR **abs/1505.04597** (2015), <http://arxiv.org/abs/1505.04597>
32. Sandler, M., Howard, A.G., Zhu, M., Zhmoginov, A., Chen, L.: Inverted residuals and linear bottlenecks: Mobile networks for classification, detection and segmentation. CoRR **abs/1801.04381** (2018), <http://arxiv.org/abs/1801.04381>
33. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition (2015)
34. Singh, P., Shankar, A.: A novel optical image denoising technique using convolutional neural network and anisotropic diffusion for real-time surveillance applications. *Journal of Real-Time Image Processing* **18**(5), 1711–1728 (Oct 2021). <https://doi.org/10.1007/s11554-020-01060-0>, <https://doi.org/10.1007/s11554-020-01060-0>
35. Tan, M., Chen, B., Pang, R., Vasudevan, V., Le, Q.V.: Mnasnet: Platform-aware neural architecture search for mobile. CoRR **abs/1807.11626** (2018), <http://arxiv.org/abs/1807.11626>
36. Tan, M., Le, Q.V.: Efficientnet: Rethinking model scaling for convolutional neural networks. CoRR **abs/1905.11946** (2019), <http://arxiv.org/abs/1905.11946>
37. Tan, M., Le, Q.V.: Efficientnetv2: Smaller models and faster training. CoRR **abs/2104.00298** (2021), <https://arxiv.org/abs/2104.00298>
38. Wang, Z., Cun, X., Bao, J., Liu, J.: Uformer: A general u-shaped transformer for image restoration. CoRR **abs/2106.03106** (2021), <https://arxiv.org/abs/2106.03106>
39. Wang, Z., Bovik, A., Sheikh, H., Simoncelli, E.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004). <https://doi.org/10.1109/TIP.2003.819861>
40. Williams, D.P., Couillard, M., Dugelay, S.: On human perception and automatic target recognition: Strategies for human-computer cooperation. In: 2014 22nd International Conference on Pattern Recognition. pp. 4690–4695 (2014). <https://doi.org/10.1109/ICPR.2014.802>
41. Wu, B., Dai, X., Zhang, P., Wang, Y., Sun, F., Wu, Y., Tian, Y., Vajda, P., Jia, Y., Keutzer, K.: Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. CoRR **abs/1812.03443** (2018), <http://arxiv.org/abs/1812.03443>
42. Wu, J., Timofte, R., Huang, Z., Van Gool, L.: On the relation between color image denoising and classification. arXiv preprint arXiv:1704.01372 (2017)
43. Xie, S., Zheng, H., Liu, C., Lin, L.: SNAS: stochastic neural architecture search. CoRR **abs/1812.09926** (2018), <http://arxiv.org/abs/1812.09926>
44. Ye, P., Li, B., Li, Y., Chen, T., Fan, J., Ouyang, W.:  $\beta$ -darts: Beta-decay regularization for differentiable architecture search (2022)
45. Zhang, H., Li, Y., Chen, H., Shen, C.: IR-NAS: neural architecture search for image restoration. CoRR **abs/1909.08228** (2019), <http://arxiv.org/abs/1909.08228>
46. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a gaussian denoiser: Residual learning of deep CNN for image denoising. CoRR **abs/1608.03981** (2016), <http://arxiv.org/abs/1608.03981>
47. Zhang, K., Zuo, W., Zhang, L.: Ffdnet: Toward a fast and flexible solution for CNN based image denoising. CoRR **abs/1710.04026** (2017), <http://arxiv.org/abs/1710.04026>

48. Zhang, X., Zhou, X., Lin, M., Sun, J.: Shufflenet: An extremely efficient convolutional neural network for mobile devices. CoRR **abs/1707.01083** (2017), <http://arxiv.org/abs/1707.01083>