This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

ISNAS-DIP: Image-Specific Neural Architecture Search for Deep Image Prior

Metin Ersin Arican^{1*}, Ozgur Kara^{1*}, Gustav Bredell², Ender Konukoglu²

¹Department of Electrical and Electronics Engineering, Bogazici University, Istanbul, Turkey

² Department of Information Technology and Electrical Engineering, ETH-Zurich, Zurich, Switzerland

¹{metin.arican, ozgur.kara}@boun.edu.tr

²{gustav.bredell, ender.konukoglu}@vision.ee.ethz.ch

Abstract

Recent works show that convolutional neural network (CNN) architectures have a spectral bias towards lower frequencies, which has been leveraged for various image restoration tasks in the Deep Image Prior (DIP) framework. The benefit of the inductive bias the network imposes in the DIP framework depends on the architecture. Therefore, researchers have studied how to automate the search to determine the best-performing model. However, common neural architecture search (NAS) techniques are resource and time-intensive. Moreover, best-performing models are determined for a whole dataset of images instead of for each image independently, which would be prohibitively expensive. In this work, we first show that optimal neural architectures in the DIP framework are image-dependent. Leveraging this insight, we then propose an image-specific NAS strategy for the DIP framework that requires substantially less training than typical NAS approaches, effectively enabling image-specific NAS. We justify the proposed strategy's effectiveness by (1) demonstrating its performance on a NAS Dataset for DIP that includes 522 models from a particular search space (2) conducting extensive experiments on image denoising, inpainting, and super-resolution tasks. Our experiments show that image-specific metrics can reduce the search space to a small cohort of models, of which the best model outperforms current NAS approaches for image restoration. Codes and datasets are available at https://github.com/ozgurkara99/ISNAS-DIP.

1. Introduction

Convolutional neural networks (CNNs) have been ubiquitously utilized in almost every field of computer vision. Particularly, researchers harness the power of CNNs in image restoration tasks [1, 2, 3, 4], which refers to the task of recovering the original image from a corrupted version. The success of CNNs comes as a result of their ability to learn a mapping from a corrupted image to its uncorrupted counterpart. However, the ground truth labels are not always available to learn such a mapping for a given domain, limiting the use of approaches under supervised settings. To tackle this problem, researchers orient their attention towards unsupervised approaches. Recent discoveries have shown that the architecture of CNNs contains an intrinsic *prior* that can be used in image restoration tasks [5, 6]. This insight led to the Deep Image Prior (DIP) framework [5], which works solely with the degraded image and can produce competitive results for image restoration tasks without a supervised training phase. It offers an alternative solution to restoration problems by suggesting a new regularizer: the network architecture itself. In addition to this empirical discovery, Rahaman et al. [7] investigated the spectral bias of neural networks towards low frequencies theoretically, which can explain the impressive performance of the DIP framework. Chakrabarty [8] further explored the underlying reason behind the success of DIP in denoising natural images. The work demonstrates that the network tends to behave similarly to a low pass filter at the early stages of iterations. Finally, DeepRED [9] merged the concept of "Regularization by Denoising" (RED) by adding explicit priors to enhance DIP.

One problem faced in the DIP framework is that the architectural design of the network has a substantial impact on the performance. Recent works attempted to automate the search process of network architecture for various tasks, which is referred to as the *Neural Architecture Search* (NAS). In the context of DIP, Chen *et al.* [10] applied NAS to the DIP framework. However, current NAS approaches come with substantial computational costs, as they require optimizing a large number of architectures to determine the optimum. Moreover, this cost prohibits determining the optimum architecture for every image; instead, existing NAS approaches search for the best architecture for a dataset of images.

Our work. In this paper, we propose novel image-

^{*}equal contribution

dependent metrics to determine optimal network architectures in the DIP framework with minimal training. Unlike previous works, we apply our metrics to DIP for finding image-specific architectures since performance is strongly dependent on the content of the image that is to be restored.

We first motivate image-specific NAS, by showing that in a given search space, there is only a small overlap of the best architectures for different images. This is illustrated in Figure 1, where the matrices show the number of overlaps between the top 10 models (of a total of 522 models) for each image for denoising and inpainting.

To identify architectures that are fitting for a specific image, we propose image-dependent metrics that measure the property of how far the power spectral density (PSD) of the generated initial output of a network is from that of the corrupted image and use it as our metric. The intuition relies on the fact that the more these two are similar, the better the model will reconstruct the image since it is closer to the solution space.



Figure 1: The overlap of the best-performing architectures between different images is shown here. The numbers indicate how many of the top-10 models of the search space for image x are also in the top-10 for image y. This is shown for the task of denoising and inpainting in (a) and (b), respectively. E.g. the value at the intersection of chest and lena in the denoising heatmap, which is 2, indicates that there are 2 models in each of the images' best-performing 10 models that are the same.

We motivate the choice of metrics by looking at the correlation between the metrics' values and image restoration performance. There is an imperfect correlation; hence we select a small cohort of architectures to optimize based on the metrics' values. A final selection is then made by selecting the model whose output is closest to the average of outputs of all models.

We conduct experiments on conventional datasets for image denoising, image inpainting, and single image superresolution tasks using the proposed strategy. For each image in the datasets, we run our ISNAS algorithm to identify the optimal image-specific models. The results demonstrate that our method is superior to the state-of-the-art work [10] in terms of its quality improvement.

The main contributions can be summarized as follows:

- We empirically show the necessity of identifying image-specific models to augment the quality of DIP.
- We present novel metrics to be used in NAS requiring only the randomly initialized CNN network. These metrics allow ranking architectures within any search space without lengthy optimization as a surrogate to their success on image restoration tasks.
- We introduce two selection procedures among a subset of models for finding optimal architectures in an unsupervised fashion for DIP.
- We generate a NAS Dataset for DIP having 522 models optimized for ten images from different domains, including image denoising and image inpainting tasks,.
- Extensive experiments on commonly used datasets and NAS Dataset for DIP validate our approach.

2. Related Work

2.1. NAS

The challenges of designing complex architectures by hand have shifted researchers' interest to the area of automatic neural architecture search [11]. One of the initial works [12] formulated NAS as a reinforcement learning problem, where better architectures are sampled by training a policy network. Building upon this work, Zoph et al. [13] applied a cell-based search which is then stacked to constitute the network. Pham et al. [14] introduced the idea of *weight sharing* where the networks are trained jointly. They show that it reduces the complexity by 1000 times. In the context of image restoration, Suganuma et al. [15] employed evolutionary search to convolutional autoencoders. Subsequently, Zhang et al. developed HiNAS [16], which utilized gradient-based search strategies and introduced a hierarchical neural architecture search for image denoising. A recent study [10] utilized a reinforcement learningbased NAS to improve the network architecture and offered a search space for upsampling blocks, of which we follow the same procedures throughout our experiments.

2.2. NAS Without Training

One of the major challenges in the NAS is its computational expense since numerous networks need to be evaluated before determining the best model. To minimize resource consumption, researchers have proposed *trainingfree* approaches. Recently, Mellor *et al.* [17] proposed to assign scores to network architectures at initialization requiring no training. Their intuition is that the closeness of binary codes associated with two inputs is an indicator of how much the network is able to separate these. They empirically show the correlation between the initial score and



Figure 2: Visualization of proposed metrics. The left block with three branches demonstrates the image-dependent metrics. All of them use the PSDs of corrupted images and random output. The right block shows the image-independent metric, which only requires a randomly initialized CNN architecture

the network's performance, although it lacks theoretical explanations. Chen *et al.* [18] attempted to bridge the deep learning theory and NAS applications by questioning and analyzing the effectiveness of two indicators on ranking the networks prior to training, namely *trainability*, and *expressivity*. Zhang *et al.* [19] went a step further and proposed a metric based on *synaptic saliency* that does scoring requiring neither training nor labels. They achieved promising results on NAS-Bench-201 [20], which is a conventional benchmark used for evaluating NAS algorithms. However, existing NAS benchmarks are mainly composed of image classification datasets.

Inspired by the previous *training-free* approaches, we adapt the idea of ranking networks at initialization to DIP settings. Furthermore, to evaluate our NAS approach, we generate our *NAS Dataset* for DIP analogous to NAS-Bench 201 [20], which is explained in Section 4.1.

3. Methodology

In a typical NAS algorithm, the main bottleneck, in terms of time, is the training of the models to compute their performance. If one can find an easy-to-calculate training-free performance predictor, this bottleneck can be eliminated. In addition, our experiments show that model selection for DIP settings should be image-dependent. In this paper, we propose several different training-free and image-dependent performance predictors and study their effectiveness.

Ulyanov *et al.* [5] observed that the architectures with better performances in DIP tend to have outputs possessing large spatial structures at the early iterations of training. One useful metric to capture the distribution of the spatial

structures is to use power spectral density (PSD). Coarse and fine textures will lead to a PSD that is concentrated on low and high frequencies, respectively.

Inspired by these, we put forward a hypothesis that if the PSD of an untrained CNN's output is *similar* to that of the image to be reconstructed, then the model will be closer to the desired solution space, hence it will facilitate the optimization and lead to better restoration results compared to others with lower similarity. In this section, we formulate different metrics to quantify the similarity between an image and the CNN's random output.

3.1. Image Dependent Metrics

It would be preferable to compute the distance using the ground truth image. However, in a practical situation, we do not have access to them. Therefore, we use the distance between the CNN's random output and the corrupted image as a proxy for the distance between the output and the ground truth image.

3.1.1 PSD DB MSE

One straightforward approach that can be used to measure the distance between PSDs is the mean square error (MSE). Generally, the PSD of an array consists of numbers that are orders of magnitudes different from each other. Thus, in MSE calculation, using the logarithm of the PSD, which we call decibel PSD, instead of directly PSD itself is better suited. To that end, we first calculate the decibel PSD's of the output of a given CNN with randomly initialized weights and the corrupted image. Then, we measure the MSE between them. A schematic representation of this metric can be seen in Fig. 2 and it is formulated as

$$\frac{1}{n} \sum_{i,j} \left(10 \cdot \log X_{i,j} - 10 \cdot \log Y_{i,j} \right)^2 \tag{1}$$

where X and Y denote the corresponding power spectral densities of the CNN's random output and a corrupted image, respectively, and n denotes the number of pixels.

3.1.2 PSD DB Strip MSE

Spatial structures and texture in an image are related to its PSD, but each frequency region of the PSD does not equally contribute to the spatial structures. Very high-frequency regions of the PSD of a corrupted image are heavily affected by noise. Hence, focusing only a band of frequencies around the center in similarity comparison is more suitable. In light of these insights, the metric is calculated as follows: First, a mask is applied to the decibel of the PSDs of the CNN's random output and the corrupted image. Then, we calculate the MSE between them. As the mask, we employ a strip having inner and outer diameter sizes of 10% and 20% of the image size, respectively, to reduce the dependency on the image size. A schematic representation of this metric can be seen in Fig. 2 and it is formulated as

$$\frac{1}{n} \sum_{i,j} \left(10 \cdot \log X_{i,j} - 10 \cdot \log Y_{i,j} \right)^2 \cdot M_{i,j}$$
 (2)

where X and Y denote the corresponding power spectral densities of the CNN's random output and a corrupted image, respectively, M denotes the mask, a 2-D array of 1s and 0s, and n denotes the number of non-zero pixels of the mask M.

3.1.3 PSD Strip Hist EMD

To make our metrics rotation invariant, we use the histograms of PSDs. In this metric, we first calculate the PSDs of the CNN's random output and a corrupted image. Then, we discard the entries of PSDs, where the corresponding entry of a mask is zero. For this, we use the mask defined in the previous metric (PSD DB Strip MSE). Afterwards, the PSDs are flattened into two 1-D arrays, which are then converted to histogram representations. Finally, we calculate the earth mover's distance (EMD) between these two histograms. The range and number of bins of the histograms are determined as a result of trials. In our experiments, we use 75 as the number of bins and 0-1 as the range of the histograms. A schematic representation of this metric can be seen in Fig. 2.

3.2. Image Independent Metrics

In our evaluation, we also included an image independent metric also using the PSD, inspired by the structural bias of CNNs as exploited by Heckel [21]. This allows us to dissect whether the contribution is due to using the PSD of the CNN-generated image or the image dependency of the metrics described above. The structural part of an image can be thought of as the low-frequency component, just as noise or corruptions are of high-frequency components. Relying on the hypothesis that if the frequency spectrum of the output of a randomly initialized CNN is concentrated on low-frequency regions, then it tends to perform better in restoration tasks such as denoising and inpainting, we propose a metric to measure the low-pass characteristic of a CNN and use it as our image independent metric.

3.2.1 99% Bandwidth (99 BW)

A straightforward method to quantify the low-frequency nature of an array is to calculate its bandwidth. We define the P% bandwidth of a 2D array as the radius of the circle containing the P% of the total energy in the PSD of the 2D array.

Obviously, the bandwidth of a CNN does not depend on the image to be reconstructed, which makes it an image independent metric. To choose a value of P, we created several outputs from randomly initialized CNNs and selected the P value resulting in the most variation in bandwidths. This was P = 99, so we used this value in our experiments.

3.3. ISNAS-DIP Overview

Given any search space, we can just use one of the metrics to calculate the fitness score of models to shrink the entire search space into a few models. As mentioned, the metrics are correlated to the performance of the restoration of an optimized network, but the correlation is imperfect. Therefore, we first calculate the metrics of the models at initialization, i.e., without any optimization of the weights. Then, we sort each model in ascending order for "Image Independent" and "Image Dependent" metrics according to their values. Finally, for each metric, we choose top-N models having the lowest metric values.

Next, we have to make a final selection between the chosen N models. We propose two *averaging techniques* to perform the selection. The first selection procedure is as follows: Optimize the chosen N models under DIP settings. Take the average of the N reconstructed images that have coefficients inversely proportional to the value of the corresponding models' metrics. Calculate the MSE score between the average and each reconstructed image. Select the model that corresponds to the reconstructed image giving the lowest MSE with average restored, that is to say, closest one to the average. We will call this technique as *full-sized averaging selection* throughout the paper.

The second technique differs in the size of the image to be optimized by each model and dramatically increase the speed of ISNAS. Before running optimization for N models, we resize the corrupted image to " 64×64 " by rescaling instead of the full-sized image to speed up the optimization process. We will call this technique *resized averaging selection* throughout the paper. By choosing the model whose output is closest to the average, we at least guarantee to exclude the worst models among N models. Note that these selection methods introduce a negligible image dependency for 99 BW metric since the final model selection among N models may depend on the content of the image.

4. Experimentation

In this section, we first describe the experiment set-up, including the search space, the datasets that we use, and the implementation details. We then continue with the analysis of the architecture selection on a small dataset, composed of both natural and medical images, denoted as *NAS Dataset for DIP*, and further evaluate image restoration performance for denoising, inpainting, and super-resolution on established datasets for the given restoration task.

4.1. Architecture Selection in NAS Dataset

4.1.1 Experiment Set-up

Search Space. Following the recent works, we select the same search space with NAS-DIP [10] throughout our experiments. The search space consists of different types of upsampling cells and random cross-level feature connections between decoder and encoder cells. Each upsampling cell is defined by five discrete attributes: Spatial feature sampling, feature transformation, kernel size, dilation rate, activation layer. If the connection scale factors are higher than $2\times$, the series of $2\times$ upsampling operations are connected consecutively (e.g for $4 \times$ connection between the decoder and encoder cell, two consecutive 2x upsampling operations are connected). Following the original work [10], each architecture is constituted of 5 encoder and 5 decoder cells. In total, we randomly selected 522 different models in the search space and recorded each model's results for image denoising and image inpainting tasks.

NAS Dataset for DIP. To evaluate how well our metrics perform a model selection in the search space, we created a *NAS Dataset for DIP* using the images in Fig 3.

The dataset is composed of 10 images, 8 are selected from either "BM3D Dataset [22]" or "Dataset of Standard 512 × 512 Grayscale Test Images"¹. The images "blood"² and "chest" [23] are appended to offer domain diversity. Each image is converted to grayscale and resized to 512×512 . Afterwards, for the image denoising task, Gaussian noise with $\sigma = 25$ is applied; for image inpainting, a



Figure 3: Images in *NAS Dataset for DIP*. Images are selected from BM3D Dataset [22], Dataset of Standard 512×512 Grayscale Test Images¹, BCCD dataset², and Chest X-Ray Images (Pneumonia) [23]

Bernoulli mask with 50% is applied. Analogous to NAS-Bench 201 [20], *NAS Dataset for DIP* allows us to measure the capability of ISNAS algorithm. We optimized the 522 models for all the generated images in this dataset for both denoising and inpainting tasks.

Implementation Details. While training the models, we use Adam as an optimizer with a constant learning rate of 0.01. We choose the stopping point as 1200 for denoising, 9500 for inpainting, and 4500 for super-resolution.

We use the final PSNR score as our evaluation metric and follow the same procedure with DIP [5], where an exponential sliding window is applied to the resulting images.

4.1.2 Architecture Selection Accuracy

Using the *NAS Dataset for DIP* and the 522 models, we empirically determined that the same model architectures behave differently to different images under the DIP setting (see Figures 1a and 1b, and Supplementary for further results).

Building on the insight that image restoration tasks benefit from image-specific neural network architectures, we investigate how accurate the suggested metrics are at identifying the best network architecture for the specific image. Fig. 4 shows the distribution of the final PSNR scores of the 522 models for both image denoising (first row) and image inpainting (second row). We observe that the found model by PSD DB Strip MSE with N=15 (the best-performing one in 15 models is depicted with red bin) is always located on the right side of the peak value for all images. In addition, the black bin shows the same score for N=5 case. Even though the latter case is worse than N=15, it is still competitive for some cases. Thus, PSD DB Strip MSE has the potential to reduce the size of the search space into only a few models.

Table 1 shows the correlation between the final PSNR scores and the values of the corresponding metrics for image denoising task. These correlations are extracted from

¹https://ccia.ugr.es/cvg/CG/base.htm

² Taken from Link and is under MIT licence



Figure 4: Histograms of the final PSNR scores of models for five images in *the NAS Dataset for DIP* for image denoising (first row) and image inpainting (second row). Non-blue colored bars represent the final PSNR scores of the best models included in top-N models found by PSD DB Strip MSE metric. The bin depicted with red denotes the best model among N=15 models, the black bin denotes the best one among N=5 models. Histograms of the remaining images and random selection results are provided in supplementary.

Images	PSE M	PSD DB MSE		PSD DB STRIP MSE		STRIP EMD	PSD 99 PER BW
	N	GT	N	GT	N	GT	N/GT
misc	-0.02	-0.21	-0.20	-0.20	-0.12	-0.12	-0.17
lena	0.15	-0.25	-0.37	-0.37	-0.26	-0.26	-0.30
barbara	0.05	-0.13	-0.20	-0.20	-0.13	-0.13	-0.18
couple	0.12	-0.24	-0.36	-0.36	-0.25	-0.25	-0.30
chest	0.25	-0.25	-0.32	-0.07	-0.29	-0.29	-0.34
man	0.10	-0.22	-0.34	-0.34	-0.24	-0.24	-0.27
palace	0.08	-0.20	-0.30	-0.30	-0.20	-0.20	-0.24
blood	0.30	-0.27	0.05	0.39	-0.28	-0.18	-0.34
f16	0.09	-0.20	-0.31	-0.31	-0.21	-0.21	-0.26
boat	0.10	-0.22	-0.34	-0.34	-0.23	-0.23	-0.27

Table 1: Kendall correlation values between the final PSNR scores and corresponding metrics for image denoising task. MSE metrics are calculated either with the noisy image (denoted as N) or with the ground truth image (denoted as GT).

the NAS Dataset for DIP. For most of the images, the metrics calculated with ground truth (GT) correlate reasonably well with model performance. It follows that the metrics can be used to reduce the significant number of models to be trained in a NAS algorithm. However, in a practical situation, we must calculate the metrics using corrupted images since ground truth images are not accessible. In this case, the correlation of PSD DB MSE metric drops, but others still remain almost the same. The comparison between PSD DB MSE and PSD DB Strip MSE shows that focusing on specific parts of the PSDs can be used to address the correlation drop and results in higher correlation than focusing on the whole PSD. Also, images from different domains respond very differently to the metrics. For example, the correlation values for the blood and chest images are dissimilar to each other. All these observations highlight the benefits of concentrating on specific regions of PSDs rather than the whole.

4.2. Experiments on Image Restoration Datasets

In this section, we run experiments on commonly used datasets for image denoising, image inpainting, and singleimage super-resolution tasks. To test our algorithm on conventional datasets, we take a subset of 5000 models from the search space and calculate the image-specific metrics for each model. For each metric, we sort the models according to their metric scores and perform selection among top-N models. Note that for the following experiments, we used the same 5000 models. In the quantitative analysis part, we describe and compare our metrics from a practical perspective.

Image Restoration Set-Up. To compare the performance of the models selected by the proposed metrics with that of state-of-the-art work [10], we evaluate on well adopted datasets. For image denoising, we use BM3D [22], Set12 [1] and CBM3D [22] datasets and apply Gaussian noise with $\sigma = 25$. For image inpainting, we use BM3D [22] and Set12 [1] datasets under 50% missing pixels setting. We use Set5 [24] and Set14 [25] datasets for the super-resolution experiments under three upsampling scales: $2\times$, $4\times$, $8\times$. We reproduce the results of NAS-DIP [10] and DIP [5] with the code provided by the authors [10, 5], respectively. This allows a fair analysis by performing optimizations under same conditions.

4.2.1 Quantitative Analysis

In Table 2, we report the PSNR scores of DIP [5], NAS-DIP [10] and our metrics. We employ full-sized averaging selection among 15 models for our metrics. For denoising, PSD DB Strip MSE metric takes first place for all datasets. Also, the idea of fighting against the correlation drop that arises when using the corrupted image while calculating the metrics is consistent with the findings presented

Datasets	DIP [5]	NAS-DIP [10]	PSD DB MSE (15 ⁿ)	PSD DB Strip MSE (15 ⁿ)	PSD Strip Hist EMD (15 ⁿ)	99 BW (15 ⁿ)		
Denoising ($\sigma = 25$)								
BM3D [22]	27.87	27.44	24.31	28.39	27.81	27.14		
Set12 [1]	27.92	26.88	23.77	28.06	27.81	26.86		
CBM3D [22]	28.93	29.13	<u>30.01</u>	30.36	29.04	28.28		
Inpainting								
BM3D [22]	31.04	30.55	32.90	32.32	32.75	30.07		
Set12 [1]	31.00	30.86	32.22	31.68	32.04	29.91		
Super-resolution								
Set5 ×2 [24]	33.19*	36.16	<u>34.83</u>	34.72	33.98	28.84		
Set5 ×4 [24]	29.89^{*}	30.66	<u>30.22</u>	30.05	30.03	25.23		
Set5 ×8 [24]	25.88^{*}	25.88	25.82	25.94	25.74	25.46		
Set14 ×2 [25]	29.80^{*}	32.11	30.74	30.87	29.71	25.63		
Set14 ×4 [25]	27.00^{*}	27.36	27.19	27.03	27.15	25.03		
Set14 ×8 [25]	24.15*	23.96	23.97	23.90	<u>24.03</u>	23.21		

Table 2: Final PSNR scores of DIP [5], NAS-DIP [10], and metrics on denoising, inpainting, and super-resolution tasks. * denotes the PSNR scores that are evaluated at the optimal-stopping point (where we have access to ground truth). ⁿ denotes that the *full-sized averaging* selection technique is used. **Bold** denotes the highest scores, <u>underline</u> denotes the 2nd highest scores along the row. Note that super-resolution scores of DIP [5] are directly taken from NAS-DIP paper [10].

in Table 2. There is a large gap between PSD DB MSE and its stripped version for gray-scale datasets since eliminating the very high regions by applying a strip mask alleviates the effects of noise on the PSD, which in turn enables finding better models. Conversely, in super-resolution, they are marginally different from each other and, one possible explanation could be that our method is prone to select models with low pass characteristics. Super-resolution requires producing fine details, however our selected model produces smoother outputs failing to produce fine details (see Supplementary), indicating the need for a different metric for super-resolution. This observation is more evident in the 99 BW metric (Table 2) since models with low pass characteristics are explicitly selected leading to the worst performance amongst the metrics. For inpainting, all the image-specific metrics have outperformed both DIP [5] and NAS-DIP [10]. Overall, the metrics give promising results for denoising and inpainting tasks. Furthermore, the performance of our image independent metric (99 BW) supports the need for image-specific models since 99 BW never outperforms any one of the stripped metrics for all image restoration tasks.

We further compare the performance of resizedaveraging selection technique with both full-sized averaging selection and a random baseline. As the random baseline, 15 models are selected randomly from the search space, then resized-averaging is applied to perform the selection. Table 3 shows the final PSNR scores of the random

			PSD DB Strip MSE					
Datasets	Random ^r	NAS-DIP	15 ^r	15 ⁿ				
Denoising ($\sigma = 25$)								
BM3D [22]	23.20	27.44	<u>27.66</u>	28.39				
Set12 [1]	22.66	26.88	27.40	28.06				
CBM3D [22]	27.86	29.13	<u>29.28</u>	30.36				
Inpainting								
BM3D [22]	32.30	30.55	33.01	32.32				
Set12 [1]	31.64	30.86	32.17	<u>31.68</u>				

Table 3: 2 different averaging techniques are used for PSD DB Strip MSE metric. ^r denotes the resized averaging selection, ⁿ denotes the full-sized averaging selection. "Random" denotes the random selection over search space. **Bold** denotes the highest scores, <u>underline</u> denotes the 2nd highest scores along the row.

baseline, NAS-DIP [10], and PSD DB Strip MSE metric evaluated under two selection procedures. It demonstrates that PSD DB Strip MSE metric performs significantly better than random selection. Additionally, owing to the imagespecific model need, it performs better for all datasets.

There is a tradeoff between the speed and the performance of averaging techniques, especially in image denois-



Figure 5: Qualitative examples of image denoising (1^{st} row) , image inpainting (2^{nd} row) , and single-image super-resolution (3^{rd} row) tasks. Each column shows the images of ground truth, corrupted image (Bicubic [26] for super-resolution), reconstructed images of DIP [5], NAS-DIP [10], and ISNAS-DIP, respectively.

ing. For a 256×256 grayscale image, the optimization speed is approximately $5 \sim 8$ it/s whereas it rises to ~ 20 it/s when the image is resized to 64×64 . For a selection among 15 models for denoising, it requires approximately 15 minutes for resized-averaging and $37.5 \sim 60$ minutes for full-sized averaging technique. These cases are much faster than NAS-DIP [10] selection since the authors claim that the search lasted 3 days for denoising in total.

We note that the reason we compared with NAS-DIP in our experiments is that they have outperformed previous learning-free DIP approaches.

4.2.2 Qualitative Analysis

Figure 5 shows qualitative samples of several images for all tasks. Generally, we observe apparent visual improvements in the quality of the restored image. For example, in denoising, fine details such as the notch on top of the fin of F16 are retrieved better with ISNAS-DIP. However, in super-resolution, since NAS-DIP [10] performs better, we do not observe much improvement. The improvements in denoising and inpainting emphasize the importance of taking into consideration the content of an input image while designing network architectures.

5. Conclusion, Limitations and Future Work

In this work we show that the optimal architecture to exploit the deep image prior is image dependent. Based on this insight, we propose several metrics that allow for a fast image-specific neural network architecture search and show that the found models outperform benchmark architectures for image restoration tasks using DIP. Furthermore, we establish a *NAS Dataset for DIP* that can be used for future research investigating image-specific NAS.

There are still some limitations that future work entails. We show that for N=15, our metrics are able to select a subset that contains optimal image-specific models. That is, we narrow down the whole search space to just 15 models. However, we are still limited in selecting the top-1 model. The current averaging technique needs N models to be optimized. We attempt to boost the optimization process by using the resized image while optimizing 15 models, which is shown to expedite the process. Yet, it does not guarantee that the selected one is the best among this subset, but is the one that is close to the average. Hence, selecting the top-1 model among a subset of models is still an open question.

One other issue is the early stopping problem of DIP. In our experiments, we use a fixed number of iterations for each restoration task. We observe a significant correlation between the proposed metrics and the optimal number of iterations for each image (see Supplementary). This means that the metrics might also be useful to determine the early stopping point for the training of selected models.

Moreover, the *NAS Dataset for DIP* experiments reveal the need for better-suited metrics for different types of images since the domains of images influence the correlations. Future work entails implementing *learned* metrics rather than hand-designed definitions as in our case. A deep learning model can be utilized in order to do so. Using a search space, a metric can be described by a *neural network*, which can be trained with the CNNs' outputs and corrupted images to obtain a better formulation of a metric.

References

- K. Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing*, 26:3142–3155, 2017.
- [2] K. Zhang, Wangmeng Zuo, Shuhang Gu, and Lei Zhang. Learning deep cnn denoiser prior for image restoration. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 2808–2817, 2017.
- [3] K. Zhang, Wangmeng Zuo, and Lei Zhang. Ffdnet: Toward a fast and flexible solution for cnn-based image denoising. *IEEE Transactions on Image Processing*, 27:4608– 4622, 2018.
- [4] Shi Guo, Zifei Yan, Kai Zhang, Wangmeng Zuo, and Lei Zhang. Toward convolutional blind denoising of real photographs. 2019 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019.
- [5] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), June 2018.
- [6] Andrew M. Saxe, Pang Wei Koh, Zhenghao Chen, Maneesh Bhand, Bipin Suresh, and Andrew Y. Ng. On random weights and unsupervised feature learning. In *Proceedings* of the 28th International Conference on International Conference on Machine Learning, ICML'11, page 1089–1096, Madison, WI, USA, 2011. Omnipress.
- [7] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, Proceedings of the 36th International Conference on Machine Learning, volume 97 of Proceedings of Machine Learning Research, pages 5301–5310. PMLR, 09–15 Jun 2019.
- [8] Prithvijit Chakrabarty and Subhransu Maji. The spectral bias of the deep image prior. *ArXiv*, abs/2107.01125, 2019.
- [9] Gary Mataev, Michael Elad, and Peyman Milanfar. Deepred: Deep image prior powered by red. *ArXiv*, abs/1903.10176, 2019.
- [10] Yun-Chun Chen, Chen Gao, Esther Robb, and Jia-Bin Huang. Nas-dip: Learning deep image prior with neural architecture search. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 442–459, Cham, 2020. Springer International Publishing.
- [11] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. J. Mach. Learn. Res., 20:55:1–55:21, 2019.
- [12] Barret Zoph and Quoc V. Le. Neural architecture search with reinforcement learning. *ArXiv*, abs/1611.01578, 2017.
- [13] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. Learning transferable architectures for scalable image recognition. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8697–8710, 2018.

- [14] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. Efficient neural architecture search via parameters sharing. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4095–4104. PMLR, 10–15 Jul 2018.
- [15] M. Suganuma, M. Ozay, and T. Okatani. Exploiting the potential of standard convolutional autoencoders for image restoration by evolutionary search. In *ICML*, 2018.
- [16] Haokui Zhang, Youming Li, Hao Chen, and Chunhua Shen. Memory-efficient hierarchical neural architecture search for image denoising. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 3654–3663, 2020.
- [17] Joseph Charles Mellor, Jack Turner, Amos J. Storkey, and Elliot J. Crowley. Neural architecture search without training. *ArXiv*, abs/2006.04647, 2021.
- [18] Wuyang Chen, Xinyu Gong, and Zhangyang Wang. Neural architecture search on imagenet in four gpu hours: A theoretically inspired perspective. In *International Conference* on Learning Representations, 2021.
- [19] Miao Zhang, Steven Su, Shirui Pan, Xiaojun Chang, Wei Huang, and Gholamreza Haffari. Differentiable architecture search without training nor labels: A pruning perspective. *ArXiv*, abs/2106.11542, 2021.
- [20] Xuanyi Dong and Yi Yang. Nas-bench-201: Extending the scope of reproducible neural architecture search. ArXiv, abs/2001.00326, 2020.
- [21] Reinhard Heckel and Mahdi Soltanolkotabi. Denoising and regularization via exploiting the structural bias of convolutional generators. *ArXiv*, abs/1910.14634, 2020.
- [22] Kostadin Dabov, Alessandro Foi, and Karen O. Egiazarian. Video denoising by sparse 3d transform-domain collaborative filtering. 2007 15th European Signal Processing Conference, pages 145–149, 2007.
- [23] Daniel Kermany. Large dataset of labeled optical coherence tomography (oct) and chest x-ray images, 2018.
- [24] Marco Bevilacqua, Aline Roumy, Christine Guillemot, and Marie-Line Alberi-Morel. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. In *BMVC*, 2012.
- [25] Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In Jean-Daniel Boissonnat, Patrick Chenin, Albert Cohen, Christian Gout, Tom Lyche, Marie-Laurence Mazure, and Larry Schumaker, editors, *Curves and Surfaces*, pages 711–730, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- [26] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38:295–307, 2016.