

Deep Network Interpolation for Continuous Imagery Effect Transition

Xintao Wang¹ Ke Yu¹ Chao Dong² Xiaoou Tang¹ Chen Change Loy³

¹CUHK - SenseTime Joint Lab, The Chinese University of Hong Kong

²SIAT-SenseTime Joint Lab, Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences

³Nanyang Technological University, Singapore

{wx016, yk017, xtang}@ie.cuhk.edu.hk chao.dong@siat.ac.cn ccloy@ntu.edu.sg



Figure 1. Deep network interpolation is capable of generating continuous imagery effect transitions. (1st row) from MSE effect to GAN effect in super-resolution; (2nd row) from Van Gogh style to Ukiyo-e style; (3rd row) from day photo to night one; (4th row) from deep depth of field (DoF) to shallow one. More applications are provided in Sec. 4. (Zoom in for best view)

Abstract

Deep convolutional neural network has demonstrated its capability of learning a deterministic mapping for the desired imagery effect. However, the large variety of user flavors motivates the possibility of continuous transition among different output effects. Unlike existing methods that require a specific design to achieve one particular transition (e.g., style transfer), we propose a simple yet universal approach to attain a smooth control of diverse imagery effects in many low-level vision tasks, including image restoration, image-to-image translation, and style transfer. Specifically, our method, namely Deep Network Interpolation (DNI), applies linear interpolation in the parameter space of two or more correlated networks. A smooth control of imagery effects can be achieved by tweaking the interpolation coefficients. In addition to DNI and its broad applications, we also investigate the mechanism of network interpolation from the perspective of learned filters.

1. Introduction

Deep convolutional neural (CNN) network has achieved a great success in many low-level vision tasks, such as image restoration [3, 18, 20, 1, 44, 2, 28], image-to-image translation [15, 37, 26, 47] and image style transfer [8, 17, 6]. For each specific task, the deep network learns a deterministic mapping and outputs a fixed image for the same inputs. However, one determined output is unable to satisfy diverse user flavors and meet the needs in various scenarios, limiting the applicability for practical use.

In many real-world applications, it is desired to have a smooth control for continuous transition among different output effects. For instance, 1) in super-resolution, models trained with the mean-square-error (MSE) loss [34] tend to produce over-smooth images while those with the generative adversarial network (GAN) [20] generate vivid details but with some unpleasant noise (e.g., Fig 1, 1st row). A balanced result between these two different effects would be more visual-pleasing with reduced artifacts. 2) Many image restoration tasks deal with multiple degradation levels, such as different noise levels and blur kernels. Most ex-

isting methods can only handle limited degradation levels. It is costly to train lots of models for continuous degradation levels in practice. Thus, a model with the flexibility of adjusting the restoration strength would expand the application coverage. 3) In artistic manipulation like image-to-image translation and image style transfer, different users have different aesthetic flavors. Achieving a smooth control for diverse effects with a sliding bar are appealing in these applications.

Several approaches have been proposed to improve the CNN’s flexibility for producing continuous transitions in different tasks. Take image style transfer as an example, adaptive scaling and shifting parameters are used in instance normalization layers [6, 12] for modeling different styles. Interpolating these normalization parameters for different styles produces the combination of various artistic styles. In order to further control the stroke size in the stylized results, a carefully-designed pyramid structure consisting of several stroke branches are proposed [16]. Though these methods are able to realize continuous transition, there are several drawbacks: 1) These careful designs are problem-specific solutions, lacking the generalizability to other tasks. 2) Modifications to existing networks are needed, thus complicate the training process. 3) There is still no effective and general way to solve the smooth control in tasks like balancing MSE and GAN effects in super-resolution.

In this paper, we address these drawbacks by introducing a more general, simple but effective approach, known as *Deep Network Interpolation* (DNI). Continuous imagery effect transition is achieved via linear interpolation in the *parameter space* of existing trained networks. Specifically, provided with a model for a particular effect \mathcal{A} , we fine-tune it to realize another relevant effect \mathcal{B} . DNI applies linear interpolation for all the corresponding parameters of these two deep networks. Various interpolated models can then be derived by a controllable interpolation coefficient. Performing feed-forward operations on these interpolated models using the same input allows us to outputs with a continuous transition between the different effects \mathcal{A} and \mathcal{B} .

Despite its simplicity, the proposed DNI can be applied to many low-level vision tasks. Some examples are presented in Fig 1. Extensive applications showcased in Sec. 4 demonstrate that deep network interpolation is *generic* for many problems. DNI also enjoys the following merits. 1) The transition effect is *smooth* without abrupt changes during interpolation. The transition can be easily controlled by an interpolation coefficient. 2) The linear interpolation operation is *simple*. No network training is needed for each transition and the computation for DNI is negligible. 3) DNI is *compatible* with popular network structures, such as VGG [31], ResNet [10] and DenseNet [11].

Our main contribution in this work is the novel notion of interpolation in parameter space, and its application in

low-level vision tasks. We demonstrate that interpolation in the parameter space could achieve much better results than mere pixel interpolation. We further contribute a systematic study that investigates the mechanism and effectiveness of parameter interpolation through carefully analyzing the filters learned.

2. Related Work

Image Restoration. CNN-based approaches have led to a series of breakthroughs for several image restoration tasks including super-resolution [3, 18, 25, 19, 33, 20, 43], denoising [1, 44], de-blocking [40, 7] and deblurring [42, 32, 28]. While most of the previous works focus on addressing one type of distortion without the flexibility of adjusting the restoration strength, there are several pioneering works aiming to handle various practical scenarios with controllable “hyper-parameters”. Zhang *et al.* [45] adopt CNN denoisers to solve image restoration tasks by manually selecting the hyper-parameters in a model-based optimization framework. However, a bank of discriminative CNN denoisers are required and the hyper-parameter selection in optimization is not a trivial task [4]. SRMD [46] proposes an effective super-resolution network handling multiple degradations by taking an degradation map as extra inputs. However, the employed dimensionality stretching strategy is problem-specific, lacking the generalizability to other tasks.

Image Style Transfer. Gatys *et al.* [8] propose the neural style transfer algorithm for artistic stylization. A number of methods are developed to further improve its performance and speed [35, 17, 21]. In order to model various/arbitrary styles in one model, several techniques are developed, including conditional instance normalization [6], adaptive instance normalization [12, 9] and whitening and coloring transforms [23]. These carefully-designed approaches are also able to achieve user control. For example, interpolating the normalization parameters of different styles produces the combination of various artistic styles. The balance of content and style can be realized by adjusting the weights of their corresponding features during mixing. In order to control the stroke size, a specially designed pyramid structure consisting of several stroke branches are further proposed [16]. In these studies, different controllable factors require specific structures and strategies.

Image-to-image Translation. Image-to-image translation [15, 37, 26, 47, 13] aims at learning to translate an image from one domain to another. For instance, from landscape photos to Monet paintings, and from smartphone snaps to professional DSLR photographs. These methods can only transfer an input image to a specific target manifold and they are unable to produce continuous translations. The controllable methods proposed in image restoration and image style transfer are problem-specific and cannot be di-

rectly applied to the different image-to-image translation task. On the contrary, the proposed DNI is capable of dealing with all these problems in a general way, regardless of the specific characteristics of each task.

Interpolation. Instead of performing parameter interpolation, one can also interpolate in the pixel space or feature space. However, it is well known that interpolating images pixel by pixel introduces ghosting artifacts since natural images lie on a non-linear manifold [41]. Upchurch *et al.* [36] propose a linear interpolation of pre-trained deep convolutional features to achieve image content changes. This method requires an optimization process when inverting the features back to the pixel space. Moreover, it is mainly designed for transferring facial attributes and not suitable for generating continuous transition effects for low-level vision tasks. More broadly, several CNN operations in the input and feature space have been proposed to increase the model’s flexibility. Concatenating extra conditions to inputs [46] or to middle features [22] alters the network behavior in various scenes. Modulating features with an affine transformation [29, 5, 38] is able to effectively incorporate other information. Different from these works, we make an attempt to investigate the manipulation in the parameter space. A very preliminary study for network interpolation is presented by ESRGAN [39], focusing to enhance competition results. DNI provides more comprehensive investigations and extends to more applications.

3. Methodology

3.1. Deep Network Interpolation

Many low-level vision tasks, *e.g.*, image restoration, image style transfer, and image-to-image translation, aim at mapping a corrupted image or conditioned image x to the desired one y . Deep convolutional neural networks are applied to directly learn this mapping function G_θ parametrized by θ as $y = G_\theta(x)$.

Consider two networks G^A and G^B with the same structure, achieving different effects \mathcal{A} and \mathcal{B} , respectively. The networks consist of common operations such as convolution, up/down-sampling and non-linear activation. The parameters in CNNs are mainly the weights of convolutional layers, called *filters*, filtering the input image or the precedent features. We assume that their parameters θ_A and θ_B have a “strong correlation” with each other, *i.e.*, the filter orders and filter patterns in the same position of G^A and G^B are similar. This could be realized by some constraints like fine-tuning, as will be analyzed in Sec. 3.2. This assumption provides the possibility for meaningful interpolation.

Our aim is to achieve a continuous transition between the effects \mathcal{A} and \mathcal{B} . We do so by the proposed Deep Network Interpolation (DNI). DNI interpolates all the corresponding parameters of these two models to derive a new interpolated

model G^{interp} , whose parameters are:

$$\theta_{interp} = \alpha \theta_A + (1 - \alpha) \theta_B, \quad (1)$$

where $\alpha \in [0, 1]$ is the interpolation coefficient. Indeed, it is a linear interpolation of the two parameter vectors θ_A and θ_B . The interpolation coefficient α controls a balance of the effect \mathcal{A} and \mathcal{B} . By smoothly sliding α , we achieve continuous transition effects without abrupt changes.

Generally, DNI can be extended for N models, denoted by G^1, G^2, \dots, G^N , whose parameters have a “close correlation” with each other. The DNI is then formulated as:

$$\theta_{interp} = \alpha_1 \theta_1 + \alpha_2 \theta_2 + \dots + \alpha_N \theta_N, \quad (2)$$

where α_i satisfy $\alpha_i \geq 0$ and $\alpha_1 + \alpha_2 + \dots + \alpha_N = 1$. In other words, it is a convex combination of the parameter vectors $\theta_1, \theta_2, \dots, \theta_N$. By adjusting $(\alpha_1, \alpha_2, \dots, \alpha_N)$, a rich and diverse effects with continuous transitions could be realized.

The interpolation is performed on all the layers with parameters in the networks, including convolutional layers and normalization layers. Convolutional layers have two parameters, namely *weights* (filters) and *biases*. The biases are added to the results after filtering operation with filters. Apart from the weights, DNI also operates on the biases, since the added biases influence the successive non-linear activation.

Batch normalization (BN) layers [14] have two kinds of parameters. 1) The statistics *running_mean* and *running_var* track the mean and variance of the whole dataset during training and are then used for normalization during evaluation. 2) the learned parameters γ and β are for a further affine transformation. During inference, all these four parameters actually could be absorbed into the precedent or successive convolutional layers. Thus, DNI also performs on normalization parameters. Instance normalization (IN) has a similar behavior as BN, except that IN uses instance statistics computed from input data in both training and evaluation. We take the same action as that for BN. In practice, the interpolation is performed on not only the weights but also the biases and further normalization layers. We believe a better interpolation scheme considering the property of different kinds of parameters is worthy of exploiting.

It is worth noticing that the choice of the network structure for DNI is flexible, as long as the structures of models to be interpolated are kept the same. Our experiments on different architectures show that DNI is compatible with popular network structures such as VGG [31], ResNet [10] and DenseNet [11]. We also note that the computation of DNI is negligible. The computation is only proportional to the number of parameters.




Figure 2. Filter correlations. The first two rows are the filters of different runs (both from scratch) for the denoising (N20) task. The filter orders and filter patterns in the same position are different. The *fine-tuned* model (N60) (3rd row) has a “strong correlation” to the pre-trained one (1st row).

3.2. Understanding Network Interpolation

We attempt to gain more understanding on network interpolation from some empirical studies. From our experiments, we observe that: 1) Fine-tuning facilitates high correlation between parameters of different networks, providing the possibility for meaningful interpolation. 2) Fine-tuned filters for a series of related tasks present continuous changes. 3) Our analyses show that interpolated filters could fit the actual learned filters well. Note that our analyses mainly focus on *filters* since most of the parameters in CNNs are in the form of filters.

We present our main observations with a representative denoising task and focus on increasing noise levels with N20, N30, N40, N50, and N60, where N20 denotes the Gaussian noise with zero mean and variance 20. In order to better visualize and analyze the filters, we adopt a three-layer network similar to SRCNN [3], where the first and last convolutional layers have 9×9 filter size. Following the notion of [3], the first and last layer layers can be viewed as a *feature extraction* and *reconstruction* layer, respectively.

Fine-tuning for inter-network correlation. Even for the same task like denoising with the N20 level, if we simply train two models from scratch, the filter orders among channels and filter patterns in the corresponding positions could be very different (Fig. 2). However, a core representation is shared between these two networks [24]. For instance, in Fig. 2, filter *c* is identical to filter *f*; filter *a* and filter *e* have a similar pattern but with different colors; filter *b* is a inverted and rotated counterpart of filter *d*.

Fine-tuning, however, can help to maintain the filters’s order and pattern. To show this, we fine-tune a pre-trained network (N20) to a relevant task (N60). It is observed that the filter orders and filter patterns are maintained (Fig. 2). The “high correlation” between the parameters of these two networks provides the possibility for meaningful interpolation. We note that besides fine-tuning, other constraints such as joint training with regularization may also achieve such inter-network correlation.

Learned filters for related tasks exhibit continuous changes. When we fine-tune several models for relevant

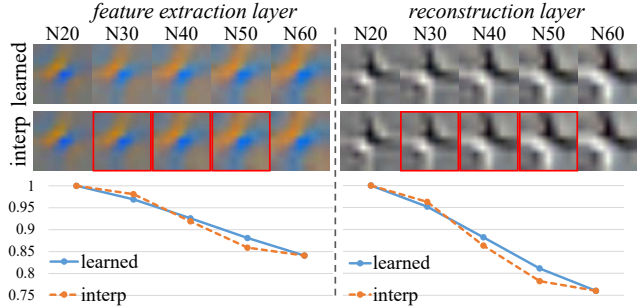


Figure 3. Filters with gradual changes. We show one representative filter for each layer. **1st row:** The fine-tuned filters for different noise level show gradual changes. **2nd row:** The interpolated filters (with red frames) from the N20 and N60 filters could visually fit those learned filters well. **3rd row:** The correlation curves for learned and interpolated filters are also very close.

tasks (N30, N40, N50, and N60) from a pre-trained one (N20), the correspondingly learned filters have intrinsic relations with a smooth transition. As shown in Fig. 3 (1st row), the trained filters show gradual changes as the noise level increases. Apart from visualization, we also calculate a correlation index ρ_{ij} to measure the correlations of filters F_i and F_j :

$$\rho_{ij} = \frac{(F_i - \bar{F}_i) \cdot (F_j - \bar{F}_j)}{\sqrt{\|F_i - \bar{F}_i\|_2} \sqrt{\|F_j - \bar{F}_j\|_2}}. \quad (3)$$

We choose this form (similar to the Pearson correlation) since it ignores the scale and shift influences and focus on the filter pattern itself. We calculate the correlation index for each filter with the first N20 filter and plot the curve (blue curve in Fig. 3). The results suggest a close relationship among the learned filters, exhibiting a gradual change as the noise level increases.

The interpolated filters fit the learned filters well. The continuous changes of learned filters suggest that *it is possible to obtain the intermediate filters by interpolating the two ends*. To further verify this observation, we perform linear interpolation between the filters from the N20 and N60 models. With optimal coefficients α , the interpolated filters could visually fit those learned filters (2nd row with red frames, Fig. 3). We further calculate the correlation index for each interpolated filter with the first N20 filter. The correlation curves for learned and interpolated filters are also very close.

The optimal α is obtained through the final performance of the interpolated network. Specifically, we perform DNI with α from 0 to 1 with an interval of 0.05. The best α for each noise level is selected based on which that makes the interpolated network to produce the highest PSNR on the test dataset.

Discussion. It is noteworthy that similar observations could also be found in several other tasks, such as super-resolution with different kernels and JPEG artifacts removal with dif-

ferent compression levels. We provide the details in the supplementary material.

The analysis above is by no means complete, but it gives a preliminary explanation behind the DNI from the filter perspective. As the network goes deeper, the non-linearity increases and the network behaviors become more complicated. However, we still observe a similar phenomenon for deeper networks. Since filter visualization is difficult in deep networks, which typically designed with a stack of convolutional layers of 3×3 kernels, we adopt the correlation index (Eq. 3) to analyze the filter correlations among models for a series of noise levels. We employ DnCNN [44] with 17 layers and analyze the 5th (the front) and 12th (the back) convolutional layers. In Fig. 4, the correlation curves show the median of correlation indexes w.r.t. the first N20 model and the correlation distribution are also plotted. Besides the high correlation among these models, we can also observe gradual changes as the noise level increases. Furthermore, the front and the back convolution layers present a similar transition trend, even their distributions highly coincide.

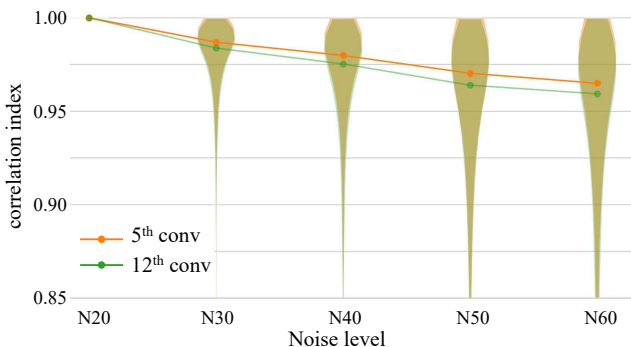


Figure 4. Filter correlation for deeper denoising networks. We show a front (5th) and a back (12th) layer. The curves present the median of correlation indexes and the correlation distribution are also plotted. Note that the two distributions highly coincide. Zoom in for subtle differences.

4. Applications

We experimentally show that the proposed DNI can be applied to extensive low-level vision tasks, *e.g.*, image restoration (Sec. 4.1), image-to-image translation (Sec. 4.2) and image style transfer (Sec. 4.3). Another example of smooth transitions on face attributes are presented in Sec. 4.4, indicating its potential for semantic changes. Due to space limitations, more examples and analyses are provided in the supplementary material and the project page¹.

4.1. Image Restoration

Balance MSE and GAN effects in super-resolution. The aim of super-resolution is to estimate a high-resolution im-

age from its low-resolution counterpart. A super-resolution model trained with MSE loss [34] tends to produce over-smooth images. We fine-tune it with GAN loss and perceptual loss [20], obtaining results with vivid details but always together with unpleasant artifacts (*e.g.*, the eaves and water waves in Fig. 5). We use dense blocks [11, 39] in the network architecture and the MATLAB bicubic kernel with a scaling factor of 4 is adopted as the down-sampling kernel.

As presented in Fig. 5, DNI is able to smoothly alter the outputs from the MSE effect to the GAN effect. With appropriate interpolation coefficient, it produces visually-pleasing results with largely reduced artifacts while maintaining the textures. We also compare it with pixel interpolation, *i.e.*, interpolating the output images pixel by pixel. However, the pixel interpolation is incapable of separating the artifacts and details. Taking the water waves for example (Fig. 5), the water wave texture and artifacts simultaneously appear and enhance during the transition. Instead, DNI first enhances the vivid water waves without artifacts and then finer textures and undesirable noise appears successively. The effective separation helps to remove the displeasing artifacts while keeping the favorable textures, superior to the pixel interpolation. This property could also be observed in the transition of animal fur in Fig. 5, where the main structure of fur first appears followed by finer structures and subtle textures.

One can also obtain several models for different mixture effects, by tuning the weights of MSE loss and GAN loss during training. However, this approach requires tuning on the loss weights and training many networks for various balances, and thus it is too costly to achieve continuous control.

Adjust denoising strength. The goal of denoising is to recover a clean image from a noisy observation. In order to satisfy various user demands, most popular image editing softwares (*e.g.*, Photoshop) have controllable options for each tool. For example, the noise reduction tool comes with sliding bars for controlling the denoising strength and the percentage of preserving or sharpening details.

We show an example to illustrate the importance of adjustable denoising strength. We are provided with a denoising model specialized in addressing a specific Gaussian noise level N40. We use DnCNN [44] as our **implementation**. As shown in Fig. 6, however, the determined outputs (with yellow frames) are not satisfactory due to the different imagery contents. In particular, the denoising strength for the grass is too strong, producing over-smooth results, while in the smooth sky region, it requires a larger strength to remove the undesired artifacts.

Existing deep-learning based approaches fail to meet this user requirements since they are trained to generate deterministic results without the flexibility to control the denoising strength. On the contrary, our proposed DNI is able to achieve adjustable denoising strength by simply tweaking

¹<https://xinntao.github.io/projects/DNI>




Figure 5. Balancing the MSE and GAN effects with DNI in super-resolution. The MSE effect is over-smooth while the GAN effect is always accompanied with unpleasant artifacts (e.g., the eaves and water waves). DNI allows smooth transition from one effect to the other and produces visually-pleasing results with largely reduced artifacts while maintaining the textures. In contrast, the pixel interpolation strategy fails to separate the artifacts and textures. (Zoom in for best view)




Figure 6. Adjustable denoising strength with DNI. One model without adjustment (with yellow frames) is unable to balance the noise removal and detail preservation. For the grass, a weaker denoising strength could preserve more textures while for the sky, the stronger denoising strength could obtain an artifact-free result, improving the visual quality (with red frames). (Zoom in for best view)

the interpolation coefficient α of different denoising models for N20, N40 and N60. For the grass, a weaker denoising strength could preserve more details while in the sky region, the stronger denoising strength could obtain an artifact-free result (red frames in Fig. 6). This example demonstrates the flexibility of DNI to customize restoration results based on the task at hand and the specific user preference.

4.2. Image-to-image Translation

Image-to-image translation aims at learning to translate an image from one domain to another. Most existing approaches [15, 37, 26, 47] can only transfer one input image to several discrete outputs, lacking continuous translation for diverse user flavors. For example, one model may be able to mimic the style of Van Gogh or Cézanne, but translating a landscape photo into a mixed style of these two painters is still challenging.

The desired continuous transition between two painters'

styles can be easily realized by DNI. The popular CycleGAN [47] is used as our implementation. We first train a network capturing characteristics of Van Gogh, and then fine-tune it to produce paintings of Cézanne's style. DNI is capable of generating various mixtures of these two styles given a landscape photo, by adjusting the interpolation coefficient. Fig. 7a presents a smooth transition from Van Gogh's style to Cézanne's style both in the palette and brush strokes. We note that DNI can be further employed to mix styles of more than two painters using Eq. 2. Results are provided in the supplementary material.

In addition to the translation between painting styles of a whole image, DNI can also achieve smooth and natural translation for a particular image region. Fig. 7b shows an example of photo enhancement to generate photos with shallower depth of field (DoF). We train one model to generate flower photos with shallow DoF and then fine-tune it with identity mapping. DNI is then able to produce con-

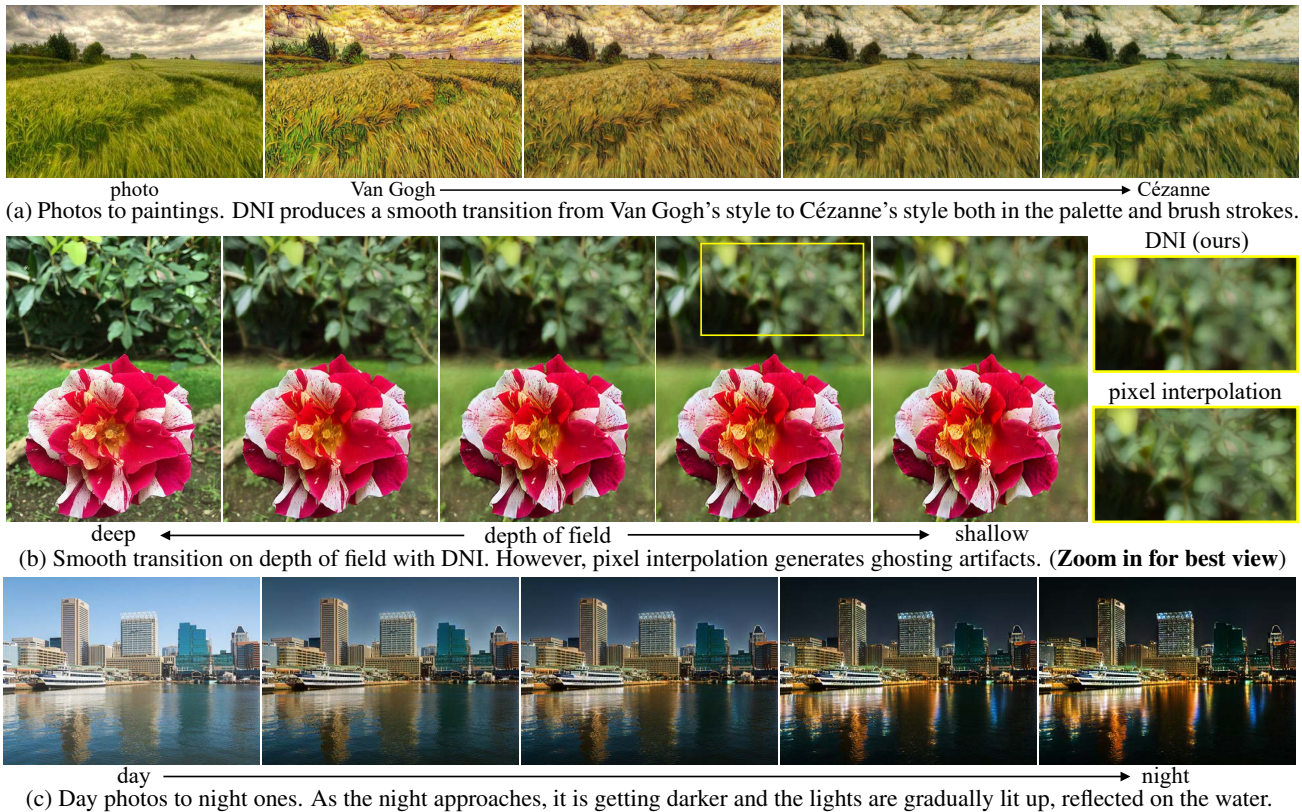


Figure 7. Several applications for image-to-image translation. (Zoom in for best view)

tinuous transitions of DoF by interpolating these two models. We also compare DNI with pixel interpolation, where the results look unnatural due to the ghosting artifacts, *e.g.*, translucent details appearing at the edge of blurry leaves.

DNI can be further applied to achieve continuous imagery translations in other dimensions such as light changes, *i.e.*, transforming the day photos to night ones. Only trained with day and night photos, DNI is capable of generating a series of images, simulating the coming of nights. In Fig. 7c, as the night approaches, it is getting darker and the lights are gradually lit up, reflected on the water.

4.3. Style Transfer

There are several controllable factors when transferring the styles of one or many pieces of art to an input image, *e.g.*, style mixture, stroke adjustment and the balance of content and style. Some existing approaches design specific structures to achieve a continuous control of these factors [16]. On the contrary, DNI offers a general way to attain the same goals without specific solutions. As shown in Fig. 8, DNI is capable of generating smooth transitions between different styles, from large to small strokes, together with balancing the content and style. Furthermore, DNI can be applied among multiple models to achieve a continuous control of various factors simultaneously. For instance, the stroke and style can be adjusted at the same time based on user flavors, as shown in Fig. 8.

Another branch of existing methods achieves a combination of various artistic styles by interpolating the parameters of instance normalization (IN) [6, 12]. These approaches can be viewed as a special case of DNI, where only IN parameters are fine-tuned and interpolated. To clarify the difference between DNI and IN interpolation, we conduct experiments with 3 settings: 1) fine-tune IN; 2) fine-tune convolutional layers and 3) fine-tune both IN and convolutional layers. Specifically, as shown in Fig. 9, we try a challenging task to fine-tune the model from generating images with mosaic styles to the one with fire styles, where the two styles look very different in both color and texture. It is observed that fine-tuning only IN is effective in color transformation, however, it is unable to transfer the fire texture effectively compared with the other two settings. This observation suggests that convolutional layers also play an important role in style modeling, especially for the textures, since IN parameters may not be effective for capturing spatial information. However, we do not claim that DNI is consistently better than IN interpolation, since IN is also effective in most cases. A more thorough study is left for future work.

4.4. Semantic Transition

Apart from low-level vision tasks, we show that DNI can also be applied for smooth transitions on face attributes, suggesting its potential for semantic adjustment. We first




Figure 8. In image style transfer, without specific structures and strategies, DNI is capable of generating smooth transitions between different styles, from large strokes to small strokes, together with balancing the content and style. (Zoom in for best view)

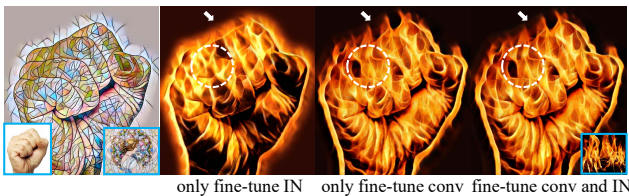


Figure 9. Fine-tuning only IN is effective in color transformation, however, it is unable to transfer the fire texture effectively.




Figure 10. Smooth transitions on face attributes with DNI.

train a DCGAN model [30] using the CelebA [27] dataset with one attribute (*e.g.*, young or male). After that, we fine-tune it to generate faces with another attribute (*e.g.*, old or female). DNI is then able to produce a series of faces with smoothly transformed attributes by interpolating these models (Fig. 10). Although neither of the interpolated models observes any data with middle attributes, the faces in middle states has an intermediate attribute and looks natural.

4.5. Limitations

We have shown that DNI offers a simple yet general approach for transition manipulation. However, DNI also has

limitations - 1) it needs fine-tuning or jointly training, otherwise the interpolated model is meaningless; 2) if the two (or several) tasks have little relation (*i.e.*, DNNs trained for different tasks), the interpolated model could not produce meaningful results. We examine the correlation indexes of successfully interpolated models, which are in $[0.689, 1)$. But for those failure cases, the correlation indexes are extremely small (<0.1). Future work will investigate how to quantify the network ‘distance’ and task ‘distance’ for DNI.

5. Conclusion

In this paper, we propose a novel notion of interpolation in the parameter space, *i.e.*, applying linear interpolation among the corresponding parameters of multiple correlated networks. The imagery effects change smoothly while adjusting the interpolation coefficients. With extensive experiments on super-resolution, denoising, image-to-image translation and style transfer, we demonstrate that the proposed method is applicable for a wide range of low-level vision tasks despite its simplicity. Compared with existing methods that achieve continuous transition by task-specific designs, our method is easy to generalize with negligible computational overhead. Future work will investigate the effects of network interpolation on high-level tasks.

Acknowledgement. This work is supported by SenseTime Group Limited, Joint Lab of CAS-HK, the General Research Fund sponsored by the Research Grants Council of the Hong Kong SAR (CUHK 14241716, 14224316, 14209217), and Singapore MOE AcRF Tier 1 (M4012082.020).

References

- [1] Harold C Burger, Christian J Schuler, and Stefan Harmeling. Image denoising: Can plain neural networks compete with BM3D? In *CVPR*, 2012. 1, 2
- [2] Chao Dong, Yubin Deng, Chen Change Loy, and Xiaoou Tang. Compression artifacts reduction by a deep convolutional network. In *ICCV*, 2015. 1
- [3] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *ECCV*, 2014. 1, 2, 4
- [4] Weisheng Dong, Lei Zhang, Guangming Shi, and Xin Li. Nonlocally centralized sparse representation for image restoration. *TIP*, 22(4):1620–1630, 2013. 2
- [5] Vincent Dumoulin, Ethan Perez, Nathan Schucher, Florian Strub, Harm de Vries, Aaron Courville, and Yoshua Bengio. Feature-wise transformations. *Distill*, 2018. <https://distill.pub/2018/feature-wise-transformations>. 3
- [6] Vincent Dumoulin, Jonathon Shlens, Manjunath Kudlur, Arash Behboodi, Filip Lemic, Adam Wolisz, Marco Molinaro, Christoph Hirsche, Masahito Hayashi, Emilio Bagan, et al. A learned representation for artistic style. In *ICLR*, 2016. 1, 2, 7
- [7] Leonardo Galteri, Lorenzo Seidenari, Marco Bertini, and Alberto Del Bimbo. Deep generative adversarial compression artifact removal. In *ICCV*, 2017. 2
- [8] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *CVPR*, 2016. 1, 2
- [9] Golnaz Ghiasi, Honglak Lee, Manjunath Kudlur, Vincent Dumoulin, and Jonathon Shlens. Exploring the structure of a real-time, arbitrary neural artistic stylization network. In *Bmvc*, 2017. 2
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2, 3
- [11] Gao Huang, Zhuang Liu, Kilian Q Weinberger, and Laurens van der Maaten. Densely connected convolutional networks. In *CVPR*, 2017. 2, 3, 5
- [12] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, 2017. 2, 7
- [13] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *ECCV*, 2018. 2
- [14] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICMR*, 2015. 3
- [15] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. 1, 2, 6
- [16] Yongcheng Jing, Yang Liu, Yezhou Yang, Zunlei Feng, Yizhou Yu, Dacheng Tao, and Mingli Song. Stroke controllable fast style transfer with adaptive receptive fields. In *ECCV*, 2018. 2, 7
- [17] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016. 1, 2
- [18] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *CVPR*, 2016. 1, 2
- [19] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *CVPR*, 2017. 2
- [20] 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 super-resolution using a generative adversarial network. In *CVPR*, 2017. 1, 2, 5
- [21] Chuan Li and Michael Wand. Combining markov random fields and convolutional neural networks for image synthesis. In *CVPR*, 2016. 2
- [22] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Diversified texture synthesis with feed-forward networks. In *CVPR*, 2017. 3
- [23] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Universal style transfer via feature transforms. In *NIPS*, 2017. 2
- [24] Yixuan Li, Jason Yosinski, Jeff Clune, Hod Lipson, and John E Hopcroft. Convergent learning: Do different neural networks learn the same representations? In *ICLR*, 2016. 4
- [25] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *CVPRW*, 2017. 2
- [26] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *NIPS*, 2017. 1, 2, 6
- [27] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, 2015. 8
- [28] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *CVPR*, 2017. 1, 2
- [29] Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. *arXiv preprint arXiv:1709.07871*, 2017. 3
- [30] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*, 2016. 8
- [31] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2, 3
- [32] Jian Sun, Wenfei Cao, Zongben Xu, and Jean Ponce. Learning a convolutional neural network for non-uniform motion blur removal. In *CVPR*, 2015. 2
- [33] Ying Tai, Jian Yang, and Xiaoming Liu. Image super-resolution via deep recursive residual network. In *CVPR*, 2017. 2
- [34] Radu Timofte, Eirikur Agustsson, Luc Van Gool, Ming-Hsuan Yang, Lei Zhang, Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, Kyoung Mu Lee, et al. Ntire 2017 challenge on single image super-resolution: Methods and results. In *CVPRW*, 2017. 1, 5

- [35] Dmitry Ulyanov, Vadim Lebedev, Andrea Vedaldi, and Victor S Lempitsky. Texture networks: Feed-forward synthesis of textures and stylized images. In *ICMR*, 2016. 2
- [36] Paul Upchurch, Jacob R Gardner, Geoff Pleiss, Robert Pless, Noah Snaveley, Kavita Bala, and Kilian Q Weinberger. Deep feature interpolation for image content changes. In *CVPR*, 2017. 3
- [37] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *CVPR*, 2018. 1, 2, 6
- [38] Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Recovering realistic texture in image super-resolution by deep spatial feature transform. In *CVPR*, 2018. 3
- [39] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. ESRGAN: Enhanced super-resolution generative adversarial networks. In *ECCVW*, 2018. 3, 5
- [40] Zhangyang Wang, Ding Liu, Shiyu Chang, Qing Ling, Yingzhen Yang, and Thomas S Huang. D3: Deep dual-domain based fast restoration of jpeg-compressed images. In *CVPR*, 2016. 2
- [41] Kilian Q Weinberger and Lawrence K Saul. Unsupervised learning of image manifolds by semidefinite programming. *IJCV*, 70(1):77–90, 2006. 3
- [42] Li Xu, Jimmy SJ Ren, Ce Liu, and Jiaya Jia. Deep convolutional neural network for image deconvolution. In *NIPS*, 2014. 2
- [43] Ke Yu, Chao Dong, Liang Lin, and Chen Change Loy. Crafting a toolchain for image restoration by deep reinforcement learning. In *CVPR*, 2018. 2
- [44] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. *TIP*, 26(7):3142–3155, 2017. 1, 2, 5
- [45] Kai Zhang, Wangmeng Zuo, Shuhang Gu, and Lei Zhang. Learning deep cnn denoiser prior for image restoration. In *CVPR*, 2017. 2
- [46] Kai Zhang, Wangmeng Zuo, and Lei Zhang. Learning a single convolutional super-resolution network for multiple degradations. In *CVPR*, 2018. 2, 3
- [47] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 1, 2, 6